

IBM System Storage SAN Volume Controller



Command-Line Interface User's Guide

Version 5.1.0

IBM System Storage SAN Volume Controller



Command-Line Interface User's Guide

Version 5.1.0

Note:

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

This edition applies to the IBM System Storage SAN Volume Controller, release 5.1.0, and to all subsequent releases and modifications until otherwise indicated in new editions. This edition replaces SC26-7903-05.

© **Copyright International Business Machines Corporation 2003, 2009.**

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

Contents

Tables xxv

About this guide xxvii

Who should use this guide xxvii

Summary of changes xxvii

Summary of changes for SC26-7903-05 and
SC26-7903-06 SAN Volume Controller

Command-Line Interface User's Guide . . . xxvii

Summary of changes for SC26-7903-04 SAN

Volume Controller Command-Line Interface
User's Guide xxxi

Emphasis xxxii

SAN Volume Controller library and related
publications xxxii

How to order IBM publications. xxxvi

How to send your comments xxxvi

Syntax diagrams xxxvii

Terminology xxxviii

CLI special characters xxxix

Using wildcards in the SAN Volume
Controller CLI xxxix

Data types and value ranges xl

CLI commands and parameters xlv

CLI flags. xlvi

CLI messages xlvii

Chapter 1. Preparing the SSH client system for the CLI 1

Preparing the SSH client system to issue CLI
commands 2

Chapter 2. Secure Shell 3

Configuring a Secure Shell client system other than
PuTTY 3

Generating an SSH key pair using PuTTY 4

Configuring a PuTTY session for the CLI. . . . 5

Creating users 6

Adding SSH keys for hosts other than the IBM
System Storage Productivity Center or the master
console 7

Chapter 3. Copying the SAN Volume Controller software upgrade files using PuTTY scp 9

Chapter 4. Audit log commands 11

catauditlog 11

dumpauiditlog 13

lsauiditlogdumps. 14

Chapter 5. User management commands 17

chauthservice. 17

chcurrentuser. 20

chuser 20

chusergrp 22

mkuser 22

mkusergrp. 24

rmuser 25

rmusergrp. 25

Chapter 6. Role-based security commands (Discontinued) 27

mkauth (Discontinued) 27

rmauth (Discontinued). 27

lsaauth (Discontinued) 27

Chapter 7. E-mail and event notification commands 29

chemail. 30

chemailserver. 31

chemailuser 32

chsnmpserver 34

chsyslogserver 35

mkemailserver 36

mkemailuser 37

mksnmpserver 38

mksyslogserver 39

rmemailserver 40

rmemailuser 41

rmsnmpserver 42

rmsyslogserver 42

sendinventoryemail. 43

setemail (Discontinued) 43

startemail 43

stopemail 44

testemail 44

Chapter 8. Cluster commands. 47

addnode 48

cfgportip 50

chcluster 51

chclusterip. 55

chiogrp 57

chnode 59

cleardumps 60

cpdumps 61

detectmdisk 63

ping. 64

rmnode. 65

rmportip 67

setclustertime. 68

setpwdreset 69

settimezone 70

startstats 70

stopcluster. 72

stopstats 74

Chapter 9. Backup and restore commands	75	Chapter 15. FlashCopy commands	141
backup	75	chfconsistgrp	141
clear	76	chfcmmap	141
help	76	mkfconsistgrp	143
restore	77	mkfcmmap	144
Chapter 10. Cluster diagnostic and service-aid commands	79	prestartfconsistgrp	146
applysoftware	79	prestartfcmmap	148
cherrstate	80	rmfconsistgrp	149
clearerrlog	81	rmfcmmap	149
dumperrlog	82	startfconsistgrp	150
finderr	82	startfcmmap	151
setevent (Discontinued)	83	stopfconsistgrp	152
setlocale	83	stopfcmmap	153
svqueryclock	84	Chapter 16. Metro Mirror and Global Mirror commands	155
writesernum	84	chpartnership	155
Chapter 11. Host commands	87	chrconsistgrp	156
addhostiogr	87	chrrelationship	156
addhostport	87	mkpartnership	158
chhost	89	mkrconsistgrp	159
mkhost	90	mkrrelationship	160
rmhost	92	rmpartnership	163
rmhostiogr	92	rmrconsistgrp	163
rmhostport	93	rmrrelationship	164
Chapter 12. Virtual disk commands	97	startrcconsistgrp	165
addvdiskcopy	97	startrcrelationship	167
chvdisk	100	stoprcconsistgrp	169
expandvdisksize	104	stoprcrelationship	171
mkvdisk	106	switchrcconsistgrp	172
mkvdiskhostmap	112	switchrcrelationship	173
recovervdisk	114	Chapter 17. Migration commands	175
recovervdiskbycluster	115	migrateexts	175
recovervdiskbyiogr	115	migratetoimage	176
repairsevdiskcopy	116	migratevdisk	177
repairvdiskcopy	117	Chapter 18. Tracing commands	179
rmvdisk	118	setdisktrace	179
rmvdiskcopy	120	settrace	180
rmvdiskhostmap	121	starttrace	182
shrinkvdisksize	121	stoptrace	182
splitvdiskcopy	124	Chapter 19. Attributes of the -filtervalue parameters	185
Chapter 13. Managed disk group commands	127	Chapter 20. Overview of the list dumps commands	193
addmdisk	127	Chapter 21. Information commands	195
chmdiskgrp	128	caterrlog	195
mkmdiskgrp	129	caterrlogbyseqnum	197
rmdisk	130	ls2145dumps	198
rmdiskgrp	132	lscimomdumps	199
Chapter 14. Managed disk commands	135	lscopystatus	200
applymdiskssoftware	135	lslcluster	200
chmdisk	136	lslclustercandidate	204
includemdisk	136	lslclusterip	205
setquorum	137	lscontroller	207
triggermdiskdump	138		

lscontrollerdependentvdisks	210
lscurrentuser	210
lsdiscoverystatus	211
lsemailserver	212
lsemaileruser	213
lserrlogbyfcconsistgrp	214
lserrlogbyfcmap	215
lserrlogbyhost	217
lserrlogbyiogrp	218
lserrlogbymdisk	219
lserrlogbymdiskgrp	220
lserrlogbynode	222
lserrlogbyrcconsistgrp	223
lserrlogbyrcrelationship	224
lserrlogbyvdisk	225
lserrlogdumps	227
lsfabric	228
lsfcconsistgrp	230
lsfcmap	232
lsfcmapcandidate	234
lsfcmapprogress	235
lsfcmapdependentmaps	236
lsfeaturedumps	237
lsfreeextents	238
lshbaportcandidate	239
lshost	240
lshostiogrp	242
lshostvdiskmap	243
lsiogrp	245
lsiogrpghost	247
lsiogrpcandidate	247
lsiostatsdumps	248
lsiotracedumps	249
lsiscsiauth	250
lslicense	252
lsmdisk	253
lsmdiskdumps	257
lsmdisklba	257
lsmdiskcandidate	259
lsmdiskextent	260
lsmdiskgrp	262
lsmdiskmember	265
lsmigrate	267
lsnode	267
lsnodecandidate	270
lsnodependentvdisks	271
lsnodevpd	273
lsportip	276
lsquorum	279
lsrconsistgrp	280
lsrrelationship	283
lsrrelationshipcandidate	286
lsrrelationshipprogress	287
lsrepairsevdiskcopyprogress	288
lsrepairvdiskcopyprogress	289
lsrmvdiskdependentmaps	290
lsroute	291
lssevdiskcopy	292
lssnmpserver	294
lssoftwaredumps	295
lssoftwareupgradestatus	296

lssshkeys (Discontinued)	296
lssyslogserver	296
lstimezones	298
lsuser	299
lsusergrp	300
lsvdisk	301
lsvdiskcopy	306
lsvdiskdependentmaps	308
lsvdiskextent	308
lsvdiskfcmapcopies	310
lsvdiskfcmapmappings	311
lsvdiskhostmap	312
lsvdisklba	313
lsvdiskmember	314
lsvdiskprogress	316
lsvdisksyncprogress	317
showtimezone	318

Chapter 22. Error log commands

(Discontinued)	321
setevent (Discontinued)	321

Chapter 23. Licensing commands . . . 323

chlicense	323
dumpinternallog	325

Chapter 24. Secure Shell key

commands (Discontinued). 327

addsshkey (Discontinued)	327
rmallsshkeys (Discontinued)	327
rmsshkey (Discontinued)	327

Chapter 25. Service mode commands 329

applysoftware	329
cleardumps	330
dumperrlog	331
exit	332

Chapter 26. Service mode information

commands 333

ls2145dumps	333
lscimomdumps	334
lsclustervpd	335
lserrlogdumps	336
lsfeaturedumps	337
lsiostatsdumps	338
lsiotracedumps	339
lsmdiskdumps	340
lsnodevpd	341
lssoftwaredumps	344

Chapter 27. Controller command . . . 347

chcontroller	347
------------------------	-----

Chapter 28. Command-line interface

messages 349

CMMVC5000I No message was found for major rc <i>MAJOR_RC</i> , minor rc <i>MINOR_RC</i> , for action/view id <i>ACTION_VIEW_ID</i>	349	CMMVC5737E The parameter <i>PARAMETER</i> has been entered multiple times. Enter the parameter only one time..	357
CMMVC5700E The parameter list is not valid.	349	CMMVC5738E The argument <i>ARGUMENT</i> contains too many characters.	357
CMMVC5701E No object ID was specified.	349	CMMVC5739E The argument <i>ARGUMENT</i> does not contain enough characters.	357
CMMVC5702E <i>VALUE</i> is below the minimum level..	350	CMMVC5740E The filter flag <i>VALUE</i> is not valid..	357
CMMVC5703E The value or list starting with <i>VALUE</i> is above the maximum permitted for that value or has exceeded the number of items allowed in a list.	350	CMMVC5741E The filter value <i>VALUE</i> is not valid..	358
CMMVC5704E <i>VALUE</i> is not divisible by the permitted step value.	350	CMMVC5742E A specified parameter is out of its valid range..	358
CMMVC5705E A required parameter is missing.	350	CMMVC5743E A specified parameter does not comply with the step value..	358
CMMVC5706E An invalid argument has been entered for the <i>PARAMETER</i> parameter.	351	CMMVC5744E Too many objects were specified in the command.	358
CMMVC5707E Required parameters are missing.	351	CMMVC5745E Too few objects were specified in the request.	359
CMMVC5708E The <i>PARAMETER</i> parameter is missing its associated arguments.	351	CMMVC5746E The requested operation cannot be applied to the object specified..	359
CMMVC5709E <i>VALUE</i> is not a supported parameter.	351	CMMVC5747E The action requested is invalid - internal error.	359
CMMVC5711E <i>VALUE</i> is not valid data..	351	CMMVC5748E The action requested is invalid - internal error.	359
CMMVC5712E Required data is missing.	352	CMMVC5749E The dump filename specified already exists.	359
CMMVC5713E Some parameters are mutually exclusive.	352	CMMVC5750E The dump file could not be created - the file system is probably full.	360
CMMVC5714E The parameter list is empty.	352	CMMVC5751E The dump file could not be written to..	360
CMMVC5715E The parameter list does not exist..	352	CMMVC5752E Request failed. The object contains child objects, these must be deleted first..	360
CMMVC5716E Non-numeric data was entered for the numeric field <i>FIELD</i> . Enter a numeric value..	353	CMMVC5753E The specified object does not exist or is not a suitable candidate.	360
CMMVC5717E No match was found for the specified unit.	353	CMMVC5754E The specified object does not exist, or the name supplied does not meet the naming rules.	360
CMMVC5718E An unexpected return code was received.	353	CMMVC5755E Cannot create as the sizes of the specified objects do not match..	361
CMMVC5719E A value of <i>VALUE</i> requires the parameter <i>PARAMETER</i> to be specified.	353	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.	361
CMMVC5721E <i>VALUE</i> is not a valid time stamp format. The valid format is MMDDHhmmYYYY.	354	CMMVC5757E Self Defining Structure (SDS) defaults not found - internal error.	361
CMMVC5722E <i>VALUE</i> is not a valid month.	354	CMMVC5758E Object name already exists.	361
CMMVC5723E <i>VALUE</i> is not a valid day.	354	CMMVC5759E An internal error has occurred - memory could not be allocated..	361
CMMVC5724E <i>VALUE</i> is not a valid hour..	354	CMMVC5760E Failed to add the node to the cluster member list.	362
CMMVC5725E <i>VALUE</i> is not a valid minute.	354	CMMVC5761E Failed to delete the node from the cluster member list..	362
CMMVC5726E <i>VALUE</i> are not valid seconds.	355	CMMVC5762E The request did not complete before the timeout period expired.	362
CMMVC5727E <i>VALUE</i> is not a valid filter.	355	CMMVC5763E The node failed to go online.	362
CMMVC5728E <i>VALUE</i> should be in the format minute:hour:day:month:weekday..	355	CMMVC5764E The mode change request is invalid - internal error.	362
CMMVC5729E One or more components in the list is not valid.	355		
CMMVC5730E <i>VALUE</i> is only valid when <i>VALUE</i> has a value of <i>VALUE</i>	355		
CMMVC5731E <i>VALUE</i> can only be entered when <i>VALUE</i> has been entered.	356		
CMMVC5732E The shared-memory interface is not available, return code <i>RETURN_CODE</i>	356		
CMMVC5733E Enter at least one parameter.	356		
CMMVC5734E A combination of values was entered that is not valid.	356		
CMMVC5735E The name entered is not valid. Enter an alphanumeric string that does not start with a number.	357		

CMMVC5765E The object specified is no longer a candidate - a change occurred during the request.	363	CMMVC5792E The action failed because the I/O group is used for recovery.	367
CMMVC5767E One or more of the parameters specified are invalid or a parameter is missing.	363	CMMVC5793E The node was not added to the cluster because the I/O group already contains a pair of nodes.	368
CMMVC5769E The requested operation requires all nodes to be online - one or more nodes are not online.	363	CMMVC5794E The action failed because the node is not a member of the cluster.	368
CMMVC5770E The ssh key file supplied is invalid.	363	CMMVC5795E The node was not deleted because a software upgrade is in progress.	368
CMMVC5771E The operation requested could not complete, usually due to child objects existing. To force the operation, specify the force flag.	363	CMMVC5796E The action failed because the I/O group that the node belongs to is unstable.	368
CMMVC5772E The operation requested could not be performed because software upgrade is in progress.	364	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.	369
CMMVC5773E The object selected is in the wrong mode to perform the requested operation.	364	CMMVC5798E The action failed because the node is offline.	369
CMMVC5774E The userid supplied is not valid.	364	CMMVC5799E The shut down was not successful because there is only one online node in the I/O group.	369
CMMVC5775E The directory attribute specified is not valid.	364	CMMVC5800E The action failed because an entity that was specified in the command does not exist.	369
CMMVC5776E The directory listing could not be retrieved.	364	CMMVC5801E The upgrade of the cluster software could not proceed because every node in the cluster must be online. Either delete the node that is offline or bring the node online and resubmit the command	370
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.	365	CMMVC5802E The upgrade of the cluster software could not proceed because there is an I/O group in the cluster that contains only one node. The software upgrade requires that each node in an I/O group be shut down and restarted. If there is only one node in an I/O group, I/O operations could be lost if I/O operations are not stopped before beginning the software upgrade.	370
CMMVC5778E Cannot create another cluster, a cluster already exists.	365	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.	370
CMMVC5780E The action could not be completed using the Remote Cluster name. Use the Remote Cluster Unique ID instead.	365	CMMVC5804E The action failed because an object that was specified in the command does not exist.	371
CMMVC5781E The cluster ID specified is invalid.	365	CMMVC5805E The progress information was not returned because the FlashCopy statistics are not ready yet.	371
CMMVC5782E The object specified is offline.	365	CMMVC5806E The action failed because an object that was specified in the command does not exist.	371
CMMVC5783E The information is not available to complete this command.	366	CMMVC5807E The action failed because the managed disk (MDisk) cannot be changed to the specified mode.	371
CMMVC5784E The cluster name specified is not unique, specify the cluster using the cluster ID.	366	CMMVC5808E The action failed because the managed disk (MDisk) does not exist.	372
CMMVC5785E The filename specified contains an illegal character.	366	CMMVC5809E The tracing of I/O operations was not started because it is already in progress.	372
CMMVC5786E The action failed because the cluster is not in a stable state.	366		
CMMVC5787E The cluster was not created because a cluster already exists.	366		
CMMVC5788E The service IP address is not valid.	367		
CMMVC5789E The cluster was not modified because the IP address, subnet mask, service address, SNMP address, or gateway address is not valid.	367		
CMMVC5790E The node was not added to the cluster because the maximum number of nodes has been reached.	367		
CMMVC5791E The action failed because an object that was specified in the command does not exist.	367		

CMMVC5810E The quorum index number for the managed disk (MDisk) was not set because the MDisk is not available, either because the MDisk is offline or because it is associated with another quorum disk.	372	CMMVC5828E The virtual disk (VDisk) was not created because the I/O group contains no nodes.. . . .	377
CMMVC5811E The quorum index number for the managed disk (MDisk) was not set because the quorum disk does not exist.. . . .	372	CMMVC5829E The image-mode or sequential-mode virtual disk (VDisk) was not created because more than one managed disk (MDisk) is specified.	377
CMMVC5812E The quorum index number for the managed disk (MDisk) was not set because the MDisk is in the wrong mode.. . . .	373	CMMVC5830E The image-mode virtual disk (VDisk) was not created because no managed disk (MDisk) was specified in the command. . .	377
CMMVC5813E The quorum index number for the managed disk (MDisk) was not set because the MDisk has a sector size that is not valid. . .	373	CMMVC5831E The virtual disk (VDisk) was not created because the preferred node for I/O operations is not part of the I/O group. . . .	377
CMMVC5814E The quorum index number for the managed disk (MDisk) was not set because quorum is not allowed on one or more associated controllers.	373	CMMVC5832E The property of the virtual disk (VDisk) was not modified because an entity that was specified in the command does not exist.. . . .	378
CMMVC5815E The managed disk (MDisk) group was not created because an entity that was specified in the command does not exist.. .	373	CMMVC5833E The property of the virtual disk (VDisk) was not modified because there are no nodes in the I/O group.	378
CMMVC5816E The action failed because an entity that was specified in the command does not exist.	374	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.. . . .	378
CMMVC5817E The specified managed disk (MDisk) group was invalid.	374	CMMVC5835E The virtual disk (VDisk) was not expanded because an entity that was specified in the command does not exist.. . . .	378
CMMVC5818E The managed disk (MDisk) group was not deleted because there is at least one MDisk in the group..	374	CMMVC5836E The virtual disk (VDisk) was not shrunk because it is locked.	379
CMMVC5819E The managed disk (MDisk) was not added to the MDisk group because the MDisk is part of another MDisk group.	374	CMMVC5837E The action failed because the virtual disk (VDisk) is part of a FlashCopy mapping.	379
CMMVC5820E The managed disk (MDisk) was not added to the MDisk group because an entity that was specified in the command does not exist..	375	CMMVC5838E The action failed because the virtual disk (VDisk) is part of a Remote Copy mapping.	379
CMMVC5821E The managed disk (MDisk) was not added to the MDisk group because not enough MDisks were included in the list.	375	CMMVC5839E The virtual disk (VDisk) was not shrunk because an object that was specified in the command does not exist..	379
CMMVC5822E The managed disk (MDisk) was not added to the MDisk group because too many MDisks were included in the list.	375	CMMVC5840E The virtual disk (VDisk) was not deleted because it is mapped to a host or because it is part of a FlashCopy or Remote Copy mapping, or is involved in an image mode migrate.	380
CMMVC5823E The managed disk (MDisk) was not deleted from the MDisk group because the MDisk is part of another MDisk group.	375	CMMVC5841E The virtual disk (VDisk) was not deleted because it does not exist.	380
CMMVC5824E The managed disk (MDisk) was not deleted from the MDisk group because it does not belong to the MDisk group..	376	CMMVC5842E The action failed because an object that was specified in the command does not exist.	380
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.	376	CMMVC5843E The virtual disk (VDisk)-to-host mapping was not created because the VDisk does not have a capacity greater than zero bytes. .	380
CMMVC5826E The virtual disk (VDisk) was not created because an entity that was specified in the command does not exist..	376	CMMVC5844E The virtual disk (VDisk)-to-host mapping was not created because the SCSI logical unit number (LUN) ID is not valid. . .	381
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. .	376	CMMVC5845E The extent was not migrated because an object that was specified in the command does not exist.	381
		CMMVC5846E The virtual disk (VDisk) was not migrated because an object that was specified in the command does not exist..	381

CMMVC5847E The virtual disk (VDisk) was not migrated because its associated managed disk (MDisk) is already in the MDisk group.. . . .	381	CMMVC5867E The action failed because the worldwide port name is already assigned or is not valid.	386
CMMVC5848E The action failed because the virtual disk (VDisk) does not exist or it is being deleted..	382	CMMVC5868E The action failed because an entity that was specified in the command does not exist.	386
CMMVC5849E The migration failed because some or all of the extents are already being migrated.	382	CMMVC5869E The host object was not renamed because the host ID or name is not valid.	387
CMMVC5850E The extent was not migrated because there is a problem with the source extents..	382	CMMVC5870E The host object was not deleted because an entity that was specified in the command does not exist.	387
CMMVC5851E The extent was not migrated because there is a problem with the target extents..	382	CMMVC5871E The action failed because one or more of the configured port names is in a mapping..	387
CMMVC5852E The migration failed because there are too many migrations in progress.	382	CMMVC5872E The port (WWPN) was not added to the host object because an object that was specified in the command does not exist..	387
CMMVC5853E The action failed because there was a problem with the group.	383	CMMVC5873E No matching WWPN.	387
CMMVC5854E The extent information was not returned because the extent is not used or does not exist.	383	CMMVC5874E The action failed because the host does not exist.	388
CMMVC5855E The extent information was not returned because the managed disk (MDisk) is not used by any virtual disk (VDisk).	383	CMMVC5875E The action failed because the virtual disk (VDisk) does not exist.	388
CMMVC5856E The action failed because the virtual disk (VDisk) does not belong to the specified managed disk (MDisk) group.	383	CMMVC5876E The virtual disk (VDisk)-to-host mapping was not created because the maximum number of mappings has been reached.	388
CMMVC5857E The action failed because the managed disk (MDisk) does not exist or it is not a member of the managed disk (MDisk) group..	384	CMMVC5877E The virtual disk (VDisk)-to-host mapping was not created because the maximum number of SCSI LUNs has been allocated.	388
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..	384	CMMVC5878E The virtual disk (VDisk)-to-host mapping was not created because this VDisk is already mapped to this host.	389
CMMVC5859E The migration did not complete because an error occurred during the migration of the last extent on an image-mode virtual disk (VDisk).	384	CMMVC5879E The virtual disk (VDisk)-to-host mapping was not created because a VDisk is already mapped to this host with this SCSI LUN.	389
CMMVC5860E The action failed because there were not enough extents in the managed disk (MDisk) group..	384	CMMVC5880E The virtual disk was not created because a capacity of zero bytes is not allowed for image mode disks.	389
CMMVC5861E The action failed because there were not enough extents on the managed disk (MDisk)..	385	CMMVC5881E The FlashCopy mapping was not created because an entity that was specified in the command does not exist..	389
CMMVC5862E The action failed because the virtual disk (VDisk) is being formatted..	385	CMMVC5882E The FlashCopy mapping was not created because a mapping for the source or target virtual disk (VDisk) already exists..	390
CMMVC5863E The migration failed because there are not enough free extents on the target managed disk (MDisk).	385	CMMVC5883E The FlashCopy mapping was not created because the recovery I/O group is associated with the source or target virtual disk (VDisk).	390
CMMVC5864E The extent information was not returned because the source extent is not used.	385	CMMVC5884E The FlashCopy mapping was not created because the source or target virtual disk (VDisk) cannot be a member of a Remote Copy mapping.	390
CMMVC5865E The action failed because the extent is out of range for the managed disk (MDisk) or virtual disk (VDisk) specified.	386	CMMVC5885E The FlashCopy mapping was not created because this source or target virtual disk (VDisk) cannot be a member of a FlashCopy mapping.	390
CMMVC5866E The action failed because the extent contains internal data..	386	CMMVC5886E The FlashCopy mapping was not created because the source or target virtual disk (VDisk) is associated with the recovery I/O group.	391

CMMVC5887E The FlashCopy mapping was not created because the source or target virtual disk (Vdisk) must not be in router mode.	391	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.	395
CMMVC5888E The action failed because an entity that was specified in the command does not exist.	391	CMMVC5906E The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the preparing state.	396
CMMVC5889E The FlashCopy mapping was not deleted because an entity that was specified in the command does not exist.	391	CMMVC5907E The FlashCopy mapping or consistency group was not started because the mapping or consistency group is already in the copying state.	396
CMMVC5890E The FlashCopy mapping or consistency group was not started because starting consistency group 0 is not a valid operation.	392	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.	396
CMMVC5891E The FlashCopy consistency group was not created because the name is not valid.	392	CMMVC5909E The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the suspended state.	396
CMMVC5892E The FlashCopy consistency group was not created because it already exists.	392	CMMVC5910E The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is in the idle state.	397
CMMVC5893E The action failed because an entity that was specified in the command does not exist.	392	CMMVC5911E The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is in the preparing state.	397
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.	393	CMMVC5912E The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is already in the stopped state.	397
CMMVC5895E The FlashCopy consistency group was not deleted because it contains mappings. To delete this consistency group, a forced deletion is required.	393	CMMVC5913E The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the preparing state.	397
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.	393	CMMVC5914E The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the prepared state.	398
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.	393	CMMVC5915E The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the copying state.	398
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.	394	CMMVC5916E The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the suspended state.	398
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.	394	CMMVC5917E The FlashCopy mapping was not created because there is no memory in which to create the bitmap.	398
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.	394	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.	399
CMMVC5901E The FlashCopy mapping was not prepared because the mapping or consistency group is already in the preparing state.	394		
CMMVC5902E The FlashCopy mapping was not prepared because the mapping or consistency group is already in the prepared state.	395		
CMMVC5903E The FlashCopy mapping was not prepared because the mapping or consistency group is already in the copying state.	395		
CMMVC5904E The FlashCopy mapping was not prepared because the mapping or consistency group is already in the suspended state.	395		

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.	399	CMMVC5937E The action failed because an object that was specified in the command does not exist.	404
CMMVC5920E The FlashCopy mapping was not created because the consistency group is not idle.	400	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.	404
CMMVC5921E The properties of the FlashCopy mapping were not modified because the consistency group is not idle.	400	CMMVC5939E The action failed because the cluster is not in a stable state.	404
CMMVC5922E The FlashCopy mapping was not created because the destination virtual disk (VDisk) is too small.	400	CMMVC5940E The cluster that contains the auxiliary virtual disk (VDisk) is unknown.	405
CMMVC5923E The FlashCopy mapping cannot be created, either because there are no online nodes in the I/O group or because there are unrecovered FlashCopy mappings or unrecovered Global Mirror or Metro Mirror relationships in the I/O group.	400	CMMVC5941E The cluster that contains the master virtual disk (VDisk) has too many consistency groups.	405
CMMVC5924E The FlashCopy mapping was not created because the source and target virtual disks (VDisks) are different sizes.	401	CMMVC5942E The cluster that contains the auxiliary virtual disk (VDisk) has too many consistency groups.	405
CMMVC5925E The remote cluster partnership was not created because it already exists.	401	CMMVC5943E The specified relationship is not valid.	405
CMMVC5926E The remote cluster partnership was not created because there are too many partnerships.	401	CMMVC5944E The specified consistency group is not valid.	405
CMMVC5927E The action failed because the cluster ID is not valid.	401	CMMVC5945E The specified master cluster is not valid.	406
CMMVC5928E The action failed because the cluster name is a duplicate of another cluster.	402	CMMVC5946E The specified auxiliary cluster is not valid.	406
CMMVC5929E The Remote Copy partnership was not deleted because it has already been deleted.	402	CMMVC5947E The specified master virtual disk (VDisk) is not valid.	406
CMMVC5930E The Remote Copy relationship was not created because an object that was specified in the command does not exist.	402	CMMVC5948E The specified auxiliary virtual disk (VDisk) is not valid.	406
CMMVC5931E The Remote Copy relationship was not created because the master or auxiliary virtual disk (VDisk) is locked.	402	CMMVC5949E The specified relationship is unknown.	406
CMMVC5932E The Remote Copy relationship was not created because the master or auxiliary virtual disk (VDisk) is a member of a FlashCopy mapping.	403	CMMVC5950E The specified consistency group is unknown.	407
CMMVC5933E The Remote Copy relationship was not created because the master or auxiliary virtual disk (VDisk) is in the recovery I/O group.	403	CMMVC5951E The operation cannot be performed because the relationship is not a stand-alone relationship.	407
CMMVC5934E The Remote Copy relationship was not created because the master or auxiliary virtual disk (VDisk) is in the router mode.	403	CMMVC5952E The relationship and consistency group have different master clusters.	407
CMMVC5935E The action failed because an object that was specified in the command does not exist.	403	CMMVC5953E The relationship and group have different auxiliary clusters.	407
CMMVC5936E The action failed because an object that was specified in the command does not exist.	404	CMMVC5954E The master and auxiliary virtual disks (VDisks) are different sizes.	407
		CMMVC5955E The maximum number of relationships has been reached.	408
		CMMVC5956E The maximum number of consistency groups has been reached.	408
		CMMVC5957E The master virtual disk (VDisk) is already in a relationship.	408
		CMMVC5958E The auxiliary virtual disk (VDisk) is already in a relationship.	408
		CMMVC5959E There is a relationship that already has this name on the master cluster.	408
		CMMVC5960E There is a relationship that already has this name on the auxiliary cluster.	409
		CMMVC5961E There is a consistency group that already has this name on the master cluster.	409

CMMVC5962E There is a consistency group that already has this name on the auxiliary cluster.	409	CMMVC5984E The dump file was not written to disk. This may be due to the file system being full.	414
CMMVC5963E No direction has been defined.	409	CMMVC5985E The action failed because the directory that was specified was not one of the following directories: /dumps, /dumps/iostats, /dumps/iotrace, /dumps/feature, /dumps/configs, /dumps/elogs, or /home/admin/upgrade.	414
CMMVC5964E The copy priority is not valid.	409	CMMVC5986E The tracing of I/O operations was not started because the virtual disk (VDisk) or managed disk (MDisk) failed to return any statistics.	415
CMMVC5965E The virtual disks (VDisks) are in different I/O groups on the local cluster.	410	CMMVC5987E <i>VALUE</i> is not a valid command line option.	415
CMMVC5966E The master virtual disk (VDisk) is unknown.	410	CMMVC5988E command should not be run by the root userid. Use the admin userid.	415
CMMVC5967E The auxiliary virtual disk (VDisk) is unknown.	410	CMMVC5989E The operation was not performed because the relationship is offline.	415
CMMVC5968E The relationship cannot be added because the states of the relationship and the consistency group do not match.	410	CMMVC5990E The FlashCopy consistency group was not stopped as there are no FlashCopy mappings within the group.	416
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.	410	CMMVC5991E The Remote Copy consistency group was not stopped as there are no Remote Copy relationships within the group.	416
CMMVC5970E The Remote Copy relationship was not created because there is not enough memory.	411	CMMVC5992E The Remote Copy consistency group was not stopped as there are no Remote Copy relationships within the group.	416
CMMVC5971E The operation was not performed because the consistency group contains no relationships.	411	CMMVC5993E The specified upgrade package does not exist.	416
CMMVC5972E The operation was not performed because the consistency group contains relationships.	411	CMMVC5994E Error in verifying the signature of the upgrade package.	416
CMMVC5973E The operation was not performed because the consistency group is not synchronized.	412	CMMVC5995E Error in unpacking the upgrade package.	417
CMMVC5974E The operation was not performed because the consistency group is offline.	412	CMMVC5996E The specific upgrade package cannot be installed over the current version.	417
CMMVC5975E The operation was not performed because the cluster partnership is not connected.	412	CMMVC5999W Featurization for this facility has not been enabled.	417
CMMVC5976E The operation was not performed because the consistency group is in the freezing state.	412	CMMVC6000W Featurization for this facility has not been enabled.	417
CMMVC5977E The operation was not performed because it is not valid given the current consistency group state.	412	CMMVC6001E The FlashCopy consistency group was not started as there are no FlashCopy mappings within the group.	418
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.	413	CMMVC6002E This command can only be run on a node that is in service mode.	418
CMMVC5980E The operation was not performed because the master and auxiliary clusters are not connected.	413	CMMVC6003E This command can not be run on a node that is in service mode.	418
CMMVC5981E The operation was not performed because the relationship is in the freezing state.	413	CMMVC6004E The delimiter value <i>VALUE</i> is invalid.	418
CMMVC5982E The operation was not performed because it is not valid given the current relationship state.	414	CMMVC6005E The view request failed as the specified object is not a member of an appropriate group.	418
CMMVC5983E dump file was not created. This may be due to the file system being full.	414	CMMVC6006E The managed disk (MDisk) was not deleted because the resource was busy.	419
		CMMVC6007E The two passwords that were entered do not match.	419
		CMMVC6008E The key already exists.	419
		CMMVC6009E Unable to malloc a block of memory in which to copy the returned data.	419

CMMVC6010E Unable to complete the command as there are insufficient free extents, or the command requested an expansion of 0 size.	419	CMMVC6032E The operation was not performed because one or more of the entered parameters is invalid for this operation.. . . .	425
CMMVC6011E This cluster is part of a remote cluster partnership. Because this upgrade package will make changes to the cluster state, it cannot be applied to the current code level until all remote cluster partnerships are deleted.	420	CMMVC6033E The action failed due to an internal error.	425
CMMVC6012W The virtualized storage capacity is approaching the amount that you are licensed to use.	420	CMMVC6034E The action failed because the maximum number of objects has been reached.. . . .	425
CMMVC6013E The command failed because there is a consistency group mismatch on the aux cluster.	420	CMMVC6035E The action failed as the object already exists.	425
CMMVC6014E The command failed because the requested object is either unavailable or does not exist.	420	CMMVC6036E An invalid action was requested.	425
CMMVC6015E A delete request is already in progress for this object.	421	CMMVC6037E The action failed as the object is not empty.	426
CMMVC6016E The action failed as there would be, or are, no more disks in the MDisk group.	421	CMMVC6038E The action failed as the object is empty.	426
CMMVC6017E A parameter or argument contains invalid characters. Ensure that all characters are ASCII.	421	CMMVC6039E The action failed as the object is not a member of a group.. . . .	426
CMMVC6018E The software upgrade pre-install process failed.	421	CMMVC6040E The action failed as the object is not a parent..	426
CMMVC6019E The software upgrade failed as a node pended as the upgrade was in progress.. . . .	422	CMMVC6041E The action failed as the cluster is full..	426
CMMVC6020E The software upgrade failed as the system was unable to distribute the software package to all nodes..	422	CMMVC6042E The action failed as the object is not a cluster member..	427
CMMVC6021E The system is currently busy performing another request. Try again later.	422	CMMVC6043E The action failed as the object is a member of a group.	427
CMMVC6022E The system is currently busy performing another request. Try again later.	422	CMMVC6044E The action failed as the object is a parent.	427
CMMVC6023E The system is currently busy performing another request. Try again later.	422	CMMVC6045E The action failed, as the -force flag was not entered.	427
CMMVC6024E The auxiliary VDisk entered is invalid.	423	CMMVC6046E The action failed as too many candidates were selected.	427
CMMVC6025E The RC consistency group Master cluster is not the local cluster.	423	CMMVC6047E The action failed as too few candidates were selected.	428
CMMVC6026E The RC consistency group is not in the stopped state.	423	CMMVC6048E The action failed as the object is busy..	428
CMMVC6027E The RC consistency group is not the primary master..	423	CMMVC6049E The action failed as the object is not ready.	428
CMMVC6028E This upgrade package cannot be applied to the current software level because it contains changes to the cluster state and there are remote cluster partnership defined.	424	CMMVC6050E The action failed as the command was busy.	428
CMMVC6029E All nodes must have identical code level before a concurrent code upgrade can be performed.	424	CMMVC6051E An unsupported action was selected.	428
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.	424	CMMVC6052E The action failed as the object is a member of a FlashCopy mapping..	429
CMMVC6031E The operation was not performed because the FlashCopy consistency group is empty..	424	CMMVC6053E An invalid WWPN was entered.	429
		CMMVC6054E The action failed as not all nodes are online.	429
		CMMVC6055E The action failed as an upgrade is in progress.	429
		CMMVC6056E The action failed as the object is too small.	429
		CMMVC6057E The action failed as the object is the target of a FlashCopy mapping.	430
		CMMVC6058E The action failed as the object is in the recovery HWS.	430
		CMMVC6059E The action failed as the object is in an invalid mode..	430
		CMMVC6060E The action failed as the object is being deleted.	430
		CMMVC6061E The action failed as the object is being resized.	430
		CMMVC6062E The action failed as the object is being moved between HWS..	431

CMMVC6063E The action failed as there are no more disks in the group.	431	CMMVC6087E Metadata recovery could not map the buffers necessary to complete the operation.	435
CMMVC6064E The action failed as the object has an invalid name.	431	CMMVC6088E The lba at which metadata recovery was requested does not contain metadata.	435
CMMVC6065E The action failed as the object is not in a group.	431	CMMVC6089E The metadata at the requested lba is flagged as invalid.	436
CMMVC6066E The action failed as the system is running low on memory.	431	CMMVC6090E The metadata header checksum verification failed.	436
CMMVC6067E The action failed as the SSH key was not found.	432	CMMVC6091E The metadata region checksum verification failed.	436
CMMVC6068E The action failed as there are no free SSH keys.	432	CMMVC6092E The metadata recovery operation was aborted.	436
CMMVC6069E The action failed as the SSH key is already registered.	432	CMMVC6093E Metadata recovery internal error - (read only)	436
CMMVC6070E An invalid or duplicated parameter, unaccompanied argument, or incorrect argument sequence has been detected. Ensure that the input is as per the help.	432	CMMVC6095E Metadata recovery encountered the end of the disk.	436
CMMVC6071E The virtual disk (VDisk)-to-host mapping was not created because the VDisk is already mapped to a host.	432	CMMVC6096E Metadata recovery encountered an error from a lower layer - (v) no resource).	436
CMMVC6073E The maximum number of files has been exceeded.	433	CMMVC6097E Metadata recovery encountered an error from a lower layer - (v) failure).	436
CMMVC6074E The command failed as the extent has already been assigned.	433	CMMVC6098E The copy failed as the specified node is the configuration node.	437
CMMVC6075E The expand failed as the last extent is not a complete extent.	433	CMMVC6100E <i>OPTION</i> not consistent with <i>ACTION</i>	437
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.	433	CMMVC6101E <i>OPTION</i> not consistent with <i>OPTION</i>	437
CMMVC6077E WARNING - Unfixed errors should be fixed before applying software upgrade. Depending on the nature of the errors, they might cause the upgrade process to fail. It is highly recommended to fix these errors before proceeding. If a particular error cannot be fixed, contact the support center.	434	CMMVC6102E <i>OPTION</i> and <i>OPTION</i> are alternatives	437
CMMVC6078E The action failed because the object is in an invalid mode.	434	CMMVC6103E Problem with <i>FILENAME: DETAILS</i>	437
CMMVC6079E Metadata recovery could not complete the operation because a parameter is invalid.	434	CMMVC6104E Action <i>ACTION</i> not run.	438
CMMVC6081E Metadata Recovery is busy processing the previous operation.	434	CMMVC6105E Different names for source <i>SOURCE_CLUSTER_NAME</i> and target <i>TARGET_CLUSTER_NAME</i> clusters	438
CMMVC6082E The attempt to abort metadata recovery failed because the previous operation has completed.	434	CMMVC6106W Target cluster has non-default id alias <i>ALIAS</i>	438
CMMVC6083E Metadata recovery could not find a valid dumpfile required for the rebuild operation.	435	CMMVC6107E <i>NUMBER_OF_OBJECTS</i> io_grp objects in target cluster; <i>NUMBER_OF_REQUIRED_OBJECTS</i> are required	438
CMMVC6084E Metadata recovery could not create/open/write the scan file, the disk might be full.	435	CMMVC6108I Disk controller system with a <i>WNN</i> of <i>WNN_VALUE</i> found.	438
CMMVC6085E Metadata recovery could not create/open/write the dump file, the disk might be full.	435	CMMVC6109E Disk controller system with a <i>WNN</i> of <i>WNN_VALUE</i> not available.	439
CMMVC6086E Metadata recovery could not create/open/write the progress file, the disk might be full.	435	CMMVC6110E Bad code level: <i>VALUE</i>	439
		CMMVC6111E The cluster code level could not be determined from <i>VALUE</i>	439
		CMMVC6112W <i>OBJECT_TYPE OBJECT_NAME</i> has a default name.	439
		CMMVC6113E The command <i>COMMAND</i> has failed with return code <i>RETURN_CODE</i>	440
		CMMVC6114E No help for action <i>ACTION</i>	440
		CMMVC6115W Feature <i>FEATURE_PROPERTY</i> mismatch: <i>VALUE</i> expected; <i>VALUE</i> found.	440
		CMMVC6116I Feature match for <i>FEATURE</i>	440
		CMMVC6117E <i>FIX_OR_FEATURE</i> is not available.	440
		CMMVC6118I <i>TYPE</i> with <i>PROPERTY PROPERTY_VALUE</i> and <i>PROPERTY PROPERTY_VALUE</i> found.	441
		CMMVC6119E <i>TYPE</i> with <i>PROPERTY PROPERTY_VALUE</i> not found.	441

CMMVC6120E Target is not the configuration node	441	CMMVC6154E Required <i>OBJECT_TYPE</i> property <i>PROPERTY_NAME</i> has a null value.	448
CMMVC6121E No cluster id or id_alias in backup configuration.	441	CMMVC6155I The command <i>COMMAND</i> processing has completed successfully.	448
CMMVC6122E No <i>TYPE</i> with <i>PROPERTY VALUE</i> is present in the table.	441	CMMVC6156W <i>COMMAND</i> processing completed with errors.	449
CMMVC6123E No <i>PROPERTY</i> for <i>TYPE NAME</i>	442	CMMVC6164E The SVCCONFIG CRON job, which runs overnight on a daily overnight, has failed.	449
CMMVC6124E No <i>TYPE</i> with <i>PROPERTY VALUE</i>	442	CMMVC6165E The target is not the original configuration node with a <i>WNN</i> of <i>WNN_VALUE</i>	449
CMMVC6125E No unique ID for <i>TYPE NAME</i>	442	CMMVC6166E The property <i>PROPERTY</i> of the object <i>OBJECT</i> has changed during svconfig restore -execute.	449
CMMVC6126E No <i>TYPE</i> with unique ID <i>VALUE</i>	442	CMMVC6181E The target cluster contains an object that has a counterpart in the configuration to be restored, and has the correct ID.	450
CMMVC6127I The SSH key <i>IDENTIFIER</i> for <i>USER</i> is already defined; the SSH key will not be restored.	442	CMMVC6182W An object that does not contribute to the fabric of the configuration cannot be restored because its configuration does not permit it to be created.	450
CMMVC6128W <i>DIRECTORY</i>	443	CMMVC6186W The IO group <i>IO_GROUP_NAME</i> has been restored with ID <i>ID_VALUE</i> instead of <i>ID_VALUE</i>	450
CMMVC6129E VDisk-to-host map objects have VDisk_UID values that are not consistent.	443	CMMVC6200E The action failed because of incompatible software.	450
CMMVC6130W Inter-cluster <i>PROPERTY VALUE</i> will not be restored.	443	CMMVC6201E The node could not be added because of incompatible software. The status code is <i>STATUS_CODE</i>	451
CMMVC6131E No location cluster information	443	CMMVC6202E The cluster was not modified because the IP address is not valid.	451
CMMVC6132E The object <i>OBJECT</i> of type <i>TYPE</i> has a property <i>PROPERTY</i> with an incorrect value <i>INCORRECT_VALUE</i> . The operation cannot proceed until the property has the correct value <i>CORRECT_VALUE</i> . Take administrative action to change the value and try again.	444	CMMVC6203E The action failed because the directory that was specified was not one of the following directories: /dumps, /dumps/iostats, /dumps/iotrace, /dumps/feature, /dumps/config, /dumps/elogs, /dumps/ec or /dumps/pl.	451
CMMVC6133E Required <i>TYPE</i> property <i>PROPERTY</i> not found	444	CMMVC6204E The action failed as the resulting disk size would be less than, or equal to, zero.	451
CMMVC6134E No argument for <i>OPTION</i>	444	CMMVC6205E Metadata recovery can not use the provided MDisk id - invalid or destroyed.	452
CMMVC6135E Argument <i>VALUE</i> for <i>OPTION</i> is not valid.	444	CMMVC6206E The software upgrade failed as a file containing the software for the specified MCP version was not found.	452
CMMVC6136W No SSH key file <i>FILENAME</i>	444	CMMVC6207E The action failed because the virtual disk (VDisk) is part of a Remote Copy mapping.	452
CMMVC6137W No SSH key file <i>FILENAME</i> ; key not restored	445	CMMVC6208E The action failed because the virtual disk (VDisk) is part of a FlashCopy mapping.	452
CMMVC6138E <i>OPTION</i> is required.	445	CMMVC6209E The FlashCopy mapping or consistency group could not be started in a reasonable time. The mapping or group is instead being prepared.	452
CMMVC6139E Incorrect XML tag nesting in <i>FILENAME</i>	445	CMMVC6210E The command has failed because a virtual medium error exists on the image mode VDisk or copy.	453
CMMVC6140E No default name for type <i>TYPE</i>	445	CMMVC6211E The command failed as a migrate to image was in progress.	453
CMMVC6141E The option <i>OPTION</i> does not support an argument.	446		
CMMVC6142E Existing <i>OBJECT_TYPE OBJECT_NAME</i> has a non-default name.	446		
CMMVC6143E The required configuration file <i>FILENAME</i> does not exist.	446		
CMMVC6144W The object with default name <i>NAME</i> has been restored as <i>SUBSTITUTE_NAME</i>	446		
CMMVC6145I First use the <i>COMMAND -prepare</i> command.	446		
CMMVC6146E Problem parsing <i>OBJECT_TYPE</i> data: <i>LINE</i>	447		
CMMVC6147E <i>TYPE NAME</i> has a name beginning with <i>PREFIX</i>	447		
CMMVC6148E Target cluster has <i>NUMBER_OF_EXISTING_OBJECTS</i> objects of type <i>TYPE</i> instead of <i>NUMBER_OF_REQUIRED_OBJECTS</i>	447		
CMMVC6149E An action is required.	447		
CMMVC6150E The action <i>ACTION</i> is not valid.	447		
CMMVC6151E The option <i>OPTION</i> is not valid.	448		
CMMVC6152E VDisk <i>VDISK_NAME</i> instance number <i>INSTANCE_NUMBER</i> is not valid.	448		
CMMVC6153E <i>OBJECT</i> not consistent with <i>ACTION</i>	448		

CMMVC6212E The command failed because data in the cache has not been committed to disk.	453	CMMVC6233E This operation cannot be performed because, either a software upgrade has not been started, or a software upgrade is in progress but is not in a state where it can be aborted.	458
CMMVC6213E You are trying to recover region data that was created by a code level different from the one you are currently running on the node.	453	CMMVC6234E The upgrade cannot be aborted because at least one node has already committed to a new code level.	458
CMMVC6214E Failed to recreate the cluster you are trying to rebuild.	454	CMMVC6235E An invalid response has been entered. The command has not been executed. Input is case sensitive. Enter either yes or no.	459
CMMVC6215E The FlashCopy mapping was not created or modified because the consistency group already contains the maximum number of mappings.	454	CMMVC6236E The command has not completed. A limited availability parameter has been entered without the required environment setting being set.	459
CMMVC6216E The Remote Copy relationship was not created because the master or auxiliary virtual disk (Vdisk) is a member of a Remote Copy mapping.	454	CMMVC6237E The command failed as the remote cluster does not support global mirror.	459
CMMVC6217E The maximum number of hosts for the cluster is already configured.	454	CMMVC6238E The copy type differs from other copies already in the consistency group.	459
CMMVC6218E The maximum number of host/I/O group pairs for the cluster is already configured.	455	CMMVC6239E The FlashCopy mapping was not prepared because the mapping or consistency group is in the stopping state. The mapping or consistency group must first complete the stop operation and then be prepared.	459
CMMVC6219E The maximum number of WWPNs for the cluster is already configured.	455	CMMVC6240E The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the stopping state.	460
CMMVC6220E The maximum number of hosts for one or more IO groups is already configured.	455	CMMVC6241E The FlashCopy mapping was not deleted because the mapping or consistency group is in the stopping state. The mapping or consistency group must be stopped first.	460
CMMVC6221E The maximum number of WWPNs for one or more IO groups is already configured.	456	CMMVC6242E The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the stopping state. The mapping or consistency group must first complete the stop operation and then be prepared.	460
CMMVC6222E The maximum number of WWPNs for the host is already configured.	456	CMMVC6243E The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is already in the stopping state.	461
CMMVC6223E The host does not belong to one or more of the IO groups specified or inferred.	456	CMMVC6244E The FlashCopy mapping was not created because the source virtual disk (Vdisk) cannot be the target for a FlashCopy mapping.	461
CMMVC6224E The host already belongs to one or more of the IO groups specified.	456	CMMVC6245E The FlashCopy mapping was not created because the source virtual disk (Vdisk) is already in the maximum number of FlashCopy mappings.	461
CMMVC6225E An IO group cannot be removed from a host because of one or more associated VDisks.	457	CMMVC6246E The FlashCopy mapping was not created because the target virtual disk (Vdisk) is already a source VDisk in a FlashCopy mapping.	462
CMMVC6226E The action was not completed because the cluster has reached the maximum number of extents in MDisk Groups.	457	CMMVC6247E The FlashCopy mapping was not created because the target virtual disk (Vdisk) is already a target VDisk in a FlashCopy mapping.	462
CMMVC6227I The package installed successfully.	457	CMMVC6248E The command failed because the authorization table is full.	462
CMMVC6228E The cluster was recovered and the CLI functionality is limited until the cause of the failure is determined and any corrective action taken. Contact IBM technical support for assistance.	457		
CMMVC6229E The action failed as the SSH key has been revoked.	458		
CMMVC6230E The action failed as the SSH key index (SSH_LABEL_ID) is invalid.	458		
CMMVC6231E The action failed as the audit table is full.	458		
CMMVC6232E This operation cannot be performed because the cluster is currently aborting the previous software upgrade command.	458		

CMMVC6249E The command failed because the authorization record was not found or is already set to the default role.	462	CMMVC6282E Sendmail error EX_NOPERM. The user does not have permission to perform the requested operation.	467
CMMVC6250E The command failed because the authorization record is not set to the default role. Use rmail to set the default role.	462	CMMVC6283E Sendmail error EX_CONFIG. There is a fatal problem with the sendmail configuration..	467
CMMVC6251E The command failed because the specified role was not found.	463	CMMVC6284E An unknown error occurred. Please ensure your SMTP server is running.	468
CMMVC6252E The command failed authorization because the session ssh key is invalid or was deleted.	463	CMMVC6285E The e-mail command timed out. Please check your e-mail server settings as listed on the SVC.	468
CMMVC6253E The task has failed because the user's role is not authorized to submit the command.	463	CMMVC6286E The e-mail service has not been enabled.	468
CMMVC6254E The command failed because the specified ssh key was not found. NOTE This command must specify an admin key.	463	CMMVC6287E The user specified does not exist..	468
CMMVC6255E The command cannot set the authorization record to the default role. Use rmail to set the default role.	463	CMMVC6288E The FlashCopy mapping or consistency group could not be started because a source VDisk is the target of another FC Map that is keeping the VDisk inaccessible.	468
CMMVC6263E The command failed because the ssh key already exists or there is a duplicate ssh key.	464	CMMVC6289E The command failed because the virtual disk (VDisk) is pending synchronization.	469
CMMVC6269E Sendmail error EX_USAGE. A command or configuration line has been used incorrectly.	464	CMMVC6290E The command failed because the virtual disk (VDisk) has image mode copies, is pending synchronization and -force has not been specified.	469
CMMVC6270E Sendmail error EX_DATAERR. Address is wrong, or the message is too large for the mailbox.	464	CMMVC6291E The command failed because the virtual disk (VDisk) is pending synchronization and -force has not been specified.	469
CMMVC6271E Sendmail error EX_NOINPUT. An input file (not a system file) did not exist or was not readable.	464	CMMVC6292E The command failed because a repair action is in progress for the virtual disk (VDisk).	470
CMMVC6272E Sendmail error EX_NOUSER. The sendmail command could not recognize a specified user.	465	CMMVC6296E One or more managed disks (MDisks) have failed validation tests. The first failing MDisk ID is <i>MDISK_ID</i>	470
CMMVC6273E Sendmail error EX_NOHOST. The sendmail command could not recognize the specified host name.	465	CMMVC6297E One or more managed disks (MDisks) have timed out during validation tests. The first failing MDisk ID is <i>MDISK_ID</i>	471
CMMVC6274E Sendmail error EX_UNAVAILABLE. A required system resource is not available.	465	CMMVC6298E The command failed because a target VDisk has dependent FlashCopy mappings..	471
CMMVC6275E Sendmail error EX_SOFTWARE. An internal software error occurred (including bad arguments).	465	CMMVC6299E The create failed because the source and target VDIsks are members of FlashCopy mappings that have different grain sizes..	471
CMMVC6276E Sendmail error EX_OSERR. A system resource error prevented the sending of an e-mail.	466	CMMVC6300E The create failed because the source and target VDIsks are members of FlashCopy mappings that belong to different I/O groups..	472
CMMVC6277E Sendmail error EX_OSFILE. Failed to open a critical system file.	466	CMMVC6301E The create failed because the specified consistency group does not exist..	472
CMMVC6278E Sendmail error EX_CANTCREAT. An output file could not be written to by sendmail.	466	CMMVC6302E The create failed because the resulting tree of FlashCopy mappings would exceed the upper limit.	472
CMMVC6279E Sendmail error EX_IOERR. A system I/O error occurred during a sendmail operation. This could be due to a disk failure.	466	CMMVC6303E The create failed because the source and target VDIsks are the same..	473
CMMVC6280E Sendmail error EX_TEMPFAIL. The sendmail command could not create a connection to a remote system.	467	CMMVC6304E The create failed because the source VDisk does not exist.	473
CMMVC6281E Sendmail error EX_PROTOCOL. The remote system returned something that was incorrect during a protocol exchange.	467	CMMVC6305E The create failed because the target VDisk does not exist.	473

CMMVC6306E The create failed because the source VDisk is the member of a FlashCopy mapping whose grain size is different to that specified.	473	CMMVC6332E The command has failed because an IPv6 e-mail server address was specified and the cluster does not have an IPv6 address..	481
CMMVC6307E The create failed because the target VDisk is the member of a FlashCopy mapping whose grain size is different to that specified.	474	CMMVC6333E The command has failed because an IPv4 e-mail server address was specified and the cluster does not have an IPv4 address..	482
CMMVC6308E The create failed because the source VDisk is the member of a FlashCopy mapping whose IO group is different to that specified.	474	CMMVC6334E The command failed as the e-mail port number supplied is invalid.	482
CMMVC6309E The create failed because the target VDisk is the member of a FlashCopy mapping whose IO group is different to that specified.	475	CMMVC6335E The command failed as the combination of parameters provided are either mutually incompatible or would leave the cluster without a functioning protocol stack. .	482
CMMVC6310E The modify failed because the specified FlashCopy mapping does not exist..	475	CMMVC6336E The virtual disk (VDisk) copy was not created because the grain size must be 32, 64, 128 or 256..	482
CMMVC6311E The command failed because the source VDisk is the target of a FlashCopy mapping that is in the specified consistency group..	475	CMMVC6337E The action failed because the warning size must be a multiple of 512 bytes. .	483
CMMVC6312E The command failed because the target VDisk is the source of a FlashCopy mapping that is in the specified consistency group..	476	CMMVC6338E The action failed because the warning size can not be larger than the virtual size.	483
CMMVC6313E The command failed because the specified background copy rate is invalid. .	476	CMMVC6339E The virtual disk (VDisk) copy was not created because the virtual size was not provided.	483
CMMVC6314E The command failed because the specified cleaning rate is not valid. . . .	476	CMMVC6340E The action failed because the value supplied for real size is not a multiple of 512 bytes.	483
CMMVC6315E The command failed because the specified grain size is not valid.	476	CMMVC6341E The action failed because the virtual disk (VDisk) copy is not space-efficient..	484
CMMVC6319E The command has failed because a combination of IPv4 and IPv6 parameters were entered.	477	CMMVC6342E The virtual disk (VDisk) copy was not shrunk because its real size cannot be less than its used size.	484
CMMVC6320E The command has failed because the IPv4 address is not valid..	477	CMMVC6343E The virtual disk (VDisk) copy was not shrunk because its real size can not be negative.	484
CMMVC6321E The command has failed because the IPv4 subnet mask is not valid.	477	CMMVC6344E The repair operation cannot start because the space-efficient virtual disk (VDisk) copy is already being repaired. . .	484
CMMVC6322E The command has failed because the IPv4 gateway address is not valid.. . .	477	CMMVC6345E The repair operation cannot start because the space-efficient virtual disk (VDisk) copy was created using -import but the cluster could not recognize its format. . . .	485
CMMVC6323E The command has failed because the IPv6 address is not valid..	478	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.	485
CMMVC6324E The command has failed because the IPv6 prefix is not valid.	478	CMMVC6347E The specific upgrade package cannot be installed on this hardware level..	485
CMMVC6325E The command has failed because the IPv6 gateway address is not valid.. . .	478	CMMVC6348E The command failed as there was not enough information provided to process successfully.	486
CMMVC6326E The command has failed because the IPv4 service mode address is not valid..	479	CMMVC6349E The command failed because the virtual disk cache has been lost. Use the force flag to acknowledge this.	486
CMMVC6327E The command has failed because the IPv6 service mode address is not valid..	479	CMMVC6350E The command failed because there is insufficient mirror bitmap space.	486
CMMVC6328E The command has failed because the console address is not valid.	480	CMMVC6351E The command failed because the virtual disk (VDisk) is not mirrored.	486
CMMVC6329E The command has failed because the IP address is not valid.	480	CMMVC6352E The command failed because the number of copies of this virtual disk (VDisk) would exceed the limit.	487
CMMVC6330E The command has failed because an IPv6 address was specified and the cluster does not have an IPv6 address..	481		
CMMVC6331E The command has failed because an IPv4 address was specified and the cluster does not have an IPv4 address..	481		

CMMVC6353E The command failed because the copy specified does not exist..	487	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..	494
CMMVC6354E The command failed because a copy is not synchronized.	487	CMMVC6402E The command failed because the managed disk (MDisk) is not in the required managed disk group..	494
CMMVC6355E The command failed because an image mode copy is not synchronized and -force was not specified.	487	CMMVC6403E The command failed because the target managed disk (MDisk) is not in the required managed disk group.	494
CMMVC6356E The command failed because a copy is not synchronized and -force was not specified.	488	CMMVC6404E The command failed because the source and target managed disk groups must be different.	494
CMMVC6357E The command failed because the copy specified is not synchronized and -force was not specified.	488	CMMVC6405E The command failed because the target copy was not specified..	495
CMMVC6358E The command failed because the copy specified is the only synchronized copy.	489	CMMVC6406E The command failed because the specified managed disk group does not exist.	495
CMMVC6359E The command failed because there are insufficient online synchronized copies.	489	CMMVC6407E The command failed because the managed disk group is invalid..	495
CMMVC6363E The command failed because the Logical Block Address (LBA) specified is invalid for this virtual disk (VDisk)..	489	CMMVC6408E The command failed because too few managed disk groups were specified.	495
CMMVC6364E The command failed because the logical block address (LBA) requested is too large for the disk..	489	CMMVC6409E The command failed because too many managed disk groups were specified..	496
CMMVC6365E The command timed out.	490	CMMVC6410E The command failed because too few managed disks (MDisks) were specified.	496
CMMVC6366E One or more nodes in the cluster has hardware that is not supported by the new software package.	490	CMMVC6411E The command failed because too many managed disks (MDisks) were specified..	496
CMMVC6367E A remote cluster is running software that is incompatible with the new software package.	490	CMMVC6412E The command failed because the managed disk group extent size is above maximum permitted size.	496
CMMVC6368E The new software package might be incompatible with the remote cluster..	491	CMMVC6413E The command failed because the managed disk (MDisk) is invalid..	497
CMMVC6369E The FlashCopy storage capacity that the cluster is using is approaching the FlashCopy storage capacity that is licensed.	491	CMMVC6414E The command failed because the managed disk (MDisk) is currently being migrated.	497
CMMVC6370E The Remote Copy storage capacity that the cluster is using is approaching the Remote Copy storage capacity that is licensed..	491	CMMVC6415E The command failed because the managed disk group warning threshold is too low.	497
CMMVC6372E The virtualized storage capacity that the cluster is using is approaching the virtualized storage capacity that is licensed..	491	CMMVC6416E The command failed because the managed disk group warning threshold is too high.	497
CMMVC6373E The virtualized storage capacity that the cluster is using exceeds the virtualized storage capacity that is licensed..	492	CMMVC6417E The command failed because the managed disk group warning threshold is invalid.	498
CMMVC6374E The FlashCopy storage capacity that the cluster is using exceeds the FlashCopy storage capacity that is licensed.	492	CMMVC6418E The command failed because the virtual disk (VDisk) is in the process of being resized..	498
CMMVC6375E The Remote Copy storage capacity that the cluster is using exceeds the Remote Copy storage capacity that is licensed.	492	CMMVC6419E The command failed because one or more of the specified managed disks (MDisks) are in the process of being deleted.	498
CMMVC6394E The command failed because an attempt to make the virtual disk cache empty took too long..	492	CMMVC6423E The Send Inventory e-mail operation failed because e-mail is not started.	498
CMMVC6399E The command failed because there is not enough memory available for reservation.	493	CMMVC6424E The Send Inventory e-mail operation failed because there are no inventory e-mail users.	499
CMMVC6400E The command failed because a specified managed disk (MDisk) is already in use.	493	CMMVC6425E The action failed because the maximum number of objects has been reached..	499
		CMMVC6426E The command failed because a specified managed disk (MDisk) is already in use.	499

CMMVC6427E The command failed because one or more of the specified managed disks (MDisks) are not in the required managed disk group.	499	CMMVC6449E The operation was not performed because the partnership owns Global or Metro Mirror relationships or consistency groups.	505
CMMVC6428E The command failed because the source managed disk (MDisk) is not in the required managed disk group.	500	CMMVC6450W A FlashCopy mapping was created but physical_flash is not enabled.	505
CMMVC6429E The command failed because the target managed disk (MDisk) is not in the required managed disk group.	500	CMMVC6451W A Global Mirror or Metro Mirror relationship was created but physical_remote is not enabled.	506
CMMVC6430E The command failed because the target and source managed disk groups must be different.	500	CMMVC6452W You are using the physical disk license scheme but the values for physical_flash and physical_remote are not set.	506
CMMVC6431E The command failed because the target copy was not specified.	500	CMMVC6453W You have disabled the physical disk license scheme but the capacity license scheme is not set.	506
CMMVC6432E The command failed because the specified managed disk group does not exist.	501	CMMVC6454E The command failed because the physical disk license scheme is not enabled.	506
CMMVC6433E The command failed because the managed disk group is invalid.	501	CMMVC6455E The command failed because a capacity license scheme parameter was specified but the physical disk license scheme is enabled.	507
CMMVC6434E The command failed because too few managed disk groups were specified.	501	CMMVC6456E The command failed because a physical disk license scheme parameter was specified but the capacity license scheme is enabled.	507
CMMVC6435E The command failed because too many managed disk groups were specified.	501	CMMVC6457E One or more quorum disks are on the specified controller.	507
CMMVC6436E The command failed because too few managed disks (MDisks) were specified.	502	CMMVC6458E The specified controller cannot support quorum disks.	508
CMMVC6437E The command failed because too many managed disks (MDisks) were specified.	502	CMMVC6459E The mkrrelationship command failed because the same VDisk was specified as the master and auxiliary VDisk.	508
CMMVC6438E The command failed because the managed disk group extent size is above maximum permitted size.	502	CMMVC6460E The command failed because the migration source is offline.	508
CMMVC6439E The command failed because the managed disk (MDisk) is invalid.	502	CMMVC6461E The command failed because starting the migration will result in VDIsks going offline in the source managed disk group.	509
CMMVC6440E The command failed because the managed disk (MDisk) is currently being migrated.	503	CMMVC6462E The command failed because starting the migration will result in VDIsks going offline because the target managed disk group is offline.	509
CMMVC6441E The command failed because the managed disk group warning threshold is too low.	503	CMMVC6463E The command failed because Starting the migration will result in VDIsks going offline because a target MDisk is offline.	509
CMMVC6442E The command failed because the managed disk group warning threshold is too high.	503	CMMVC6464E The Create FlashCopy mapping task cannot be initiated because the size of the source VDisk is being changed by a previously submitted task.	510
CMMVC6443E The command failed because the managed disk group warning threshold is invalid.	503	CMMVC6465E The Create FlashCopy mapping task cannot be initiated because the size of the target VDisk is being changed by a previously submitted task.	510
CMMVC6444E The command failed because the virtual disk (VDisk) is in the process of being resized.	504	CMMVC6466E The Create FlashCopy mapping task cannot be initiated because an identical map already exists.	510
CMMVC6445E The command failed because one or more of the specified managed disks (MDisks) are in the process of being deleted.	504	CMMVC6467E The Create FlashCopy mapping task cannot be initiated because a FlashCopy map with the same target VDisk already exists in the consistency group.	510
CMMVC6446E The command failed because the managed disk groups have different extent sizes.	504		
CMMVC6447E The command failed because the virtual disk (VDisk) is currently being migrated.	504		
CMMVC6448E Deleting this node will cause data loss for resources associated with the IO group of this node.	505		

CMMVC6468E The Start or Prepare FlashCopy mapping task cannot be initiated because the target VDisk is the source of a different FlashCopy map that is being restored.	511	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.	515
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.	511	CMMVC6486E The task cannot be initiated because the user table is full.	516
CMMVC6470E The Start or Prepare FlashCopy mapping task cannot be initiated because the target VDisk is being used by a different FlashCopy map.	511	CMMVC6487E The task cannot be initiated because the user name that you have specified already exists.	516
CMMVC6471E The Create cluster partnership task cannot be initiated because a cluster in the existing partnership has a downlevel software version that does not support this configuration.	512	CMMVC6488E The task cannot be initiated because you have specified a user group ID that is not correct.	516
CMMVC6472E The Create cluster partnership task cannot be initiated because the remote cluster with which you are attempting to create a partnership has a downlevel software version that does not support this configuration.	512	CMMVC6489E The task cannot be initiated because you have specified more than one password.	516
CMMVC6473E The partnership task cannot be initiated because the supported maximum number of accessible remote clusters would be exceeded.	513	CMMVC6490E The task cannot be initiated because you have specified both a user group and the use of the remote authentication service.	517
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.	513	CMMVC6491E The task cannot be initiated because you have specified the use of the remote authentication service but you have not specified an ssh key and password.	517
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.	513	CMMVC6492E The task cannot be initiated because you have specified a local user but you have not specified a user group.	517
CMMVC6478E The Enable remote authentication service task cannot be initiated because the server settings are not configured.	514	CMMVC6493E The task cannot be initiated because the user that you have specified is not defined.	517
CMMVC6479E The task cannot be initiated because the user group table is full.	514	CMMVC6494E The task cannot be initiated because you cannot remove a default user.	518
CMMVC6480E The task cannot be initiated because the user group that you have specified is not defined.	514	CMMVC6495E The task cannot be initiated because the user superuser must be a local user.	518
CMMVC6481E The Modify user group task cannot be initiated because you have specified a default user group.	514	CMMVC6496E The task cannot be initiated because you cannot remove the superuser password.	518
CMMVC6482E The Delete user group task cannot be initiated because you have specified a default user group.	515	CMMVC6497E The task cannot be initiated because the user that you have specified does not have a password defined.	518
CMMVC6483E The task cannot be initiated because the user group name that you have specified already exists.	515	CMMVC6498E The task cannot be initiated because the user that you have specified does not have an ssh key defined.	518
CMMVC6484E The task cannot be initiated because the role that you have specified is not supported.	515	CMMVC6499E The task has failed because the ssh key that you have specified is already defined for another user.	519
		CMMVC6500E The action failed because the source and destination virtual disk (VDisk) are the same.	519
		CMMVC6501E The action failed because the node hardware is incompatible with the current I/O group member.	519
		CMMVC6502E The FlashCopy mapping was not prepared because preparing consistency group 0 is not a valid operation.	519
		CMMVC6503E The FlashCopy mapping or consistency group was not stopped because stopping consistency group 0 is not a valid operation.	519

CMMVC6504E The task cannot be initiated because the ssh key file that you have specified does not contain a valid ssh key..	520	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: * : , " ' %	526
CMMVC6505E The task has failed because an error has occurred while communicating with the authentication service..	520	CMMVC6525E The password that you have entered is not valid. The password cannot begin or end with a space character.	526
CMMVC6506E The task has failed because a timeout has occurred while communicating with the authentication service..	520	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.	527
CMMVC6507E The task has failed because the authentication service reports an incorrect user name or password..	520	CMMVC6528E The command cannot be initiated because the MDisk mode is not set to Array..	527
CMMVC6508E The task has failed because the authentication service reports that the authentication token has expired.	521	CMMVC6529E The command cannot be initiated because the maximum supported number of MDisks already exists.	527
CMMVC6509E The task has failed because the user name is not configured on the San Volume Controller (SVC) cluster.	521	CMMVC6530E The command cannot be initiated because the maximum supported number of arrays already exists.	528
CMMVC6510E The task has failed because the password is not correct for the user name that you are using..	522	CMMVC6531E The command cannot be initiated because array scrubbing is not supported for the RAID geometry of the array that you have specified.	528
CMMVC6511E The task has failed because San Volume Controller (SVC) is not configured correctly to use the authentication service. .	522	CMMVC6532E The command cannot be initiated because there is insufficient free memory that is available to the I/O group..	528
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.	522	CMMVC6533E The command cannot be initiated because the specified array member does not exist in the selected array.	528
CMMVC6513E The task has failed because you cannot activate a quorum disk until all of the quorum disks have been initialized..	523	CMMVC6534E The command cannot be initiated because the LDisk that you have specified does not exist.	529
CMMVC6514E The task has failed because the disk that you have selected to activate is not online.	523	CMMVC6535E The command cannot be initiated because there are insufficient available LDisks to configure an array using the RAID geometry that you have specified.	529
CMMVC6515E The task has failed because at least one quorum disk is in the Excluded state..	523	CMMVC6536E The command cannot be initiated because you have specified more LDisks than the specified RAID geometry permits.	529
CMMVC6516E The command has failed because an IPv4 cluster address cannot be removed while remote IPv4 services are configured. .	523	CMMVC6537E The command cannot be initiated because the LDisk that you have specified has a Use property that is not supported for the task.	529
CMMVC6517E The command has failed because an IPv6 cluster address cannot be removed while remote IPv6 services are configured. .	524	CMMVC6538E The command cannot be initiated because at least one of the LDisks that you have specified has a Use property that is not Candidate.	530
CMMVC6518E The task has failed because no roles are defined for the current user on the SAN Volume Controller (SVC) cluster.	524	CMMVC6539E The command cannot be initiated because the array does not have sufficient redundancy. The RAID state of the array is Degraded.	530
CMMVC6519E The task has failed because you cannot change the user group of the 'superuser' account to anything other than 'SecurityAdmin'. .	525	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.	530
CMMVC6520E You cannot use this task to modify the properties of the current user because those properties are only defined by an authentication service.	525		
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. .	525		
CMMVC6522E Authorization has failed.	526		
CMMVC6523E The URL that you have entered is not valid.	526		

CMMVC6541E The task cannot be initiated because the virtual capacity that you have requested for the space-efficient VDisk is larger than the maximum capacity that is supported for the extent size..	531	CMMVC6557E The task cannot be initiated because the file that you have specified is too large.	536
CMMVC6542E The remote authentication task has failed..	531	CMMVC6558E The command cannot be initiated because it might cause VDisks to go Offline. Refer to the SVC Command-Line Interface (CLI) command help for this command..	536
CMMVC6543E The task cannot be initiated because you can only specify a direct-attached managed drive when you submit the task.	531	CMMVC6559E The Add or Change e-mail 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'.	536
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..	532	CMMVC6560E The command has failed because the specified IP address is already in use by the cluster.	537
CMMVC6545E The Apply Drive Software task has failed to access the software download image..	532	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.	537
CMMVC6546E A device error was detected during the Apply Drive Software task.	532	CMMVC6562E The requested size exceeds the maximum supported value.	537
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.	532	CMMVC6563E The command has failed because a user that you have specified is not configured to receive e-mail notifications..	538
CMMVC6548E The Download FPGA firmware task cannot be initiated because the drive that you have selected is a managed disk.	533	CMMVC6564E You cannot make this user a remote user because the password type is not valid for a remote user.	538
CMMVC6549E The Authentication task has failed because the authentication service URL that you have specified is not a valid URL..	533	CMMVC6565E The command has failed because the specified node is offline..	538
CMMVC6550E The Authentication task has failed because the network address that is specified in the authentication service URL cannot be resolved..	533	CMMVC6566E The command cannot be submitted because specifying the -failover parameter requires that you also specify either the -name, -iscsialias or -noiscsialias parameter..	538
CMMVC6551E The Authentication task has failed because the combination of user name and password that is defined in the SVC for authorization by the authentication service is not defined on the authentication service.	534	CMMVC6567E The Apply Drive Software task has failed because no download images were found in the package file.	539
CMMVC6552E The Authentication task has failed because an SSL connection could not be established with the authentication service.	534	CMMVC6568E The Apply Drive Software task has failed because no download images were found in the package file for this drive type..	539
CMMVC6553E The task cannot be initiated because at least one quorum disk is not in the correct state..	534	CMMVC6569E The Apply Drive Software task has failed because no download images were found in the package file of this software type.	539
CMMVC6554E The Authentication task has failed because the user name that was received from the authentication service is not a valid SVC user name..	535	CMMVC6570E The command was not initiated because the cache mode of the virtual disk (VDisk) is already in the state that you had requested.	540
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.	535	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.	540
CMMVC6556E The task cannot be initiated because an error has occurred while attempting to read a file.	535	CMMVC6572E The command has failed because the I/O group that manages the virtual disk (VDisk) that you specified is not stable.	540
		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.	541

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.	541
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.	541
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.	542
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.	542
CMMVC6578E The command has failed because the iSCSI name is already assigned or is not valid.	542
CMMVC6579E The command cannot be initiated because the cluster Ethernet port 1 must always be fully configured in either the IPv4 or IPv6 format.	542

CMMVC6580E The command cannot be initiated because the iSCSI alias that you specified contained either leading or trailing space characters.	543
CMMVC6581E The command has failed because the iSCSI qualified name (IQN) is already assigned or is not valid.	543
CMMVC6582E The task has failed because the iSCSI host that you specified is not mapped to an I/O group.	543

Appendix. Accessibility 545

Notices 547

Trademarks	549
----------------------	-----

Index 551

Tables

1.	SAN Volume Controller library	xxxiii	10.	Relationship between the <i>rate</i> , data rate and grains per second values	146
2.	Other IBM publications	xxxv	11.	stoprconsistgrp consistency group states	170
3.	IBM documentation and related Web sites	xxxvi	12.	stoprrelationship consistency group states	171
4.	ip_address_list formats	56	13.	Valid filter attributes	185
5.	Memory required for VDisk Mirroring and Copy Services	58	14.	Attribute values	202
6.	Relationship between the <i>rate</i> value and the data copied per second	100	15.	MDisk output	255
7.	Relationship between the <i>rate</i> value and the data copied per second	111	16.	lsmdisklba command output	259
8.	Number of extents reserved by extent size	138	17.	svcinfo lsnod attribute values	269
9.	Relationship between the <i>rate</i> , data rate and grains per second values	143	18.	lsrconsistgrp command output values	282
			19.	lsrrelationship command attributes and values	285
			20.	lsvdisklba command output scenarios	314

About this guide

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

Who should use this guide

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

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

Summary of changes

This document contains terminology, maintenance, and editorial changes.

Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change. This summary of changes describes new functions that have been added to this release.

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

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

New information

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

- Cluster commands:
 - cfgportip
 - chclusterip
 - ping
 - rmpportip
- Cluster diagnostic and service-aid commands:
 - svctask applysoftware
 - svctask setlocale
 - svctask writesernum
- E-mail and event notification commands:
 - chemailserver
 - chsnmpserver
 - chsyslogserver
 - mkemailserver
 - mksnmpserver
 - mksyslogserver

- rmailserver
- rsmnmpserver
- rmsyslogserver
- Information commands:
 - lsclusterip
 - lscurrentuser
 - lsemailserver
 - lsiscsiauth
 - lsmdiskdumps
 - lsnodedependentvdisks
 - lspportip
 - lsquorum
 - lsrmdiskdependentmaps
 - lsroute
 - lssnmpserver
 - lssyslogserver
 - lsuser
 - lsusergrp
 - lsvdiskfcmappcopies
- Managed disk commands:
 - applymdisksoftware
 - triggermdiskdump
- Service mode information commands:
 - lsmdiskdumps
- User management commands:
 - chauthservice
 - chcurrentuser
 - chuser
 - chusergrp
 - mkuser
 - mkusergrp
 - rmuser
 - rmusergrp

Changed information

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

- Cluster commands:
 - chcluster
 - chnode
 - cleardumps
 - cpdumps
 - rmnode
 - setpwdreset
 - startstats
 - stopcluster

- Cluster diagnostic and service-aid commands:
 - applysoftware
 - setlocale
 - setevent
 - writesernum
- E-mail and event notification commands:
 - chemail
 - chemailuser
 - mkemailuser
 - rmemailuser
- FlashCopy[®] commands:
 - mkfcmap
 - prestartfcconsistgrp
 - prestartfcmap
 - rmfcmap
 - startfcconsistgrp
 - startfcmap
 - stopfcconsistgrp
 - stopfcmapping
- Host commands:
 - addhostport
 - chhost
 - mkhost
 - rmhostiogr
 - rmhostport
- Information commands:
 - lscluster
 - lsclustercandidate
 - lsclusterip
 - lscontroller
 - lsemailuser
 - lsfcmap
 - lshost
 - lshostvdiskmap
 - lsioatdumps
 - lsmdisk
 - lsmdiskdumps
 - lsmdiskgrp
 - lsmdisklba
 - lsnode
 - lsnodecandidate
 - lsnodevpd
 - lssoftwareupgradestatus
 - lsvdiskhostmap
 - lsvdisksyncprogress

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

The following SAN Volume Controller commands have been discontinued:

- Cluster diagnostic and service-aid commands:
 - svcservicetask addnode
 - svcservicetask applysoftware
 - svcservicetask rmnode
 - svcservicetask setlocale
 - svcservicetask writesernum
- E-mail and event notification commands:
 - setemail
- Error log commands:
 - svcservicetask cherrstate
 - svcservicetask clearerrlog
 - svcservicetask dumperrlog
 - svcservicetask finderr
 - svcservicetask setevent

- Information commands:
 - lsnodes
 - lsshkeys
- Role-based security commands:
 - mkauth
 - rmath
 - lsauth
- Secure Shell key commands:
 - addsshkey
 - rmallsshkeys
 - rmsshkey

Summary of changes for SC26-7903-04 SAN Volume Controller Command-Line Interface User's Guide

This topic describes the changes that have been made to the SAN Volume Controller *Command-Line Interface User's Guide* since the previous edition (SC26-7903-03).

New information

The following new SAN Volume Controller commands have been added for edition SC26-7903-04:

- Information commands:
 - lscimomdumps
- Service-mode information commands:
 - lscimomdumps
- Virtual disk commands:
 - recovervdisk
 - recovervdiskbycluster
 - recovervdiskbyiogrp

Changed information

The following commands have been updated for edition SC26-7903-04:

- addmdisk
- addvdiskcopy
- backup
- chcluster
- chcontroller
- chemail
- chfcmap
- chlicense
- chpartnership
- chvdisk
- clear
- cleardumps
- cpdumps

- lscluster
- lscontroller
- lsfabric
- lsiogrp
- lslicense
- lsnode
- lsnodevpd
- lsvdisk
- lsvdiskcopy
- migrateexts
- migratetoimage
- migratevdisk
- mkfcmap
- mkvdisk
- repairvdiskcopy
- rmmdisk
- rmnode
- rmpartnership
- rmvdisk
- setemail
- setevent

Emphasis

Different typefaces are used in this guide to show emphasis.

The following typefaces are used to show emphasis:

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

SAN Volume Controller library and related publications

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

SAN Volume Controller Information Center

The IBM System Storage SAN Volume Controller Information Center contains all of the information that is required to install, configure, and manage the SAN Volume

Controller. The information center is updated between SAN Volume Controller product releases to provide the most current documentation. The information center is available at the following Web site:

<http://publib.boulder.ibm.com/infocenter/svcic/v3r1m0/index.jsp>

SAN Volume Controller library

Table 1 lists and describes the publications that make up the SAN Volume Controller library. Unless otherwise noted, these publications are available in Adobe® portable document format (PDF) from the following Web site:

www.ibm.com/storage/support/2145

Table 1. SAN Volume Controller library

Title	Description	Order number
<i>IBM System Storage SAN Volume Controller Planning Guide</i>	This guide introduces the SAN Volume Controller and lists the features that you can order. It also provides guidelines for planning the installation and configuration of the SAN Volume Controller.	GA32-0551
<i>IBM System Storage SAN Volume Controller Model 2145-CF8 Hardware Installation Guide</i>	This guide provides the instructions that the IBM service representative uses to install the hardware for SAN Volume Controller model 2145-CF8.	GC52-1356
<i>IBM System Storage SAN Volume Controller Model 2145-8A4 Hardware Installation Guide</i>	This guide provides the instructions that the IBM service representative uses to install the hardware for SAN Volume Controller model 2145-8A4.	GC27-2219
<i>IBM System Storage SAN Volume Controller Model 2145-8G4 Hardware Installation Guide</i>	This guide provides the instructions that the IBM service representative uses to install the hardware for SAN Volume Controller model 2145-8G4.	GC27-2220
<i>IBM System Storage SAN Volume Controller Models 2145-8F2 and 2145-8F4 Hardware Installation Guide</i>	This guide provides the instructions that the IBM service representative uses to install the hardware for SAN Volume Controller models 2145-8F2 and 2145-8F4.	GC27-2221

Table 1. SAN Volume Controller library (continued)

Title	Description	Order number
<i>IBM System Storage SAN Volume Controller Software Installation and Configuration Guide</i>	This guide provides guidelines for configuring your SAN Volume Controller. Instructions for backing up and restoring the cluster configuration, using and upgrading the SAN Volume Controller Console, using the CLI, upgrading the SAN Volume Controller software, and replacing or adding nodes to a cluster are included.	SC23-6628
<i>IBM System Storage SAN Volume Controller CIM Agent Developer's Guide</i>	This guide describes the concepts of the Common Information Model (CIM) environment. Steps about using the CIM agent object class instances to complete basic storage configuration tasks, establishing new Copy Services relationships, and performing CIM agent maintenance and diagnostic tasks are included.	SC23-6665
<i>IBM System Storage SAN Volume Controller Command-Line Interface User's Guide</i>	This guide describes the commands that you can use from the SAN Volume Controller command-line interface (CLI).	SC26-7903
<i>IBM System Storage SAN Volume Controller Host Attachment Guide</i>	This guide provides guidelines for attaching the SAN Volume Controller to your host system.	SC26-7905
<i>IBM System Storage SAN Volume Controller Troubleshooting Guide</i>	This guide describes the features of each SAN Volume Controller model, explains how to use the front panel, and provides maintenance analysis procedures to help you diagnose and solve problems with the SAN Volume Controller.	GC27-2227
<i>IBM System Storage SAN Volume Controller Hardware Maintenance Guide</i>	This guide provides the instructions that the IBM service representative uses to service the SAN Volume Controller hardware, including the removal and replacement of parts.	GC27-2226
<i>IBM System Storage SAN Volume Controller Master Console Guide</i>	This guide describes how to install, maintain, and service the master console.	GC27-2223

Table 1. SAN Volume Controller library (continued)

Title	Description	Order number
<i>IBM Systems Safety Notices</i>	This guide contains translated caution and danger statements. Each caution and danger statement in the SAN Volume Controller documentation has a number that you can use to locate the corresponding statement in your language in the <i>IBM Systems Safety Notices</i> document.	G229-9054

Other IBM publications

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

Table 2. Other IBM publications

Title	Description	Order number
<i>IBM System Storage Productivity Center Introduction and Planning Guide</i>	This guide introduces the IBM System Storage Productivity Center hardware and software.	SC23-8824
<i>Read This First: Installing the IBM System Storage Productivity Center</i>	This guide describes how to install the IBM System Storage Productivity Center hardware.	GI11-8938
<i>IBM System Storage Productivity Center User's Guide</i>	This guide describes how to configure the IBM System Storage Productivity Center software.	SC27-2336
<i>IBM System Storage Multipath Subsystem Device Driver User's Guide</i>	This guide describes the IBM System Storage Multipath Subsystem Device Driver for IBM System Storage products and how to use it with the SAN Volume Controller.	GC52-1309
<i>Implementing the IBM System Storage SAN Volume Controller V4.3</i>	This IBM Redbooks® publication is a detailed technical guide to the IBM System Storage SAN Volume Controller. It provides a high-level overview of storage virtualization and the SAN Volume Controller architecture, discusses implementing and configuring the SAN Volume Controller, tells you how to migrate existing storage to the SAN Volume Controller, and discusses different supported migration activities.	SG24-6423

IBM documentation and related Web sites

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

Table 3. IBM documentation and related Web sites

Web site	Address
Support for SAN Volume Controller (2145)	www.ibm.com/storage/support/2145
Support for IBM System Storage and IBM TotalStorage® products	www.ibm.com/storage/support/
IBM Publications Center	www.ibm.com/shop/publications/order/
IBM Redbooks publications	www.redbooks.ibm.com/

Related accessibility information

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

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

How to order IBM publications

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

The IBM Publications Center offers customized search functions to help you find the publications that you need. Some publications are available for you to view or download at no charge. You can also order publications. The publications center displays prices in your local currency. You can access the IBM Publications Center through the following Web site:

www.ibm.com/shop/publications/order/

How to send your comments

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

- E-mail

Submit your comments electronically to the following e-mail address:
starpubs@us.ibm.com

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

- Mail


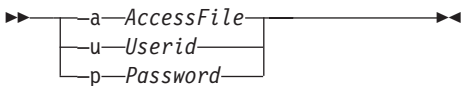

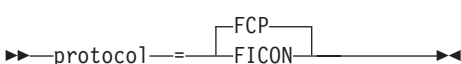
Fill out the Readers' Comments form (RCF) at the back of this book. If the RCF has been removed, you can address your comments to:

International Business Machines Corporation
RCF Processing Department
Department 61C

Syntax diagrams

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

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

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

Element	Syntax	Description
Repeatable keyword or value		Represents a parameter or argument that you can specify more than once. A repeatable keyword or value is represented by an arrow returning to the left above the keyword or value.
Variable		Represents the value that you need to supply for a parameter or argument, such as a file name, user name, or password. Variables are in italics.
Space separator		Adds a blank space on the main path line to separate keywords, parameters, arguments, or variables from each other.
Quotation mark delimiters		Indicates the start and end of a parameter or argument that contains multiple values. Enclose one or more name-value pairs in a set of double quotation marks for a particular parameter or argument. If the value of a parameter or name-value pair contains a blank or white space, enclose the entire value in a set of single quotation marks.
Equal-sign operator		Separates a name from its value in a name-value pair.
Syntax fragment		Breaks up syntax diagrams that are too long, too complex, or repetitious. The fragment name is inserted in the main diagram, and the actual fragment is shown below the main diagram.

Terminology

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

Name	Object type
Host	host
Virtual disk	vdisk
Virtual disk copy	vdiskcopy
Space-efficient virtual disk copy	sevdiskcopy

Name	Object type
Managed disk	mdisk
Managed disk group	mdiskgrp
I/O group	iogrp
Node	node
Cluster	cluster
Controller	controller
IBM FlashCopy mapping	fcmap
FlashCopy consistency group	fcconsistgrp
Metro Mirror or Global Mirror relationship	rcrelationship
Metro Mirror or Global Mirror consistency group	rcconsistgrp
Unsupported/unknown object	unknown

CLI special characters

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

minus (-) sign

Flags are prefixed with a - (minus) sign. Flags define the action of a command or modify the operation of a command. You can use multiple flags, followed by parameters, when you issue a command. The - character cannot be used as the first character of an object name.

vertical bar (|)

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

Using wildcards in the SAN Volume Controller CLI

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

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

1. Running the command while logged onto the node.

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

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

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

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

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

2. Running the command through Secure Shell (SSH), for example from a host. This method is slightly more complicated because the host shell processes the command line before it is passed through SSH to the shell on the cluster. This means an extra layer of protection is required around the wildcard as the host shell will strip off any protecting quotes, and if the wildcard is exposed to the cluster shell, this will result in the wildcard being expanded in the cluster shell. To prevent expansion, issue the following command in one of its formats:

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

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

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

Data types and value ranges

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

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

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

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

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

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

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

Data types	Value ranges
stats_time_arg	The frequency at which statistics are gathered. Valid values are 1 to 60 minutes in increments of 1 minute.
directory_arg	<p>Specifies a directory and or file name filter within the specified directory. The following directory values are valid:</p> <ul style="list-style-type: none"> • /dumps • /dumps/audit • /dumps/cimom • /dumps/configs • /dumps/eglogs • /dumps/feature • /dumps/iostats • /dumps/iotrace • /home/admin/upgrade <p>The file name filter can be any valid file name with or without the wildcard '*'.</p> <p>The file name filter can be appended to the end of one of the previous directory values.</p>
locale_arg	<p>The cluster locale setting. Valid values are 0 and 3.</p> <ul style="list-style-type: none"> • 0 US English (default) • 3 Japanese
key_arg	A user-defined identifier for an SSH key. It consists of a string of up to 30 characters.
user_arg	Specifies the user: either admin or service.
copy_rate	A numeric value of 0 - 100.
copy_type	Specifies the Mirror copy type: Metro or Global.

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

CLI commands and parameters

CLI commands and parameters are represented in the syntax diagram.

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

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

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

A valid parameter meets the following requirements:

- Parameters can be entered in any order.

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

```
svctask chuser -usergrp=-usergrp -- -password
```
- The symbol '--' is valid as the final word on the command line.

Examples that are valid:

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

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

Examples that are invalid:

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

CLI flags

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

-? or -h

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

-nomsg

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

```
svctask mkmdiskgrp -ext 16
```

it displays:

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

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

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

the following information is displayed:

```
6
```


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

CLI messages

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

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

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

Chapter 1. Preparing the SSH client system for the CLI

Before you can issue command-line interface (CLI) commands from the host to the cluster, you must prepare the Secure Shell (SSH) client system.

Microsoft Windows operating systems

The IBM System Storage Productivity Center (SSPC) and the master console for the SAN Volume Controller include the PuTTY client program, which is a Microsoft® Windows® SSH client program. The PuTTY client program can be installed on your SSPC or master console server in one of the following ways:

- If you purchased the SSPC or the master console hardware option from IBM, the PuTTY client program has been preinstalled on the hardware.
- You can use the master console software installation CD to install the PuTTY client program. The SSPC, master console hardware option, and the software-only master console each provide this CD.
- You can use the separate PuTTY client program-installation wizard, **putty-version-installer.exe**. You can download the PuTTY client program from the following Web site:

www.chiark.greenend.org.uk/~sgtatham/putty/

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

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

www.openssh.org/windows.html

IBM AIX operating systems

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

You can also obtain the AIX installation images from IBM developerWorks® at the following Web site:

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

Linux operating systems

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

www.openssh.org/portable.html

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

www.openssh.org/portable.html

Preparing the SSH client system to issue CLI commands

To issue command-line interface (CLI) commands to the cluster from a host, you must prepare the Secure Shell (SSH) client on the host so that the host is accepted by the SSH server on the cluster.

To use a host that requires a different type of SSH client, OpenSSH for example, follow the instructions for that software.

Perform the following steps to enable your host to issue CLI commands:

1. For the IBM System Storage Productivity Center or master console and Windows hosts:
 - a. Generate an SSH key pair using the PuTTY key generator.
 - b. Store the SSH clients public key on the cluster (using a browser that points to the SAN Volume Controller Console).
 - c. Configure the PuTTY session for the CLI.
2. For other types of hosts:
 - a. Follow the instructions that are specific to the SSH client to generate an SSH key pair.
 - b. Store the SSH clients public key on the cluster (using a Web browser to point to the SAN Volume Controller Console or the CLI from an already established host).
 - c. Follow the instructions that are specific to the SSH client to establish an SSH connection to the SAN Volume Controller cluster.

Chapter 2. Secure Shell

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

Overview

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

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

Authenticating SSH logins

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

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

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

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

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

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

Configuring a Secure Shell client system other than PuTTY

If you use a Secure Shell (SSH) client system other than the PuTTY client, you must configure that client system before you can access the command-line interface (CLI).

Perform the tasks that are equivalent to the following steps to configure your non-PuTTY SSH client system:

1. Install the SSH client software on the computer that will host the master console.

Note: This step is not required for a master console that has the preinstalled PuTTY client.

2. Generate SSH keys on the SSH client system.
3. Configure the PuTTY session, if required, on the SSH client system.
4. If the client system is the master console, copy the private key into the SAN Volume Controller installation directory. If the client system is not the master console, store the private key on the SSH client system.

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 to cause data to be lost.

5. Copy the SSH public key to the master console.
6. Store the SSH client public key on the SAN Volume Controller cluster.

You perform step 6 to store the SSH client public key on the SAN Volume Controller when you complete the creation of the SAN Volume Controller cluster. After you define a cluster to the SAN Volume Controller Console and therefore enable SSH communication to the cluster, you can store additional SSH client public keys on the cluster. You can store additional keys through the SAN Volume Controller Console or the CLI.

Generating an SSH key pair using PuTTY

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

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

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

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

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

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

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

- c. Click **Save**.
6. Save the private key by performing the following steps:
 - a. Click **Save private key**. The PuTTYgen Warning panel is displayed.
 - b. Click **Yes** to save the private key without a passphrase.
 - c. Type `icat` as the name of the private key, and specify the location where you want to save the private key. For example, you can create a directory on your computer called `keys` to store both the public and private keys. It is recommended that you save your public and private keys in the same location.
 - d. Click **Save**.
7. Close the PuTTY Key Generator window.

Configuring a PuTTY session for the CLI

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

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

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

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

You have now configured a PuTTY session for the CLI.

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

Creating users

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

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

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

This task assumes that you have already launched the SAN Volume Controller Console. Complete the following steps to create either a local or remote user:

1. Click **Manage Authentication** → **Users** in the portfolio. The Viewing Users panel is displayed.
2. Select **Create a User** from the task list and click **Go**. The Creating a User panel is displayed.
3. Enter a name for the user.
4. Enter a password for the user. The password cannot start or end with a blank character. The password can consist of a string of 6 - 64 printable ASCII characters.
5. Enter the SSH key file that is associated with the user. Click **Browse** to select the file. An SSH key is needed if this user plans to use the command-line interface to manage the cluster. Any SAN Volume Controller users that use the remote authentication service and require SSH keys to access the command-line interface must have the same password on the cluster and the remote authentication service. In addition the user group that the user belongs to must be visible to the remote authentication service. The remote visibility setting instructs SAN Volume Controller to check the remote authentication service for that user's group information to determine the user's role on the cluster.
6. Select the appropriate authentication type for the user. Select **Remote** if the user is authenticated to the cluster by a remote authentication service. Select **Local** if the user is authenticated to the cluster using cluster authentication methods.

| **Note:** Local is the default setting for the authentication type.

- | 7. If you selected to create a local user, you must also specify the user group that
| the user belongs to. The user group defines roles that provide the user with
| access to specific operations on the cluster.
- | 8. Click **OK**.

Adding SSH keys for hosts other than the IBM System Storage Productivity Center or the master console

You can add Secure Shell (SSH) keys on other hosts.

Perform the following steps to add SSH keys on hosts:

1. Generate the public-private key pair on each host that you want to use the SAN Volume Controller command-line interface. See the information that came with your SSH client for specific details about using the key generation program that comes with your SSH client.
2. Copy the public keys from each of these hosts to the IBM System Storage Productivity Center or the master console.
3. Use the PuTTY secure copy function to copy these public keys from the IBM System Storage Productivity Center or the master console to the cluster.
4. Repeat 3 for each public key that you copied in step 2.

Chapter 3. Copying the SAN Volume Controller software upgrade files using PuTTY scp

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

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

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

Perform the following steps to use the pscp application:

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

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

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

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

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

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

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

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

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

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

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

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

- c. Then repeat step 6.

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 SAN Volume Controller Console.

The audit log entries provide the following information:

- Identity of the user who issued the action command
 - From the command-line interface, the user name (administrator or service), and the label that is associated with the user's public SSH key in the authorized keys file
 - From the native Web pages, the user's identity (admin[web] or service[web]) according to which user name the user authenticated with
 - From the SAN Volume Controller Console, the user's identity (administrator), the label that is associated with the CIMOM key in the authorized keys file, and the user name that has been recorded by the CIMOM when the SAN Volume Controller Console user authenticated with the CIMOM
- The name of the actionable command
- The timestamp of when the actionable command was issued on the configuration node
- The parameters which were issued with the actionable command

The following commands are not documented in the audit log:

- svctask dumpconfig
- svctask cpdumps
- svctask clear.dumps
- svctask finderr
- svctask dumperrlog
- svctask dumpinternallog

The following items are also not documented in the audit log:

- Commands that fail are not logged
- A result code of 0 (success) or 1 (success in progress) is not logged
- Result object ID of node type (for the addnode command) is not logged
- Views are not logged

catauditlog

Use the **catauditlog** command to display the in-memory contents of the audit log.

Syntax

```
▶▶ svcinfo — — catauditlog — —————▶
```

```
▶ —————▶
```

```
└─ -first — number_of_entries_to_return ─┘
```

Parameters

-first *number_of_entries_to_return*

(Optional) Specifies the number of most recent entries to display.

Description

This command lists a specified number of the most recently audited commands.

The in-memory portion of the audit log can hold approximately 1 MB of audit information. Depending on the command text size and the number of parameters, 1 MB records approximately 6000 commands.

Once the in-memory audit log has reached its maximum capacity, the log is written to a local file on the configuration node in the `/dumps/audit` directory. The **svcinfo 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 need any special command to decode it.

The in-memory log entries are reset and cleared automatically, ready to start accumulating new commands. The on-disk portion of the audit log can then be analyzed at a later date.

The `lsauditlogdumps` command can be used to list the files that are on the disk.

The in-memory portion of the audit log can be transferred to an on-disk file using the **svctask dumpauditlog** command. This action clears the in-memory portion of the log.

In the following example, the user has specified that they want to list the 15 most recent audit log entries.

An invocation example

```
svcinfo catauditlog -delim : -first 15
```

The resulting output

```
audit_seq_no:timestamp:cluster_user:ssh_label:icat_user:result:res_obj_id
:action_cmd
125:060311111800:admin:Joe::0::svctask rmsshkey -key label47 -user admin
126:060311111800:admin:Joe::0::svctask addsshkey -label label48 -file
/home/Joe/id_rsa.pub -user admin
127:060311111800:admin:Joe::0::svctask rmsshkey -key label48 -user admin
128:060311111800:admin:Joe::0::svctask addsshkey -label label49 -file
/home/Joe/id_rsa.pub -user admin
129:060311111800:admin:Joe::0::svctask rmsshkey -key label49 -user admin
130:060311134617:admin:Joe::0::svctask chmdisk -name ca-0 1
131:060311134617:admin:Joe::0::svctask chmdisk -name ca-1 2
132:060311134617:admin:Joe::0::svctask chmdisk -name ca-2 3
133:060311134617:admin:Joe::0::svctask chmdisk -name cb-0 4
134:060311134617:admin:Joe::0::svctask chmdisk -name cb-1 5
135:060311134617:admin:Joe::0::svctask chmdisk -name cb-2 6
136:060311134617:admin:Joe::0::svctask chmdisk -name cc-0 7
137:060311134617:admin:Joe::0::svctask chmdisk -name cc-1 8
138:060311134617:admin:Joe::0::svctask chmdisk -name cc-2 9
139:060311134632:admin:Joe::0::svctask mkmdiskgrp -name custa-mdisks -ext 512
-mdisk ca-0:ca-1:ca-2
```

dumpauditlog

Use the **dumpauditlog** command to reset or clear the contents of the in-memory audit log. The contents of the audit log are sent to a file in the `/dumps/audit` directory on the current configuration node.

Syntax

▶— svctask — — dumpauditlog — —————▶

Parameters

There are no parameters.

Description

This command dumps the contents of the audit log to a file on the current configuration node. It also clears the contents of the audit log. This command is logged as the first entry in the new audit log.

Audit log dumps are automatically maintained in the `/dumps/audit` directory. The local file system space is used by audit log dumps and is limited to 200 MB on any node in the cluster. The space limit is maintained automatically by deleting the minimum number of old audit log dump files so that the `/dumps/audit` directory space is reduced below 200 MB. This deletion occurs once per day on every node in the cluster. The oldest audit log dump files are considered to be the ones with the lowest audit log sequence number. Also, audit log dump files with a cluster ID number that does not match the current one are considered to be older than files that match the cluster ID, regardless of sequence number.

Other than by running dumps (or copying dump files among nodes), you cannot alter the contents of the audit directory. Each dump file name is generated automatically in the following format:

```
auditlog_firstseq_lastseq_timestamp_clusterid
```

where

- *firstseq* is the audit log sequence number of the first entry in the log
- *lastseq* is the audit sequence number of the last entry in the log
- *timestamp* is the timestamp of the last entry in the audit log that is being dumped
- *clusterid* is the cluster ID at the time that the dump was created

The audit log dump files names cannot be changed.

The audit log entries in the dump files contain the same information as displayed by the **svcinfo catauditlog** command; however, the **svctask dumpauditlog** command displays the information with one field per line. The **svcinfo lsauditlogdumps** command displays a list of the audit log dumps that are available on the nodes in the cluster.

An invocation example

```
svctask dumpauditlog
```

The resulting output

lsauditlogdumps

The **lsauditlogdumps** command generates a list of the audit log dumps that are available on the nodes in the cluster.

Syntax

```

▶▶ svcinfo — lsauditlogdumps — [ -nohdr ]
▶ [ -delim delimiter ] [ node_id | node_name ]

```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. 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.

node_id | *node_name*

(Optional) Specifies the node ID or name to list the available dumps of the given type. If you do not specify a node, the files on the current configuration node are displayed.

Description

This command lists the dump files that are in the **/dumps/audit** directory on the specified node, or on the configuration node if a node is not specified.

The cluster automatically creates the audit log. The audit log can also be created manually by issuing the **svctask dumpauditlog** command. The audit log comprises the files that are listed by the **svcinfo lsauditlogdumps** command. These files are limited to approximately 200 MB on each node in the cluster, at which point the oldest files are automatically deleted. When the configuration node changes to a different node in the cluster, any old audit log files are left on the former configuration node. As with other types of dumps, you can retrieve those files using the **cpdumps** command.

An invocation example

```
svcinfo lsauditlogdumps
```


The resulting output

```
id auditlog_filename
0 auditlog_0_229_060311234532_0000020060013d8a
1 auditlog_230_475_060312234529_0000020060013d8a
2 auditlog_476_491_060313234527_0000020060013d8a
```

Audit log dump file contents

```
...
Auditlog Entry:23
Audit Sequence Number :138
Timestamp :Sat Mar 11 13:46:17 2006
:Epoch + 1142084777
SVC User :admin
SSH Label :Joe
ICAT User :
Result Object ID :
Result Code :0
Action Command :svctask chmdisk -name cc-2 9
Auditlog Entry:24
Audit Sequence Number :139
Timestamp :Sat Mar 11 13:46:32 2006
:Epoch + 1142084792
SVC User :admin
SSH Label :Joe
ICAT User :
Result Object ID :
Result Code :0
Action Command :svctask mkmdiskgrp -name custa-mdisks -ext
512 -mdisk ca-0:ca-1:ca-2
...
```

Chapter 5. User management commands

You can use the command-line interface (CLI) to configure remote authentication service and manage users and user groups on the cluster.

The following user management commands are available:

chauthservice

Configures the remote authentication service of the cluster.

chcurrentuser

Changes the attributes of the current user.

chuser Changes the attributes of an existing user.

chusergrp

Changes the attributes of an existing user group.

lscurrentuser

Displays the name and role of the logged-in user. See the “lscurrentuser” on page 210 command in the Chapter 21, “Information commands,” on page 195 section.

lsuser Lists the users who have been created on the cluster. See the “lsuser” on page 299 command in the Chapter 21, “Information commands,” on page 195 section.

lsusergrp

Lists the user groups that have been created on the cluster. See the “lsusergrp” on page 300 command in the Chapter 21, “Information commands,” on page 195 section.

mkuser

Creates either a local or a remote user to access a SAN Volume Controller cluster.

mkusergrp

Creates a new user group.

rmuser

Removes a user.

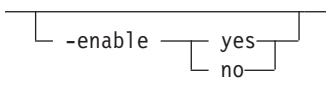
rmusergrp

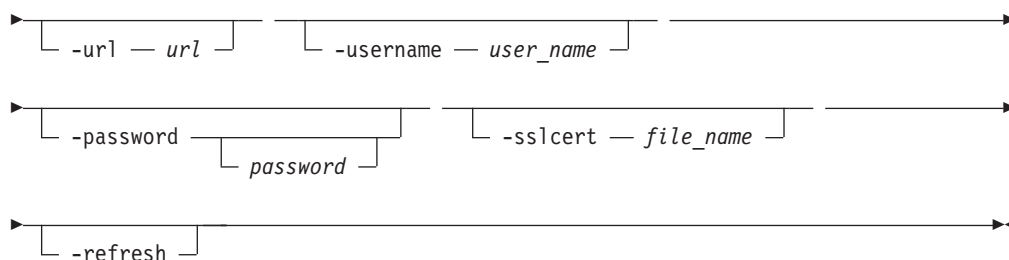
Removes a user group.

chauthservice

The chauthservice command can be used to configure the remote authentication service of the cluster.

Syntax

►► svctask — — chauthservice — 



Parameters

-enable *yes | no*

(Optional) Enables or disables the SAN Volume Controller system's use of the remote authentication server. When the **enable** parameter is set to no, remote authentications are failed by the system, but local authentications continue to operate normally.

-url *url*

(Optional) Specifies the Web address of the remote authentication service. This must be a valid IPv4 or IPv6 network address. You can use the following characters: a - z, A - Z, 0 - 9, -, ~, :, [,], %, or /. The maximum length of the Web address is 100 characters.

-username *user_name*

(Optional) Specifies the HTTP basic authentication user name. The user name cannot start or end with a blank. The user name can consist of a string of 1 - 64 ASCII characters with the exception of the following characters: %, ", * .

-password *password*

(Optional) Specifies the HTTP basic authentication user password. The password cannot start or end with a blank. It must consist of a string of 6 - 64 printable ASCII characters. The *password* variable is optional. If you do not provide a password, the system prompts you and does not display the password that you type.

-sslcert *file_name*

(Optional) Specifies the name of the file that contains the SSL certificate, in privacy enhanced mail (PEM) format, for the remote authentication server.

-refresh

(Optional) Causes the SAN Volume Controller to invalidate any remote user authorizations that are cached on the cluster. Use this when you modify user groups on the authentication service and want the change to immediately take effect on the SAN Volume Controller.

Description

This command sets the attributes of the remote authentication service on the cluster. It is not necessary to disable the remote authentication service to change its attributes. To disable the remote authentication service in a controlled manner when it is not available, use the **enable** parameter with the no option.

When the authentication service is enabled or the configuration is changed, the cluster does not test whether the remote authentication system is operating correctly. To establish whether the system is operating correctly, issue the command `svcinfo lscurrentuser` for a remotely authenticated user. If the output lists the user's roles obtained from the remote authentication server, remote authentication is

| operating successfully. If the output is an error message, remote authentication is
| not working correctly, and the error message describes the problem.

| If you are using the `url` parameter, the Web address can have either of the
| following formats:

- | • `http://network_address:http remote authentication service port`
| `number/path_to_service`
- | • `https://network_address:https remote authentication service port`
| `number/path_to_service`

| The network address must be an IPv4 or IPv6 address. Do not use the
| corresponding host name.

| For example, if the system network IPv4 address is 9.71.45.108, you could enter
| either of the following corresponding addresses:

| `http://9.71.45.108:16310/TokenService/services/Trust`
| `https://9.71.45.108:16311/TokenService/services/Trust`

| **An invocation example**

| To fully configure and enable the authentication service:

| `svctask chauthservice -url https://9.71.45.108:16311/TokenService/services/Trust`
| `-sslcert /tmp/sslCACert.pem -username admin -password password -enable yes`

| **The resulting output**

| No feedback

| **An invocation example**

| To disable remote authentication:

| `svctask chauthservice -enable no`

| **The resulting output**

| No feedback

| **An invocation example**

| To switch to an HTTPS connection to the authentication service:

| `svctask chauthservice -url https://9.71.45.108:16311/TokenService/services/Trust`
| `-sslcert /tmp/ssl_cert.pem`

| **The resulting output**

| No feedback

| **An invocation example**

| To refresh the SAN Volume Controller remote authorization cache:

| `svctask chauthservice -refresh`


| **The resulting output**

| No feedback

chcurrentuser

The chcurrentuser command changes the attributes of the current user.

Syntax



Parameters

-password *cleartext_password*

(Optional) Specifies the new password to be associated with the current user. The password cannot start or end with a blank. It must consist of a string of 6 - 64 printable ASCII characters. You can optionally specify the password with the **password** parameter. If you do not specify the password, the system prompts you for it before running the command and does not display the password that you type. Either the **password** parameter or the **nopassword** parameter can be set.

-nopassword

(Optional) Specifies that the user's password is to be deleted.

-keyfile *sshkey_filename*

(Optional) Specifies the name of the file that contains the Secure Shell (SSH) public key. Either the **keyfile** parameter or the **nokey** parameter can be set.

-nokey

(Optional) Specifies that the user's SSH key is to be deleted.

Description

Use the chcurrent user command to modify the attributes of the current user.

An invocation example

```
svctask chcurrentuser -password secret -nokey
```


The resulting output


No feedback

chuser

The chuser command changes the attributes of an existing user.

Syntax





Parameters

-password *cleartext_password*

(Optional) Specifies the new password to be associated with the user. The password CANNOT start or end with a blank. It must consist of a string of 6 - 64 printable ASCII characters. You can optionally specify the password with the **password** parameter. If you do not specify the password, the system prompts you for it before running the command and does not display the password that you type. Either the **password** parameter or the **nopassword** parameter can be set.

-nopassword

(Optional) Specifies that the user's password is to be deleted.

-keyfile *sshkey_filename*

(Optional) Specifies the name of the file that contains the Secure Shell (SSH) public key. Either the **keyfile** parameter or the **nokey** parameter can be set.

-nokey

(Optional) Specifies that the user's SSH key is to be deleted.

-remote **yes** | **no**

(Optional) Specifies whether the user authenticates to the cluster using a remote authentication service. Either the **yes** or **no** option must be set.

-usergrp *group_id_or_name*

(Optional) Specifies the new group for the user.

user_id_or_name

(Required) Specifies the user whose attributes are to be changed.

Description

Use the `chuser` command to modify the attributes of an existing user.

You must have the Security Administrator role to create, delete, or change a user.

Only use the **usergrp** parameter for local users. If you change a user from local to remote, the user's association with any group is removed.

If you change a user from remote to local, a user group must be specified. If you change a user from local to remote, the user must have both a password and an SSH key.

If you use the **keyfile** parameter, the SSH key file should be placed in the `/tmp` directory before running this command. When you run the command, the SSH key is copied into cluster state and activated for the user, and the input file is deleted.

An invocation example

```
svctask chuser -remote no -usergrp Monitor -nokey jane
```

The resulting output

No feedback

chusergrp

The chusergrp command changes the attributes of an existing user group.

Syntax

```
svctask -- chusergrp -- [-role -- role_name] --
                        [-remote -- yes | no] --
                        groupid_or_name
```

Parameters

-role *role_name*

(Optional) Specifies the role to be associated with users that belong to this group. One of the following roles must be selected: Monitor, CopyOperator, Service, Administrator, or SecurityAdmin.

-remote **yes** | **no**

(Optional) Specifies whether this user group should be used to set the role of remote users. Either the yes or no option must be set.

groupid_or_name

(Required) Specifies the user group whose attributes are to be changed.

Description

Use the chusergrp command to modify the attributes of an existing user group.

You must have the Security Administrator role to create, delete, or change a user.

The roles of the default groups cannot be changed.

An invocation example

```
svctask chusergrp -role Administrator admin
```

The resulting output


No feedback

mkuser

The mkuser command creates either a local or a remote user to access a SAN Volume Controller cluster.

Syntax

```
svctask -- mkuser -- -name -- user_name --
        [-remote --
         -usergrp -- group_name_or_id]
```

Parameters

-name *user_name*

(Required) Specifies the unique user name. The user name cannot start or end with a blank. The user name must consist of a string of 1 - 256 ASCII characters, with the exception of the following characters: %: ",*' .

-remote | **-usergrp** *group_name_or_id*

(Required) Specifies whether the user authenticates to the cluster using a remote authentication service or cluster authentication methods. The **usergrp** parameter specifies the name or ID of the user group with which the local user is to be associated. Either the **remote** parameter or the **usergrp** parameter must be set.

-password *cleartext_password*

(Optional) Specifies the password to be associated with the user. The password cannot start or end with a blank. It must consist of a string of 6 - 64 printable ASCII characters. You can optionally specify the password with the **password** parameter. If you do not specify the password, the system prompts you for it before running the command and does not display the password that you type.

-keyfile *sshkey_filename*

(Optional) Specifies the name of the file that contains the Secure Shell (SSH) public key.

Description

The `mkuser` command creates a new local or remote user to access a cluster. The command returns the ID of the created user.

You must have the Security Administrator role to create, delete, or change a user.

If you create a local user, you must specify the existing user group that the user belongs to. All local users must have a group. The user group defines roles that provide the user with access to specific operations on the cluster. You must also specify either the **keyfile** or **password** parameter, or both.

If you create a remote user, you must specify both the **keyfile** and **password** parameters. Remote users have their groups defined by the remote authentication service.

If you use the **keyfile** parameter, the SSH key file should be placed in the `/tmp` directory before running this command. When you run the command, the SSH key is copied into cluster state and activated for the user, and the input file is deleted.

An invocation example

```
svctask mkuser -name jane -usergrp Service -password secret
```

The resulting output

```
User, id [1], successfully created
```

mkusergrp

The mkusergrp command creates a new user group.

Syntax

```
svctask -- mkusergrp -- -name -- group_name -- -role -- role_name --  
-remote
```

Parameters

-name *group_name*

(Required) Specifies the unique user group name. The group name cannot start or end with a blank. The group name must consist of a string of 1 - 64 ASCII characters, with the exception of the following characters: %:"*'

-role *role_name*

(Required) Specifies the role to be associated with all users that belong to this user group. One of the following roles must be selected: Monitor, CopyOperator, Service, Administrator, or SecurityAdmin.

-remote

(Optional) Specifies that this user group should be used to set the role of remote users. This is disabled by default.

Description

The mkusergrp command creates a new user group to organize users of the SAN Volume Controller cluster by role. Use the lsusergrp command to view a list of user groups that have been created on the cluster.

You must have the Security Administrator role to create, delete, or change a user group.

Each user group has one role that determines the role of users that belong to that group. Use the **role** parameter to specify one of the following roles for the user group:

Monitor

Users can issue the following commands: svctask finderr, dumperrlog, dumpinternallog, chcurrentuser, ping, svcconfig backup, and all of the information commands.

Copy Operator

Users can issue the following commands: prestartfcconsistgrp, startfcconsistgrp, stopfcconsistgrp, chfcconsistgrp, prestartfcmap, startfcmap, stopfcmap, chfcmap, startrcconsistgrp, stoprcconsistgrp, switchrcconsistgrp, chrconsistgrp, startrcrelationship, stoprcrelationship, switchrcrelationship, chrrelationship, and chpartnership. In addition, users can issue all of the commands allowed by the Monitor role.

Service

Users can issue the following commands: applysoftware, setlocale, addnode, rmnode, cherrstate, writesernum, detectmdisk, includemdisk,

clearerrlog, cleardumps, settimezone, stopcluster, startstats, stopstats, and settime. In addition, users can issue all of the commands allowed by the Monitor role.

Administrator

Users can issue all commands except: chauthservice, mkuser, rmuser, chuser, mkusergrp, rmusergrp, chusergrp, and setpwdreset

Security Administrator

Users can issue all commands.

The command returns the ID of the created user group.

An invocation example

```
svctask mkusergrp -name support -role Service
```

The resulting output

```
User Group, id [5], successfully created
```

rmuser

The rmuser command deletes a user.

Syntax

```
►► svctask — — rmuser — — user_id_or_name —————►►
```

Parameters

user_id_or_name
(Required) Specifies the user to be removed.

Description

Use the rmuser command to delete a user.

You must have the Security Administrator role to create, delete, or change a user.

An invocation example

```
svctask rmuser jane
```

The resulting output

```
No feedback
```

rmusergrp

The rmusergrp command deletes a user group.

Syntax

```
►► svctask — — rmusergrp — — group_id_or_name —————►►  
                                  └-force-
```

Parameters

-force

(Optional) Specifies that the user group should be deleted even if there are users in the group.

group_id_or_name

(Required) Specifies the ID or name of the user group to be removed.

Description

Use the `rmusergrp` command to delete a user group.

You must have the Security Administrator role to create, delete, or change a user group.

User groups with users cannot normally be deleted. If you use the **force** parameter, the group is deleted and all of the users in that group are assigned to the Monitor group. Default user groups cannot be deleted, even if the **force** parameter is set.

An invocation example

```
svctask rmusergrp support
```

The resulting output

No feedback

Chapter 6. Role-based security commands (Discontinued)

| **Attention:** The role-based security commands are discontinued. Use the user
| management commands to configure remote authentication service and manage
| users and user groups on the cluster.

mkauth (Discontinued)

| **Attention:** The mkauth command is discontinued. Use the user management
| commands to configure remote authentication service and manage users and user
| groups on the cluster.

rmauth (Discontinued)

| **Attention:** The rmath command is discontinued. Use the user management
| commands to configure remote authentication service and manage users and user
| groups on the cluster.

lsauth (Discontinued)

| **Attention:** The lsauth command is discontinued. Use the user management
| commands to configure remote authentication service and manage users and user
| groups on the cluster.

Chapter 7. E-mail and event notification commands

You can use the command-line interface (CLI) to enable your system to send notifications.

The following e-mail and event notification commands are available:

chcluster

Enables the inventory notification function by specifying how often notifications are sent to recipients. See the “chcluster” on page 51 command in the Chapter 8, “Cluster commands,” on page 47 section.

chemail

Sets or modifies contact information for e-mail event notifications.

chemailserver

Modifies the parameters of an existing e-mail server object.

chemailuser

Allows you to modify the e-mail recipient’s information.

chsnmpserver

Modifies the parameters of an existing SNMP server.

chsyslogserver

Modifies the parameters of an existing syslog server.

lscluster

Provides information about the addition of e-mail notification recipients to the cluster. See the “lscluster” on page 200 command in the Chapter 21, “Information commands,” on page 195 section.

lsemailserver

Returns a concise list or a detailed view of e-mail servers that are configured on the cluster. See the “lsemailserver” on page 212 command in the Chapter 21, “Information commands,” on page 195 section.

lsemailuser

Generates a report that lists the users (by user ID or type of user [local or support]) who are receiving e-mail event and inventory notifications. See the “lsemailuser” on page 213 command in the Chapter 21, “Information commands,” on page 195 section.

lssnmpserver

Returns a concise list or a detailed view of SNMP servers that are configured on the cluster. See the “lssnmpserver” on page 294 command in the Chapter 21, “Information commands,” on page 195 section.

lssyslogserver

Returns a concise list or a detailed view of syslog servers that are configured on the cluster. See the “lssyslogserver” on page 296 command in the Chapter 21, “Information commands,” on page 195 section.

mkemailserver

Creates an e-mail server object that describes a remote Simple Mail Transfer Protocol (SMTP) e-mail server.

mkemailuser
 Allows you to specify the user, the receiving server address, the type of e-mail event notification, and to enable inventory notification in the e-mail event notification facility.

mksnmpserver
 Creates an SNMP server to receive notifications.

mksyslogserver
 Creates a syslog server to receive notifications.

rmemailserver
 Deletes the specified e-mail server object.

rmemailuser
 Removes an existing e-mail recipient from the e-mail event notification function.

rmsnmpserver
 Deletes the specified SNMP server.

rmsyslogserver
 Deletes the specified syslog server.

sendinventoryemail
 Sends an inventory e-mail notification to all inventory e-mail recipients.

setemail
 The setemail command is discontinued. E-mail notification can be configured using the following commands: svctask mkemailserver, svctask chemailserver, svctask rmemailserver, svctask chemail, and svcinfo lsemailserver.

startemail
 Starts the e-mail notification function.

stopemail
 Disables the e-mail notification function.

testemail
 Sends a test e-mail to a specified user or to all users using the e-mail notification function.


chemail

The chemail command can be used to set or modify contact information for e-mail event notifications. To modify settings, at least one of the parameters must be specified.

Syntax

```

▶▶▶ svctask — — chemail — — — —————▶
                                └─ -reply — reply_email_address ─┘
▶
└─ -contact — contact_name ─┘ —————▶
▶
└─ -primary — primary_telephone_number ─┘ —————▶
  
```

Parameters

-reply *reply_email_address*

(Optional) Specifies the e-mail address to which a reply is sent.

-contact *contact_name*

(Optional) Specifies the name of the person to receive the e-mail.

-primary *primary_telephone_number*

(Optional) Specifies the primary contact telephone number.

-alternate *alternate_telephone_number*

(Optional) Specifies the alternate contact telephone number that is used when you cannot reach the primary contact on the primary phone.

-location *location*

(Optional) Specifies the physical location of the system that is reporting the error. The *location* value must not contain punctuation or any other characters that are not alphanumeric or spaces.

Description

This command sets or modifies contact information that is used by the e-mail event notification facility.

Note: If you are starting the e-mail event notification facility, the **reply**, **contact**, **primary**, and **location** parameters are required. If you are modifying contact information used by the e-mail event notification facility, at least one of the parameters must be specified.

An invocation example

```
svctask chemail -primary 0441234567 -location 'room 256 floor 1 IBM'
```


The resulting output

[No feedback]

chemailserver

The `chemailserver` command modifies the parameters of an existing e-mail server object.

Syntax



Parameters

-name *server_name*

(Optional) Specifies a unique name to assign to the e-mail server object. The name must be a 1- through 15-character string, and cannot start with a hyphen or number. When specifying a server name, `emailserver` is a reserved word.

-ip *ip_address*

(Optional) Specifies the IP address of the e-mail server object. This must be a valid IPv4 or IPv6 address. IPv6 addresses can be zero compressed.

-port *port*

(Optional) Specifies the port number for the e-mail server. This must be a value of 0 - 65535. The default value is 25.

email_server_name | *email_server_id*

(Required) Specifies the name or ID of the server object to be modified.

Description

Use this command to change the settings of an existing e-mail server object. The e-mail server object describes a remote Simple Mail Transfer Protocol (SMTP) e-mail server.

You must specify either the current name or the ID of the object returned at creation time. Use the `svcinfolmailserver` command to obtain this ID.

An invocation example

```
svctask chemailserver -name newserver 0
```

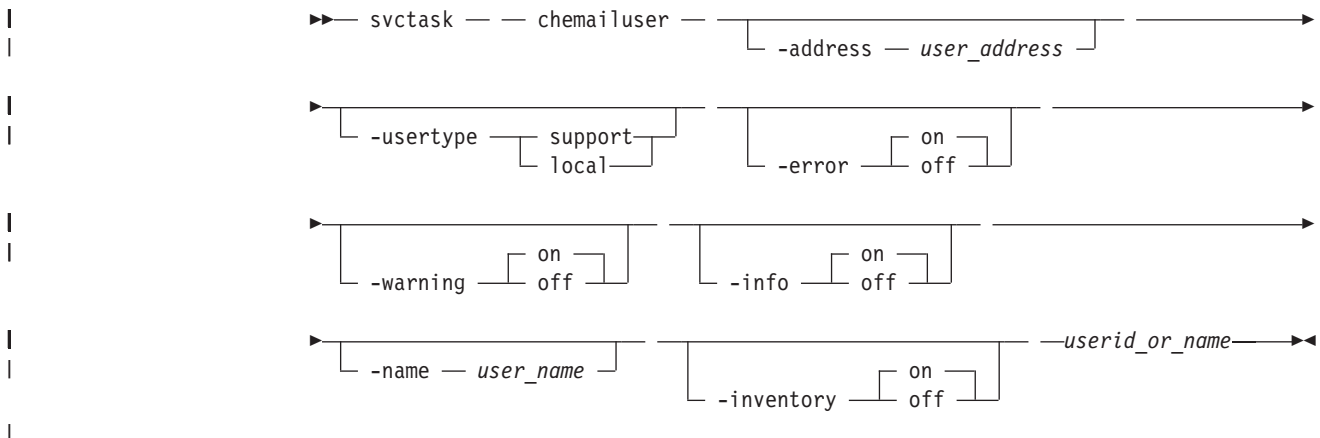
The resulting output

```
none
```

chemailuser

The `chemailuser` command modifies the settings that are defined for an e-mail recipient.

Syntax



Parameters

-address *user_address*

(Optional) Specifies the e-mail address of the person receiving the e-mail or inventory notifications, or both. The *user_address* value must be unique.

-usertype support | local

(Optional) Specifies the type of user, either local or support, based on the following definitions:

support

Address of the support organization that provides vendor support.

local All other addresses.

-error on | off

(Optional) Specifies whether the recipient receives error-type event notifications. Set to on, error-type event notifications are sent to the e-mail recipient. Set to off, error-type event notifications are not sent to the recipient.

-warning on | off

(Optional) Specifies whether the recipient receives warning-type event notifications. Set to on, warning-type event notifications are sent to the e-mail recipient. Set to off, warning-type event notifications are not sent to the recipient.

-info on | off

(Optional) Specifies whether the recipient receives informational event notifications. Set to on, informational event notifications are sent to the e-mail recipient. Set to off, informational event notifications are not sent to the recipient.

-name *user_name*

(Optional) Specifies the user name of the new e-mail event notification recipient. The *user_name* value must be unique, must not contain spaces, and must not contain all numbers. The name **emailuser n** , where n is a number, is reserved and cannot be specified as one of your user names.

-inventory on | off

(Optional) Specifies whether this recipient receives inventory e-mail notifications.

userid_or_name

(Required) Specifies the e-mail recipient for whom you are modifying settings.

Description

This command modifies the settings that are established for an e-mail recipient. Standard rules regarding names apply; therefore, it is not possible to change a name to **emailuser n** , where n is a number.

Note: Before the **usertype** parameter can be set to support, the **-warning** and **-info** flags must be set to off.

An invocation example

The following example modifies e-mail settings for e-mail recipient **manager2008**:

```
svctask chemailuser -usertype local manager2008
```

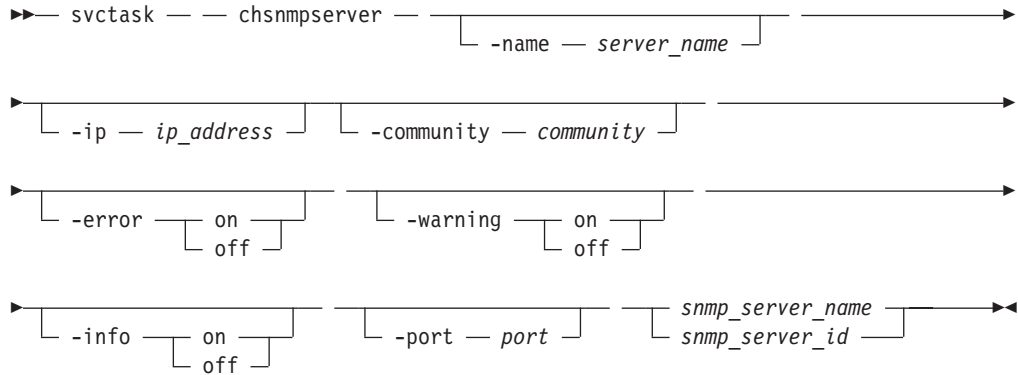
The resulting output

[no feedback]

chsnmpserver

The chsnmpserver command modifies the parameters of an existing SNMP server.

Syntax



Parameters

-name *server_name*

(Optional) Specifies a name to assign to the SNMP server. The name must be unique. When specifying a server name, *snmp* is a reserved word.

-ip *ip_address*

(Optional) Specifies an IP address to assign to the SNMP server. This must be a valid IPv4 or IPv6 address.

-community *community*

(Optional) Specifies the community name for the SNMP server.

-error on | off

(Optional) Specifies whether the server receives error notifications. Set to *on*, error notifications are sent to the SNMP server. Set to *off*, error notifications are not sent to the SNMP server.

-warning on | off

(Optional) Specifies whether the server receives warning notifications. Set to *on*, warning notifications are sent to the SNMP server. Set to *off*, warning notifications are not sent to the SNMP server.

-info on | off

(Optional) Specifies whether the server receives information notifications. Set to *on*, information notifications are sent to the SNMP server. Set to *off*, information notifications are not sent to the SNMP server.

-port *port*

(Optional) Specifies the remote port number for the SNMP server. This must be a value of 1 - 65535.

snmp_server_name | *snmp_server_id*

(Required) Specifies the name or ID of the server to be modified.

Description

Use this command to change the settings of an existing SNMP server. You must specify either the current name of the server or the ID returned at creation time. Use the `svcinfo lssnmpserver` command to obtain this ID.

An invocation example

```
svctask chsnmpserver -name newserver 0
```

The resulting output

```
none
```

chsyslogserver

The chsyslogserver command modifies the parameters of an existing syslog server.

Syntax

```
svctask -- chsyslogserver -- [-name server_name] [-ip ip_address] [-facility facility] [-error on | off] [-warning on | off] [-info 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 fully-expanded format. Servers configured with facility values of 4 - 7 receive syslog messages in concise format.

-error on | off

(Optional) Specifies whether the server receives error notifications. Set to on, error notifications are sent to the syslog server. Set to off, error notifications are not sent to the syslog server.

-warning on | off

(Optional) Specifies whether the server receives warning notifications. Set to on, warning notifications are sent to the syslog server. Set to off, warning notifications are not sent to the syslog server.

-info on | off

(Optional) Specifies whether the server receives information notifications. Set to on, information notifications are sent to the syslog server. Set to off, information notifications are not sent to the syslog server.

syslog_server_name | syslog_server_id

(Required) Specifies the name or ID of the server to be modified.

Description

Use this command to change the settings of an existing syslog server. You must specify either the current name of the server or the ID returned at creation time. Use the `svcinfolssyslogserver` command to obtain this ID.

An invocation example

```
svctask chsyslogserver -facility 5 2
```

The resulting output

```
none
```

mkemailserver

The `mkemailserver` command creates an e-mail server object that describes a remote Simple Mail Transfer Protocol (SMTP) e-mail server.

Syntax

```
svctask -- mkemailserver -- [-name -- server_name] --  
-- -ip -- ip_address -- [-port -- port] --
```

Parameters

-name *server_name*

(Optional) Specifies a unique name to assign to the e-mail server object. The name must be a 1- through 15-character string, and cannot start with a hyphen or number. If a name is not specified, then a system default of `emailservern` is applied, where *n* is the object ID. When specifying a server name, `emailserver` is a reserved word.

-ip *ip_address*

(Required) Specifies the IP address of a remote e-mail server. This must be a valid IPv4 or IPv6 address. IPv6 addresses can be zero compressed.

-port *port*

(Optional) Specifies the port number for the e-mail server. This must be a value of 0 - 65535. The default value is 25.

Description

This command creates an e-mail server object that represents the SMTP server. The SAN Volume Controller uses the e-mail server to send event notification and inventory e-mails to e-mail users. It can transmit any combination of error, warning, and informational notification types.

The SAN Volume Controller supports up to six e-mail servers to provide redundant access to the external e-mail network. The e-mail servers are used in turn until the e-mail is successfully sent from the SAN Volume Controller. The attempt is successful when the SAN Volume Controller gets a positive acknowledgement from an e-mail server that the e-mail has been received by the server.

An invocation example

```
svctask mkemailserver -ip 2.2.2.2 -port 78
```

The resulting output

```
Emailserver id [2] successfully created
```

mkemailuser

The `mkemailuser` command adds a recipient of e-mail event and inventory notifications to the e-mail event notification facility. You can add up to twelve recipients, one recipient at a time.

Syntax

```
svctask -- mkemailuser -- [-name user_name] --
--address user_address -- -usertype support|local --
[-error on|off] [-warning on|off]
[-info on|off] [-inventory on|off]
```

Parameters

-name *user_name*

(Optional) Specifies the name of the person who is the recipient of e-mail event notifications. The *user_name* value must be unique, must not contain spaces, and must not contain only numbers. If you do not specify a user name, the system automatically assigns a user name in the format of `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 e-mail address of the person receiving the e-mail event or inventory notifications, or both. The *user_address* value must be unique.

-usertype **support** | **local**

(Required) Specifies the type of user, either customer or support, based on the following definitions:

support

Address of the support organization that provides vendor support.

local All other addresses.

-error **on** | **off**

(Optional) Specifies whether the recipient receives error-type event notifications. Set to `on`, error-type event notifications are sent to the e-mail recipient. Set to `off`, error-type event notifications are not sent to the recipient. The default value is `on`.

-warning on | off

(Optional) Specifies whether the recipient receives warning-type event notifications. Set to on, warning-type event notifications are sent to the e-mail recipient. Set to off, warning-type event notifications are not sent to the recipient. The default value is on.

-info on | off

(Optional) Specifies whether the recipient receives informational event notifications. Set to on, informational event notifications are sent to the e-mail recipient. Set to off, informational event notifications are not sent to the recipient. The default value is on.

-inventory on | off

(Optional) Specifies whether this recipient receives inventory e-mail notifications. The default value is off.

Description

This command adds e-mail recipients to the e-mail event and inventory notification facility. You can add up to twelve recipients, one recipient at a time. When an e-mail user is added, if a user name is not specified, a default name is allocated by the system. This default name has the form of emailuser1, emailuser2, and so on. E-mail notification starts when you process the startemail command.

Note: Before the **usertype** parameter can be set to support, the **-warning** and **-info** flags must be set to off.

An invocation example

```
svctask mkemailuser -address manager2008@ibm.com -error on -usertype local
```


The resulting output

```
email user, id [2], successfully created
```

mksnmpserver

The mksnmpserver command creates an SNMP server to receive notifications.

Syntax



Parameters

-name server_name

(Optional) Specifies a unique name to assign to the SNMP server. If a name is

not specified, then a system default of `snmpn` is applied, where *n* is the ID of the server. When specifying a server name, `snmp` is a reserved word.

-ip *ip_address*

(Required) Specifies the IP address of the SNMP server. This must be a valid IPv4 or IPv6 address.

-community *community*

(Optional) Specifies the community name for the SNMP server. If you do not specify a community name, then the default name of `public` is used.

-error on | off

(Optional) Specifies whether the server receives error notifications. Set to `on`, error notifications are sent to the SNMP server. Set to `off`, error notifications are not sent to the SNMP server. The default value is `on`.

-warning on | off

(Optional) Specifies whether the server receives warning notifications. Set to `on`, warning notifications are sent to the SNMP server. Set to `off`, warning notifications are not sent to the SNMP server. The default value is `on`.

-info on | off

(Optional) Specifies whether the server receives information notifications. Set to `on`, information notifications are sent to the SNMP server. Set to `off`, information notifications are not sent to the SNMP server. The default value is `on`.

-port *port*

(Optional) Specifies the remote port number for the SNMP server. This must be a value of 1 - 65535. The default value is 162.

Description

This command creates an SNMP server to receive notifications.

SAN Volume Controller supports a maximum of 6 SNMP servers.

An invocation example

```
svctask mksnmpserver -ip 2.2.2.2 -port 78
```

The resulting output


```
SNMP Server id [2] successfully created
```

mksyslogserver

The `mksyslogserver` command creates a syslog server to receive notifications.

Syntax

```
svctask -- mksyslogserver -- [-name -- server_name] --
--ip -- ip_address [-facility -- facility] [-error [on | off]]
```



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 `syslogn` is applied, where *n* is the ID of the server. When specifying a server name, `syslog` is a reserved word.

-ip *ip_address*

(Required) Specifies the IP address of the syslog server. This must be a valid IPv4 or IPv6 address.

-facility *facility*

(Optional) Specifies the facility number used in syslog messages. This number identifies the origin of the message to the receiving server. Servers configured with facility values of 0-3 receive syslog messages in fully-expanded format. Servers configured with facility values of 4-7 receive syslog messages in concise format. The default value is 0.

-error on | off

(Optional) Specifies whether the server receives error notifications. Set to `on`, error notifications are sent to the syslog server. Set to `off`, error notifications are not sent to the syslog server. The default value is `on`.

-warning on | off

(Optional) Specifies whether the server receives warning notifications. Set to `on`, warning notifications are sent to the syslog server. Set to `off`, warning notifications are not sent to the syslog server. The default value is `on`.

-info on | off

(Optional) Specifies whether the server receives information notifications. Set to `on`, information notifications are sent to the syslog server. Set to `off`, information notifications are not sent to the syslog server. The default value is `on`.

Description

This command creates a syslog server to receive notifications. The syslog protocol is a client-server standard for forwarding log messages from a sender to a receiver on an IP network. Syslog can be used to integrate log messages from different types of systems into a central repository.

SAN Volume Controller supports a maximum of 6 syslog servers.

An invocation example

```
svctask mksyslogserver -ip 1.2.3.4
```

The resulting output

```
Syslog Server id [2] successfully created
```

rmemailserver

The `rmemailserver` command deletes the specified e-mail server object.

Syntax

```
svctask -- rmailtoserver -- email_server_name | email_server_id
```

Parameters

email_server_name | *email_server_id*

(Required) Specifies the name or ID of the e-mail server object to be deleted.

Description

Use this command to delete an existing e-mail server object that describes a remote Simple Mail Transfer Protocol (SMTP) e-mail server. You must specify either the current name or the ID of the object returned at creation time. Use the `svcinfo lmailtoserver` command to obtain this ID.

Note: E-mail service stops when the last e-mail server is removed. Use the `svctask startmailto` command to reactivate the e-mail and inventory notification function after at least one e-mail server has been configured.

An invocation example

```
svctask rmailtoserver email4
```

The resulting output

```
none
```

rmmailtouser

The **rmmailtouser** command allows you to remove a previously defined e-mail recipient from your system.

Syntax

```
svctask -- rmailtouser -- userid_or_name
```

Parameters

userid_or_name

(Required) Specifies the user ID or user name of the e-mail recipient to remove.

Description

This command removes an existing e-mail recipient from the system.

An invocation example

The following example removes e-mail recipient **manager2008**:

```
svctask rmailtouser manager2008
```

The resulting output

```
[No feedback]
```

An invocation example

The following example removes e-mail recipient 2:

```
svctask rmemailuser 2
```

The resulting output

[No feedback]

rmsnmpserver

The rmsnmpserver command deletes the specified SNMP server.

Syntax

```
svctask -- rmsnmpserver -- [ snmp_server_name | snmp_server_id ]
```

Parameters

snmp_server_name | *snmp_server_id*

(Required) Specifies the name or ID of the SNMP server to be deleted.

Description

Use this command to delete an existing SNMP server. You must specify either the current name of the server or the ID returned at creation time. Use the `svcinfolssnmpserver` command to obtain this ID.

An invocation example

```
svctask rmsnmpserver snmp4
```

The resulting output

none

rmsyslogserver

The rmsyslogserver command deletes the specified syslog server.

Syntax

```
svctask -- rmsyslogserver -- [ syslog_server_name | syslog_server_id ]
```

Parameters

syslog_server_name | *syslog_server_id*

(Required) Specifies the name or ID of the syslog server to be deleted.

Description

Use this command to delete an existing syslog server. You must specify either the current name of the server or the ID returned at creation time. Use the `svcinfolssyslogserver` command to obtain this ID.

An invocation example

```
svctask rmsyslogserver 2
```

The resulting output

none

sendinventoryemail

The **sendinventoryemail** command sends an inventory e-mail notification to all e-mail recipients who are enabled to receive inventory e-mail notifications. There are no parameters for this command.

Syntax

▶▶— svctask — — sendinventoryemail —————▶▶

Parameters

There are no parameters for this command.

Description

This command sends an inventory e-mail notification to all e-mail recipients who are enabled to receive inventory e-mail notifications. This command fails if the **startemail** command has not been processed and at least one e-mail recipient using the e-mail event and inventory notification facility has not been set up to receive inventory e-mail notifications. This command also fails if the e-mail infrastructure has not been set up.

An invocation example

In the following example, you send an inventory e-mail notification to all e-mail recipients who are enabled to receive them:

```
svctask sendinventoryemail
```

The resulting output

[No feedback]

setemail (Discontinued)

Attention: The **setemail** command is discontinued. E-mail notification can be configured using the following commands: **svctask mkemailserver**, **svctask chemailserver**, **svctask rmemailserver**, **svctask chemail**, and **svcinfolsemailserver**.

startemail

The **startemail** command activates the e-mail and inventory notification function. There are no parameters for this command.

Syntax

▶▶— svctask — — startemail —————▶▶

Parameters

There are no parameters for this command.

Description

This command enables the e-mail event notification service. No e-mails are sent to users until the **startemail** command has been run and at least one user has been defined to the system.

An invocation example

In the following example, you are starting the e-mail error notification service.

```
svctask startemail
```

The resulting output

```
[No feedback]
```

stopemail

The **stopemail** command stops the e-mail and inventory notification function. There are no parameters for this command.

Syntax

```
svctask -- stopemail
```

Parameters

There are no parameters for this command.

Description

This command stops the e-mail error notification function. No e-mails are sent to users until the **startemail** command is reissued.

An invocation example

In the following example, you have stopped the e-mail and inventory notification function:

```
svctask stopemail
```

The resulting output

```
[No feedback]
```

testemail

The **testemail** command allows you to send an e-mail notification to one user of the e-mail notification function or to all users of the e-mail notification function to ensure that the function is operating correctly.

Syntax

```
svctask -- testemail userid_or_name | -all
```

Parameters

userid_or_name

(Required if you do not specify **-all**) Specifies the user ID or user name of the e-mail recipient that you want to send a test e-mail to. You cannot use this parameter with the **-all** parameter. The *userid_or_name* value must not contain spaces.

-all

(Required if you do not specify *userid_or_name*) Sends a test e-mail to all e-mail users configured to receive notification of events of any notification type. No attempt is made to send the test e-mail to an e-mail user who does not have any notification setting set to *on*.

Description

This command sends test e-mails to the specified e-mail users. The e-mail recipient expects to receive the test e-mail within a specified service time. If the e-mail is not received within the expected time period, the recipient must contact the administrator to ensure that the e-mail settings for the user are correct. If there is still a problem, you must contact the IBM Support Center.

The e-mail recipient uses the test e-mail to check that the SMTP name, the IP address, the SMTP port, and the user address are valid.

An invocation example

The following example sends a test e-mail to the user ID **manager2008**:

```
svctask testemail manager2008
```

The resulting output

[No feedback]

Chapter 8. Cluster commands

Cluster commands are used to monitor and modify clusters.

A cluster consists of up to four pairs of nodes that provide a single configuration and service interface. There are a number of cluster commands available for various tasks.

addnode

Adds a new (candidate) node to an existing cluster.

cfgportip

Assigns an IP address to each node ethernet port for iSCSI I/O.

chcluster

Modifies the attributes of an existing cluster.

chclusterip

Modifies the IP configuration parameters for the cluster.

chiogrp

Modifies the name of an I/O group or the amount of memory that is available for Copy Services or VDisk mirroring.

chnode

Changes the name assigned to a node.

cleardumps

Cleans the various dump directories on a specified node.

cpdumps

Copies dump files from a nonconfiguration node onto the configuration node.

detectmdisk

Rescans the fibre-channel network for new managed disks.

ping Used to diagnose IP configuration problems by checking whether the specified IP address is accessible from the configuration node.

rmnode

Deletes a node from the cluster.

rmportip

Removes an iSCSI IP address from a node ethernet port.

setclustertime

Sets the time for the cluster.

setpwdreset

Changes the status of the password-reset feature for the display panel.

settimezone

Sets the time zone for the cluster.

startstats

Starts the collection of per-node statistics for virtual disks (VDisks), managed disks (MDisks), and nodes.

stopcluster

Shuts down a single node or the entire cluster.

stopstats

Stops the collection of per-node statistics for virtual disks (VDisks), managed disks (MDisks), and nodes.

addnode

You can use the addnode command to add a new (candidate) node to an existing cluster. You can enter this command any time after a cluster has been created. If you are adding a new node to a cluster, you must ensure that the model type of the new node is supported by the SAN Volume Controller software version of the cluster. If the model type is not supported by the cluster software, you must upgrade the cluster to a software version that supports the model type of the new node.

Syntax

```
svctask -- addnode -- [-panelname -- panel_name] -- [-wwnodename -- wwnn_arg] -- [-name -- new_name_arg] -- [-iogrp -- [iogroup_name | iogroup_id]]
```

Parameters

-panelname *panel_name*

(Required if you do not specify the **-wwnodename** parameter) Specifies the node that you want to add to a cluster by the name that is displayed on the display panel. You cannot use this parameter with the **-wwnodename** parameter.

-wwnodename *wwnn_arg*

(Required if you do not specify the **-panelname** parameter) Specifies the node that you want to add to the cluster by the worldwide node name (WWNN). You cannot use this parameter with the **-panelname** parameter.

-name *new_name_arg*

(Optional) Specifies a name for the node that you want to add to the cluster.

-iogrp *iogroup_name* | *iogroup_id*

(Required) Specifies the I/O group to which you want to add this node.

Description

This command adds a new node to the cluster. You can obtain a list of candidate nodes (those that are not already assigned to a cluster) by typing `svcinfo lsnodecandidate`.

Before you add a node to the cluster, you must check to see if any of the following conditions are true. If the following conditions exist, failure to follow the procedures that are documented here might result in the corruption of all data that is managed by the cluster.

- Is the new node being used to replace a failed node in the cluster?
- Does the node being added to the cluster use physical node hardware that has been used as a node in another cluster, and are both clusters recognized by the same hosts?

If any of the previous conditions are true, you must take the following actions:

1. Add the node to the same I/O group that it was previously in. You can use the command-line interface command `svcinfo lsnode` or the SAN Volume Controller Console to determine the WWNN of the cluster nodes.
2. Shut down all of the hosts that use the cluster, before you add the node back into the cluster.
3. Add the node back to the cluster before the hosts are restarted. If the I/O group information is unavailable or it is inconvenient to shut down and restart all of the hosts that use the cluster, you can do the following:
 - a. On all of the hosts that are connected to the cluster, unconfigure the fibre-channel adapter device driver, the disk device driver, and the multipathing driver before you add the node to the cluster.
 - b. Add the node to the cluster and then reconfigure the fibre-channel adapter device driver, the disk device driver, and multipathing driver.

If you are adding a new node to a cluster, take the following actions:

1. Ensure that the model type of the new node is supported by the SAN Volume Controller software version of the cluster. If the model type is not supported by the cluster software, you must upgrade the cluster to a software version that supports the model type of the new node.
2. Record the node serial number, the WWNN, all 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 cluster.

Other considerations when you add a node to a cluster:

When you add a node to the cluster using the `svctask addnode` command or the cluster GUI, you must confirm whether the node has previously been a member of the cluster. If it has, follow one of these two procedures:

- Add the node to the same I/O group that it was previously in. You can determine the WWNN of the nodes in the cluster using the `svcinfo lsnode` command.
- If you cannot determine the WWNN of the nodes in the cluster, call the support team to add the node back into the cluster without corrupting the data.

When a node is added to a cluster, it displays a state of adding. It can take as long as 30 minutes for the node to be added to the cluster, particularly if the software version of the node has changed.

Attention: If the node remains in the adding state for more than 30 minutes, contact your support representative to assist you in resolving this issue.

Optionally, you can assign a name to the new node. You can use this name in subsequent commands to refer to the node, instead of using the node ID. If you assign a label, this label is displayed as the node name from then on. If you do not assign a label, the default label is `nodeX`, where `X` is the node ID.

An invocation example

```
svctask addnode -wwnodename 5005076801e08b -iogrp io_grp0
```

The resulting output

```
Node, id [6], successfully added
```

cfgportip

The `cfgportip` command assigns an IP address to each node ethernet port for iSCSI I/O.

Syntax

```
svctask -- cfgportip -- -node [ node_name | node_id ]
      -ip [ ipv4addr | -ip_6 [ ipv6addr ] ] -gw [ ipv4gw | -gw_6 [ ipv6gw ] ]
      -mask [ subnet_mask | -prefix_6 [ prefix ] ] -failover [ ] --port_id [ ]
```

Parameters

-node *node_name* | *node_id*

(Required) Specifies which node has the ethernet port that the IP address is being assigned to.

-ip *ipv4addr*

(Required if you do not use **ip_6**) Sets the IPv4 address for the ethernet port. You cannot use this parameter with the **ip_6** parameter.

-ip_6 *ipv6addr*

(Required if you do not use **ip**) Sets the IPv6 address for the ethernet port. You cannot use this parameter with the **ip** parameter.

-gw *ipv4addr*

(Required if you do not use **gw_6**) Sets the IPv4 gateway IP address. You cannot use this parameter with the **gw_6** parameter.

-gw_6 *ipv6gw*

(Required if you do not use **gw**) Sets the IPv6 default gateway address for the port. You cannot use this parameter with the **gw** parameter.

-mask *subnet_mask*

(Required if you do not use **prefix_6**) Sets the IPv4 subnet mask. You cannot use this parameter with the **prefix_6** parameter.

-prefix_6 *prefix*

(Required if you do not use **mask**) Sets the IPv6 prefix. You cannot use this parameter with the **mask** parameter.

-failover

(Optional) Specifies that the IP address belongs to the partner node in the I/O group. If the partner node is offline, the address is configured and presented by this node. When another node comes online in the I/O group, the failover address is presented by that node.

port_id

(Required) Specifies which port (1 or 2) to apply changes to.

Description

This command assigns either an IPv4 or IPv6 address to a specified ethernet port of a node. The IP address is used for iSCSI I/O. Use the `svctask chclusterip` command to assign cluster IP addresses.

For an IPv4 address, the **ip**, **mask**, and **gw** parameters are required. All of the IPv4 IP parameters must be specified to assign an IPv4 address to an ethernet port.

For an IPv6 address, the **ip_6**, **prefix_6**, and **gw_6** parameters are required. All of the IPv6 IP parameters must be specified to assign an IPv6 address to an ethernet port.

Use the `svcinfo lspportip` command with the optional **ethernet_port_id** parameter to list the port IP addresses for the specified port.

An invocation example for IPv4

```
svctask cfgportip -node 1 -ip 9.8.7.1 -gw 9.0.0.1 -mask 255.255.255.0 1
```

The resulting output

No feedback

An invocation example for IPv6

```
svctask cfgportip -node 1 -ip_6 3:3:0:4::0 -gw_6 ffe8::0 -prefix_6 64 2
```

The resulting output

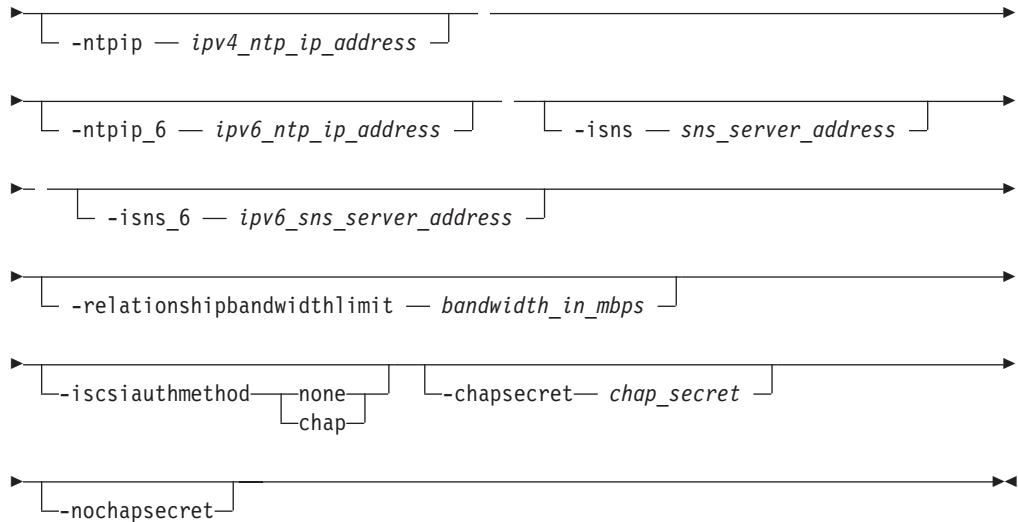
No feedback

chcluster

The `chcluster` command modifies the attributes of an existing cluster. You can enter this command any time after a cluster has been created. All the parameters that are associated with this command are optional. However, you must specify one or more parameters with this command.

Syntax

```
▶▶ svctask -- chcluster -- -- [-name cluster_name]
▶ [-servicepwd password] -- [-speed fabric_speed]
▶ [-alias id_alias] -- [-icatip icat_console_ip_address]
▶ [-invemailinterval interval]
▶ [-gmlinktolerance link_tolerance]
▶ [-gminterdelaysimulation inter_cluster_delay_simulation]
▶ [-gmintradelaysimulation intra_cluster_delay_simulation]
▶ [-icatip_6 ipv6_icat_ip_address]
```



Parameters

-name *cluster_name*
 (Optional) Specifies a new name for the cluster.

Important: The iSCSI Qualified Name (IQN) for each node is generated using the cluster and node names. If you are using the iSCSI protocol, changing either name also changes the IQN of all of the nodes in the cluster and might require reconfiguration of all iSCSI-attached hosts.

-servicepwd *password*
 (Optional) Specifies a new service user password. You can specify this parameter with or without the password. If the parameter is not followed by a password, you are prompted for the password. When you type the password in response to the prompt, the password is not displayed.

Note: Only a user with administrator authority can change the password.

-speed *fabric_speed*
 (Optional) Specifies the speed of the fabric to which this cluster is attached. Valid values are 1 or 2 (GB).

Attention: Changing the speed on a running cluster breaks I/O service to the attached hosts. Before changing the fabric speed, stop I/O from active hosts and force these hosts to flush any cached data by unmounting volumes (for UNIX® host types) or by removing drive letters (for Windows host types). Some hosts might need to be rebooted to detect the new fabric speed.

The fabric speed setting applies only to the SAN Volume Controller model 2145-8F2 in a cluster. The SAN Volume Controller models 2145-8A4, 2145-CF8, 2145-8G4, and 2145-8F4 automatically negotiate the fabric speed on a per-port basis.

-alias *id_alias*
 (Optional) Specifies an alternate name that does not change the basic ID for the cluster, but does influence the VDisk_UID of every **vdiskhostmap**, both existing and new. These objects appear to have been created for a cluster whose ID matches the alias. Therefore, changing the cluster alias causes loss of host VDisk access, until each host rescans for VDIsks presented by the cluster.

- icatip** *icat_console_ip_address*
 (Optional) Specifies the new IP address that is used by the cluster. 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.
- invemailinterval** *interval*
 (Optional) Specifies the interval at which inventory e-mails are sent to the designated e-mail recipients. The interval range is 0 to 15. The interval is measured in days. Setting the value to 0 turns the inventory e-mail notification function off.
- gmlinktolerance** *link_tolerance*
 (Optional) Specifies the length of time, in seconds, for which an inadequate intercluster link is tolerated for a Global Mirror operation. The parameter accepts values from 60 to 400 seconds in steps of 10 seconds. The default is 300 seconds. You can disable the link tolerance by entering a value of zero (0) for this parameter.
- gminterdelaysimulation** *inter_cluster_delay_simulation*
 (Optional) Specifies the intercluster delay simulation, which simulates the Global Mirror round trip delay between two clusters, in milliseconds. The default is 0; the valid range is 0 to 100 milliseconds.
- gmintradelaysimulation** *intra_cluster_delay_simulation*
 (Optional) Specifies the intracluster delay simulation, which simulates the Global Mirror round trip delay in milliseconds. The default is 0; the valid range is 0 to 100 milliseconds.
- icatip_6** *icat_console_ipv6_address*
 (Optional) Specifies the new IPv6 address that is used by the cluster. If you specify this parameter, it overwrites any existing **-icatip** address. The format of the IPv6 address must be one of the following:
- 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 cluster to immediately start using that NTP server as its time source. To stop using the NTP server as a time source, invoke the **-ntpip** parameter with a zero address, as follows:
 svctask chcluster -ntpip 0.0.0.0
- ntpip_6** *ipv6_ntp_ip_address*
 (Optional) Specifies the IPv6 address for the NTP server. Configuring an NTP server address causes the cluster to immediately start using that NTP server as its time source. To stop using the NTP server as a time source, invoke the **-ntpip_6** parameter with a zero address, as follows:
 svctask chcluster -ntpip_6 0::0
- isns** *sns_server_address*
 (Optional) Specifies the IPv4 address for the iSCSI storage name service (SNS).

| **-isns_6** *ipv6_sns_server_address*

| (Optional) Specifies the IPv6 address for the iSCSI SNS.

| **-relationshipbandwidthlimit** *bandwidth_in_mbps*

| (Optional) Specifies the new background copy bandwidth in megabytes per
| second (MBps), from 1 - 1000. The default is 25 MBps. This parameter operates
| cluster-wide and defines the maximum background copy bandwidth that any
| relationship can adopt. The existing background copy bandwidth settings
| defined on a partnership continue to operate, with the lower of the partnership
| and VDisk rates attempted.

| **Note:** Do not set this value higher than the default without establishing that
| the higher bandwidth can be sustained.

| **-iscsiauthmethod** **none** | **chap**

| (Optional) Sets the authentication method for the iSCSI communications of the
| cluster. The **iscsiauthmethod** value can be none or chap.

| **-chapsecret** *chap_secret*

| (Optional) Sets the Challenge Handshake Authentication Protocol (CHAP)
| secret to be used to authenticate the cluster via iSCSI. This parameter is
| required if the **iscsiauthmethod** **chap** parameter is specified. The specified
| CHAP secret cannot begin or end with a space.

| **-nochapsecret**

| (Optional) Clears any previously set CHAP secret for iSCSI authentication. This
| parameter is not allowed if the **chapsecret** parameter is specified.

Description

This command modifies specific features of a cluster. Multiple features can be changed by issuing a single command.

Using the **-ntpip** or **-ntpip_6** parameter allows the cluster to use an NTP server as an outside time source. The cluster adjusts the system clock of the configuration node according to time values from the NTP server. The clocks of the other nodes are updated from the configuration node's clock. In the NTP mode, the `svctask setclustertime` command is disabled.

All command parameters are optional; however, you must specify at least one parameter.

| **Modifying a password:** To change the service user password, issue the **svctask**
| **chcluster -servicepwd** *password* command.

| **Note:** If you do not want the password to display as you enter it on the command
| line, omit the new password. The command line tool then prompts you to
| enter and confirm the password without the password being displayed.

| Use the `svctask chclusterip` command to modify the cluster IP address and service
| IP address.

An invocation example

| `svctask chcluster -ntpip 9.20.165.16 -relationshipbandwidthlimit 40`

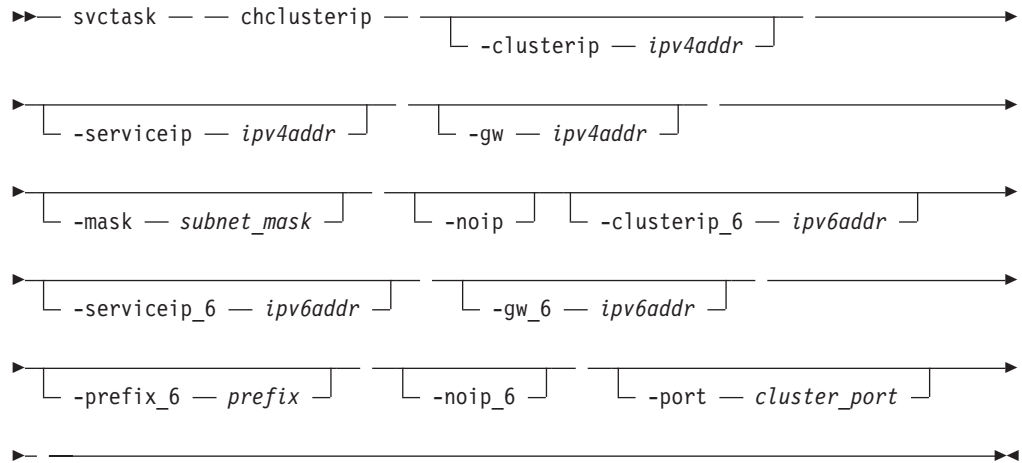
The resulting output

No feedback

chclusterip

The `chclusterip` command modifies the IP configuration parameters for the cluster.

Syntax



Parameters

-clusterip *ipv4addr*

(Optional) Changes the IPv4 cluster IP address. When you specify a new IP address for a cluster, the existing communication with the cluster is broken.

-serviceip *ipv4addr*

(Optional) Changes the IPv4 service IP address. This is the address that is used to manage any node that has to be started in service mode. A node that is operating in service mode does not operate as a member of the cluster.

-gw *ipv4addr*

(Optional) Changes the IPv4 default gateway IP address of the cluster.

-mask *subnet_mask*

(Optional) Changes the IPv4 subnet mask of the cluster.

-noip

(Optional) Unconfigures the IPv4 stack on the specified port, or both ports if none is specified.

-clusterip_6 *ipv6addr*

(Optional) Sets the IPv6 cluster address for the port.

-serviceip_6 *ipv6addr*

(Optional) Sets the IPv6 service address for the port.

-gw_6 *ipv6addr*

(Optional) Sets the IPv6 default gateway address for the port.

-prefix_6 *prefix*

(Optional) Sets the IPv6 prefix.

-noip_6

(Optional) Unconfigures the IPv6 stack on the specified port, or both ports if none is specified.

-port *cluster_port*

(Optional) Specifies which port (1 or 2) to apply changes to. This parameter is required unless the **noip** or **noip_6** parameter is used.

Description

This command modifies IP configuration parameters for the cluster. The first time you configure a second port, all IP information is required. Port 1 on the cluster must always have one stack fully configured.

There are two active cluster ports on the configuration node. There are also two active service ports on any node that is in service mode.

If the cluster IP address is changed, the open command-line shell closes during the processing of the command. You must reconnect to the new IP address if connected through that port.

The **noip** and **noip_6** parameters can be specified together only if the **port** is also specified. The **noip** and **noip_6** parameters cannot be specified with any parameters other than **port**. Port 1 must have an IPv4 or IPv6 cluster address. The configuration of port 2 is optional.

Service IP addresses for all ports and stacks are initialized to DHCP. A service IP address is always configured.

Modifying an IP address: List the IP address of the cluster by issuing the `svcinfolcluster` command. Modify the IP address by issuing the `svctask chclusterip` command. You can either specify a static IP address or have the system assign a dynamic IP address.

Table 4 provides IP address formats that are supported.

Table 4. *ip_address_list* formats

IP type	<i>ip_address_list</i> format
IPv4 (no port set, SVC uses default)	1.2.3.4
IPv4 with specific port	1.2.3.4:22
Full IPv6, default port	1234:1234:abcd:0123:0000:0000:7689:6576
Full IPv6, default port, leading zeros suppressed	1234:1234:abcd:123:0:0:7689:6576
Full IPv6 with port	[1234:1234:abcd:0123:0000:0000:7689:6576]:23
Zero-compressed IPv6, default port	1234:1234:abcd:123::7689:6576
Zero-compressed IPv6 with port	[1234:1234:abcd:123::7689:6576]:23

An invocation example

```
svctask chclusterip -clusterip 9.20.136.5 -gw 9.20.136.1 -mask 255.255.255.0 -port 1
```


The resulting output

No feedback

chiogrp

The `chiogrp` command modifies the name of an I/O group, or the amount of memory that is available for Copy Services or VDisk mirroring operations.

Syntax



Parameters

-name *new_name*

(Optional) Specifies the name to assign to the I/O group. The **-name** parameter cannot be specified with the **-feature**, **-size**, or **-kb** parameters.

-feature **flash** | **remote** | **mirror**

(Optional) Specifies the feature to modify the amount of memory for: Copy Services or VDisk mirroring. You must specify this parameter with the **-size** parameter. You cannot specify this parameter with the **-name** parameter.

Note: Specifying **remote** changes the amount of memory that is available for Metro Mirror or Global Mirror processing. Any VDisk that is in a Metro Mirror or Global Mirror relationship uses memory in its I/O group, including master and auxiliary VDIsks, and VDIsks that are in inter-cluster or intra-cluster relationships.

-size *memory_size*

(Optional) Specifies the amount of memory that is available for the specified Copy Services or VDisk mirroring function. Valid input is **0** or any integer. The default unit of measurement for this parameter is megabytes (MB); you can use the kilobytes **-kb** parameter to override the default. You must specify this parameter with the **-feature** parameter. You cannot specify this parameter with the **-name** parameter.

-kb

(Optional) Changes the units for the **-size** parameter from megabytes (MB) to kilobytes (KB). If you specify this parameter, the **-size** *memory_size* value must be any number divisible by 4. You must specify this parameter with the **-feature** and **-size** parameters. You cannot specify this parameter with the **-name** parameter.

io_group_id | *io_group_name*

(Required) Specifies the I/O group to modify. You can modify an I/O group by using the **-name** or the **-feature** parameter.

Description

The **chiogrp** command modifies the name of an I/O group or the amount of memory that is available for Copy Services or VDisk mirroring. You can assign a

name to an I/O group or change the name of a specified I/O group. You can change the amount of memory that is available for Copy Services or VDisk mirroring operations by specifying the **-feature flash | remote | mirror** parameter, and a memory size. For VDisk mirroring and Copy Services (FlashCopy, Metro Mirror, and Global Mirror), memory is traded against memory that is available to the cache. The amount of memory can be decreased or increased. Consider the following memory sizes when you use this command:

- The default memory size for FlashCopy is 20 MB.
- The default memory size for Metro Mirror and Global Mirror is 20 MB.
- The default memory size for mirrored VDIs is 0 MB.
- The maximum memory size that can be specified for FlashCopy is 512 MB.
- The maximum memory size that can be specified for Metro Mirror and Global Mirror is 512 MB.
- The maximum memory size that can be specified for mirrored VDIs is 512 MB.
- The maximum combined memory size across all features is 512 MB.

Table 5 demonstrates the amount of memory required for VDisk mirroring and Copy Services. Each 1 MB of memory provides the following VDisk capacities and grain sizes:

Table 5. Memory required for VDisk Mirroring and Copy Services

Feature	Grain size	1 MB of memory provides the following VDisk capacity for the specified I/O group
Metro Mirror and Global Mirror	256 KB	2 TB of total Metro Mirror and Global Mirror VDisk capacity
FlashCopy	256 KB	2 TB of total FlashCopy source VDisk capacity
FlashCopy	64 KB	512 GB of total FlashCopy source VDisk capacity
Incremental FlashCopy	256 KB	1 TB of total Incremental FlashCopy source VDisk capacity
Incremental FlashCopy	64 KB	256 GB of total Incremental FlashCopy source VDisk capacity
VDisk mirroring	256 KB	2 TB of mirrored VDIs

For multiple FlashCopy targets, you must consider the number of mappings. For example, for a mapping with a 256 KB grain size, 8 KB of memory allows one mapping between a 16 GB source VDisk and a 16 GB target VDisk. Alternatively, for a mapping with a 256 KB grain size, 8 KB of memory allows two mappings between one 8 GB source VDisk and two 8 GB target VDIs.

When you create a FlashCopy mapping, if you specify an I/O group other than the I/O group of the source VDisk, the memory accounting goes towards the specified I/O group, not towards the I/O group of the source VDisk.

An invocation example

```
svctask chiogrp -name testiogrpone io_grp0
```

The resulting output

No feedback

An invocation example for changing the amount of FlashCopy memory in io_grp0 to 30 MB

```
svctask chiogrp -feature flash -size 30 io_grp0
```

The resulting output

No feedback

chnode

You can use the `chnode` command to change the name that is assigned to a node and other options. The name can then be used when running subsequent commands. All the parameters that are associated with this command are optional. However, you must specify one or more parameters with this command.

Syntax

```
svctask -- chnode [ -iscsialias alias ] [ -noiscsialias ] [ -failover ] [ -name new_node_name ] [ node_name | node_id ]
```

Parameters

-iscsialias *alias*

(Optional) Specifies the iSCSI name of the node. The maximum length is 79 characters.

-noiscsialias

(Optional) Clears any previously set iSCSI name for this node. This parameter cannot be specified with the **iscsialias** parameter.

-failover

(Optional) Specifies that the name or iSCSI alias being set is the name or alias of the partner node in the I/O group. When there is no partner node, the values set are applied to the partner node when it is added to the cluster. If this parameter is used when there is a partner node, the name or alias of that node changes.

-name *new_node_name*

(Optional) Specifies the name to assign to the node.

Important: The iSCSI Qualified Name (IQN) for each node is generated using the cluster and node names. If you are using the iSCSI protocol, changing either name also changes the IQN of all of the nodes in the cluster and might require reconfiguration of all iSCSI-attached hosts.

node_name | *node_id*

(Required) Specifies the node to be modified. The variable that follows the parameter is either:

- The node name that you assigned when you added the node to the cluster.

- The node ID that is assigned to the node (not the worldwide node name).

Description

If the **failover** parameter is not specified, this command changes the name or iSCSI alias of the node. The name can then be used to identify the node in subsequent commands.

The **failover** parameter is used to specify values that are normally applied to the partner node in the I/O group. When the partner node is offline, the iSCSI alias and IQN are assigned to the remaining node in the I/O Group. The iSCSI host data access is then preserved. If the partner node is offline when these parameters are set, the node they are set on handles iSCSI I/O requests to the iSCSI alias specified, or the IQN that is created using the node name. If the partner node in the I/O group is online when these parameters are set, the partner node handles iSCSI requests to the iSCSI alias specified, and its node name and IQN change.

An invocation example

```
svctask chnode -name testnodeone nodeone
```

The resulting output

No feedback

cleardumps

The **cleardumps** command cleans the various dump directories on a specified node.

Syntax

```
svctask --cleardumps -- -prefix -- directory_or_file_filter
```

Parameters

-prefix *directory_or_file_filter*

(Required) Specifies the directory, files, or both to be cleaned. If a directory is specified, with no file filter, all relevant dump or log files in that directory are cleaned. You can use the following directory arguments (filters):

- **/dumps** (cleans all files in all subdirectories)
- **/dumps/cimom**
- **/dumps/configs**
- **/dumps/elogs**
- **/dumps/feature**
- **/dumps/iostats**
- **/dumps/iotrace**
- **/dumps/mdisk**
- **/home/admin/upgrade**

In addition to the directory, you can specify a filter file. For example, if you specify `/dumps/elogs/*.txt`, all files in the `/dumps/elogs` directory that end in `.txt` are cleaned.

Note: The following rules apply to the use of wildcards with the SAN Volume Controller CLI:

- The wildcard character is an asterisk (*).
- The command can contain a maximum of one wildcard.
- With a wildcard, you must use double quotation marks (" ") around the filter entry, such as in the following entry:

```
>svctask clear.dumps -prefix "/dumps/elogs/*.txt"
```

node_id | *node_name*

(Optional) Specifies the node to be cleaned. The variable that follows the parameter is either:

- The node name, that is, the label that you assigned when you added the node to the cluster
- The node ID that is assigned to the node (not the worldwide node name).

Description

This command deletes all the files that match the `directory/file_filter` argument on the specified node. If no node is specified, the configuration node is cleaned.

You can clean all the dumps directories by specifying `/dumps` as the directory variable.

You can clean all the files in a single directory by specifying one of the directory variables.

You can list the contents of these directories on the given node by using the `svcinfo lsxxxxdumps` commands.

You can use this command to clean specific files in a given directory by specifying a directory or file name. You can use the wildcard character as part of the file name.

Note: To preserve the configuration and trace files, any files that match the following wildcard patterns are not cleaned:

- `*svc.config*`
- `*.trc`
- `*.trc.old`

An invocation example

```
svctask clear.dumps -prefix /dumps/configs
```

The resulting output


No feedback

cpdumps

The `cpdumps` command copies dump files from a nonconfiguration node onto the configuration node.

Note: In the rare event that the /dumps directory on the configuration node is full, the copy action ends when the directory is full and provides no indicator of a failure. Therefore, clear the /dumps directory after migrating data from the configuration node.

Syntax



Parameters

-prefix *directory* | *file_filter*

(Required) Specifies the directory, or files, or both to be retrieved. If a directory is specified with no file filter, all relevant dump or log files in that directory are retrieved. You can use the following directory arguments (filters):

- /dumps (retrieves all files in all subdirectories)
- /dumps/audit
- /dumps/cimom
- /dumps/configs
- /dumps/elogs
- /dumps/feature
- /dumps/iostats
- /dumps/iotrace
- /dumps/mdisk
- /home/admin/upgrade

In addition to the directory, you can specify a file filter. For example, if you specified /dumps/elogs/*.txt, all files in the /dumps/elogs directory that end in .txt are copied.

Note: The following rules apply to the use of wildcards with the SAN Volume Controller CLI:

- The wildcard character is an asterisk (*).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, you must surround the filter entry with double quotation marks (""), as follows:

```
>svctask clearcpdumps -prefix "/dumps/elogs/*.txt"
```

node_id | *node_name*

(Required) Specifies the node from which to retrieve the dumps. The variable that follows the parameter can be one of the following:

- The node name, or label that you assigned when you added the node to the cluster
- The node ID that is assigned to the node (not the worldwide node name).

If the node specified is the current configuration node, no file is copied.

Description

This command copies any dumps that match the directory or file criteria from the given node to the current configuration node.

You can retrieve dumps that were saved to an old configuration node. During failover processing from the old configuration node to another node, the dumps that were on the old configuration node are not automatically copied. Because access from the CLI is only provided to the configuration node, cluster files can only be copied from the configuration node. This command enables you to retrieve files and place them on the configuration node so that you can then copy them.

You can view the contents of the directories by using the `svcinfo lsxxxxdumps` commands.

An invocation example

```
svctask cpdumps -prefix /dumps/configs nodeone
```

The resulting output

No feedback

detectmdisk

The **detectmdisk** command allows you to manually rescan the fibre-channel network for any new managed disks (MDisks) that might have been added and to rebalance MDisk access across all available controller device ports.

Syntax

```
▶— svctask — — detectmdisk —————▶
```

Description

This command causes the cluster to rescan the fibre-channel network. The rescan discovers any new MDisks that have been added to the cluster and rebalances MDisk access across the available controller device ports. This command also detects any loss of controller port availability, and updates the SAN Volume Controller configuration to reflect any changes.

Note: Although it might appear that the **detectmdisk** command has completed, some extra time might be required for it to run. The **detectmdisk** is asynchronous and returns a prompt while the command continues to run in the background. You can use the **lsdiscoverystatus** command to list the discovery status.

In general, the cluster automatically detects disks when they appear on the network. However, some fibre-channel controllers do not send the required SCSI primitives that are necessary to automatically discover the new disks.

If you have attached new storage and the cluster has not detected it, you might need to run this command before the cluster detects the new disks.

When back-end controllers are added to the fibre-channel SAN and are included in the same switch zone as a cluster, the cluster automatically discovers the back-end

controller and determines what storage is presented to it. The SCSI LUs that are presented by the back-end controller are displayed as unmanaged MDisks. However, if the configuration of the back-end controller is modified after this has occurred, the cluster might be unaware of these configuration changes. Run this command to rescan the fibre-channel network and update the list of unmanaged MDisks.

Note: The automatic discovery that is performed by the cluster does not write to an unmanaged MDisk. Only when you add an MDisk to an MDisk group, or use an MDisk to create an image mode virtual disk, is the storage actually used.

To identify the available MDisks, issue the **svctask detectmdisk** command to scan the fibre-channel network for any MDisks. When the detection is complete, issue the **svcinfo lsmdiskcandidate** command to show the unmanaged MDisks; these MDisks have not been assigned to an MDisk group. Alternatively, you can issue the **svcinfo lsmdisk** command to view all of the MDisks.

If disk controller ports have been removed as part of a reconfiguration, the SAN Volume Controller detects this change and reports the following error because it cannot distinguish an intentional reconfiguration from a port failure:

```
1630 Number of device logins reduced
```

If the error persists and redundancy has been compromised, the following more serious error is reported:

```
1627 Insufficient redundancy in disk controller connectivity
```

You must issue the **svctask detectmdisk** command to force SAN Volume Controller to update its configuration and accept the changes to the controller ports.

Note: Only issue the **svctask detectmdisk** command when all of the disk controller ports are working and correctly configured in the controller and the SAN zoning. Failure to do this could result in errors not being reported.

An invocation example

```
svctask detectmdisk
```

The resulting output

```
No feedback
```

ping

The ping command can be used to diagnose IP configuration problems by checking whether the specified IP address is accessible from the configuration node.

Syntax

```
➤➤ svctask -- ping [ ipv4_address | ipv6_address ]
```

Parameters

ipv4_address | *ipv6_address*

(Required) Specifies the cluster IP address.

Description

This command checks whether the specified IP address is accessible from the configuration node. The ping takes place only from the configuration node. It can be useful for diagnosing problems where the configuration node cannot be reached from a specific management server.

An invocation example

```
svctask ping 9.20.136.11
```

The resulting output

```
PING 9.20.136.11 (9.20.136.11) 56(84) bytes of data.  
64 bytes from 9.20.136.11: icmp_seq=1 ttl=249 time=0.690 ms  
64 bytes from 9.20.136.11: icmp_seq=2 ttl=249 time=0.382 ms  
64 bytes from 9.20.136.11: icmp_seq=3 ttl=249 time=0.311 ms  
--- 9.20.136.11 ping statistics ---  
3 packets transmitted, 3 received, 0% packet loss, time 2001ms  
rtt min/avg/max/mdev = 0.311/0.461/0.690/0.164 ms
```

rmnode

The `rmnode` command deletes a node from the cluster. You can enter this command any time after a cluster has been created.

Syntax

```
svctask -- rmnode [-force] [node_name | node_id]
```

Parameters

`-force`

(Optional) Overrides the checks that this command runs. The parameter overrides the following two checks:

- If the command results in VDIs going offline, the command fails unless the **force** parameter is used.
- If the command results in a loss of data because there is unwritten data in the write cache that is contained only within the node to be removed, the command fails unless the **force** parameter is used.

If you use the **force** parameter as a result of an error about VDIs going offline, you force the node removal and run the risk of losing data from the write cache. The **force** parameter should always be used with caution.

`node_name | node_id`

Specifies the node to be deleted. The value for this parameter can be one of the following:

- The node name that you assigned when you added the node to the cluster
- The node ID that is assigned to the node (not the worldwide node name).

Description

This command removes a node from the cluster. This makes the node a candidate to be added back into this cluster or into another cluster. After the node is deleted, the other node in the I/O group enters write-through mode until another node is added back into the I/O group.

By default, the **rmnode** command flushes the cache on the specified node before the node is taken offline. In some circumstances, such as when the system is already degraded (for example, when both nodes in the I/O group are online and the virtual disks within the I/O group are degraded), the system ensures that data loss does not occur as a result of deleting the only node with the cache data.

The cache is flushed before the node is deleted to prevent data loss if a failure occurs on the other node in the I/O group.

To take the specified node offline immediately without flushing the cache or ensuring data loss does not occur, run the **rmnode** command with the **-force** parameter.

Prerequisites:

Before you issue the **rmnode** command, perform the following tasks and read the following Attention notices to avoid losing access to data:

1. Determine which virtual disks (VDisks) are still assigned to this I/O group by issuing the following command. The command requests a filtered view of the VDisks, where the filter attribute is the I/O group.

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

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

Note: Any VDisks that are assigned to the I/O group that this node belongs to are assigned to the other node in the I/O group; the preferred node is changed. You cannot change this setting back.

2. Determine the hosts that the VDisks are mapped to by issuing the **svcinfo lsvdiskhostmap** command.
3. Determine if any of the VDisks that are assigned to this I/O group contain data that you need to access:
 - If you *do not* want to maintain access to these VDisks, go to step 5.
 - If you *do* want to maintain access to some or all of the VDisks, back up the data or migrate the data to a different (online) I/O group.
4. Determine if you need to turn the power off to the node:
 - If this is the last node in the cluster, you do not need to turn the power off to the node. Go to step 5.
 - If this is *not* the last node in the cluster, turn the power off to the node that you intend to remove. This step ensures that the Subsystem Device Driver (SDD) does not rediscover the paths that are manually removed before you issue the delete node request.
5. Update the SDD configuration for each virtual path (vpath) that is presented by the VDisks that you intend to remove. Updating the SDD configuration removes the vpaths from the VDisks. Failure to update the configuration can result in data corruption. See the *Multipath Subsystem Device Driver: User's Guide* for details about how to dynamically reconfigure SDD for the given host operating system.
6. Quiesce all I/O operations that are destined for the node that you are deleting. Failure to quiesce the operations can result in failed I/O operations being reported to your host operating systems.

Attention:

1. Removing the last node in the cluster destroys the cluster. Before you delete the last node in the cluster, ensure that you want to destroy the cluster.
2. If you are removing a single node and the remaining node in the I/O group is online, the data can be exposed to a single point of failure if the remaining node fails.
3. This command might take some time to complete since the cache in the I/O group for that node is flushed before the node is removed. If the **-force** parameter is used, the cache is not flushed and the command completes more quickly. However, if the deleted node is the last node in the I/O group, using the **-force** option results in the write cache for that node being discarded rather than flushed, and data loss can occur. The **-force** option should be used with caution.
4. If both nodes in the I/O group are online and the VDisks are already degraded before deleting the node, redundancy to the VDisks is already degraded and loss of access to data and loss of data might occur if the **-force** option is used.

Notes:

1. If you are removing the configuration node, the **rmnode** command causes the configuration node to move to a different node within the cluster. This process might take a short time: typically less than a minute. The cluster IP address remains unchanged, but any SSH client attached to the configuration node might need to reestablish a connection. The SAN Volume Controller Console reattaches to the new configuration node transparently.
2. If this is the last node in the cluster or if it is currently assigned as the configuration node, all connections to the cluster are lost. The user interface and any open CLI sessions are lost if the last node in the cluster is deleted. A time-out might occur if a command cannot be completed before the node is deleted.

An invocation example

```
svctask rmnode 1
```


The resulting output

```
No feedback
```

rmportip

The `rmportip` command removes an iSCSI IP address from a node ethernet port.

Syntax



Parameters

-failover

(Optional) Specifies that the failover IP address information be removed for the specified port.

-ip_6 *ipv6addr*

(Optional) Specifies that the IPv6 address be removed for the specified port. If this parameter is not used, the IPv4 address is removed by default.

-node *node_name* | *node_id*

(Required) Specifies the node with the ethernet port that the IP address is being removed from.

port *id*

(Required) Specifies which port (1 or 2) to apply changes to.

Description

This command removes an IPv4 or IPv6 address from an ethernet port of a node.

An invocation example for IPv4

```
svctask rmpoortip -node 1 1
```

The resulting output

No feedback

An invocation example for IPv6

```
svctask rmpoortip -node 1 -ip_6 2
```

The resulting output

No feedback

setclustertime

The **setclustertime** command allows you to set the time for the cluster.

Syntax

```
►► svctask — — setclustertime — — -time — time_value —————►►
```

Parameters

-time *time_value*

(Required) Specifies the time to which the cluster must be set. This must be in the following format:

MMDDHHmmYYYY

Description

This command sets the time for the cluster.

An invocation example

```
svctask setclustertime -time 040509142003
```

The resulting output

No feedback

setpwdreset

Use the **setpwdreset** command to view and change the status of the password-reset feature for the display panel.

Syntax

```
svctask setpwdreset [-disable | -enable | -show]
```

Parameters

-disable

Disables the password-reset feature that is available through the front panel menu system.

-enable

Enables the password-reset feature that is available through the front panel menu system.

-show

Displays the status of the password-reset feature, which is either enabled or disabled.

Description

The front panel menu system provides an option to reset the cluster superuser password. This option resets the password to a random string that is displayed on the front panel. You can then use this password to access the system. You can change the password at the next login.

Issue the `svctask setpwdreset` command to view and change the status of the password-reset feature for the display panel. Passwords can consist of the following characters: A - Z, a - z, 0 - 9, and underscore (`_`). Make a careful note of the cluster superuser password, because without it, you cannot access the cluster.

This command allows you access in case the cluster superuser password is forgotten. If you leave this feature enabled, you should ensure adequate physical security to the cluster hardware.

You can view or change the status of this feature.

An invocation example

```
svctask setpwdreset -show
```

The resulting output

```
Password status: [1]
```

This output means that the password or reset feature that is available through the front panel menu system is enabled. If the password status is [0], this feature is disabled.

settimezone

Use the **settimezone** command to set the time zone for the cluster.

Syntax

```
svctask — — settimezone — — -timezone — timezone_arg —————>
```

Parameters

-timezone *timezone_arg*
Specifies the time zone to set for the cluster.

Description

This command sets the time zone for the cluster. Use the **-timezone** parameter to specify the numeric ID of the time zone that you want to set. Issue the **svcinfolistimezones** command to list the time-zones that are available on the cluster. A list of valid time-zones settings are displayed in a list.

The time zone that this command sets will be used when formatting the error log that is produced by issuing the following command:

```
svctask dumperrlog
```

Note: If you have changed the timezone, you must clear the error log dump directory before you can view the error log through the Web application.

Issue the **svcinfolistimezone** command to display the current time-zone settings for the cluster. The cluster ID and its associated time-zone are displayed. Issue the **svctask setclustertime** command to set the time for the cluster.

An invocation example

```
svctask settimezone -timezone 5
```

The resulting output

No feedback

startstats

| Use the startstats command to start the collection of per-node statistics for virtual
| disks (VDisks), managed disks (MDisks), and nodes.

Syntax

```
svctask — — startstats — — -interval — time_in_minutes —————>
```

Parameters

-interval *time_in_minutes*
Specifies the time in minutes. This is the time interval between the gathering of statistics, from 1 to 60 minutes in increments of 1 minute.

Description

The startstats command starts the collection of statistics that are generated on each node. Statistics are collected at the end of each sampling period as specified by the **-interval** parameter. These statistics are written to a file. A new file is created at the end of each sampling period. Separate files are created for MDisks, VDisks and node statistics.

The files generated are written to the /dumps/iostats directory.

A maximum of 16 files are stored in the directory at any one time for each statistics file type, for example:

```
Nm_stats_nodepanelname_date_time  
Nv_stats_nodepanelname_date_time  
Nn_stats_nodepanelname_date_time
```

Statistics files are created for all time intervals. Before the 17th file for each type is created, the oldest file of that type is deleted.

These files can be listed by using the svcinfo lsiostatsdumps command.

The following naming convention is used for these files:

```
stats_type_stats_nodepanelname_date_time
```

Where *stats_type* is Nm for MDisks, Nv for VDisks, and Nn for node statistics. *nodepanelname* is the current configuration node panel name, *date* is in the format of yymmdd, and *time* is in the format of hhmmss.

The following is an example of an MDisk statistics file name:

```
Nm_stats_000229_031123_072426
```

The following is an example of a VDisk statistics file name:

```
Nv_stats_000229_031123_072426
```

The following is an example of a node statistics file name:

```
Nn_stats_000229_031123_072426
```

Statistics are collected for each MDisk and recorded in the *Nm_stats_nodepanelname_date_time* file, including the following statistical information:

- The number of SCSI read and write commands that are processed during the sample period
- The number of blocks of data that are read and written during the sample period
- Per MDisk, cumulative read and write external response times in milliseconds
- Per MDisk, cumulative read and write queued response times

Statistics are collected for each VDisk and recorded in the *Nv_stats_nodepanelname_date_time* file, including the following statistical information:

- The total number of processed SCSI read and write commands
- The total amount of read and written data
- Cumulative read and write response time in milliseconds

- Statistical information about the read/write cache usage
- Global Mirror statistics including latency

Statistics are collected for the node from which the statistics file originated and recorded in the `Nn_stats_nodepanelname_date_time` file, including the following statistical information:

- Usage figure for the node from which the statistic file was obtained
- The amount of data transferred to and received from each port on the node to other devices on the SAN
- Statistical information about communication to other nodes on the fabric

An invocation example

```
svctask startstats -interval 25
```

The resulting output

No feedback

stopcluster

The **stopcluster** command allows you to shut down a single node or the entire cluster in a controlled manner. When you issue this command, you are prompted with a confirmation of intent to process the command.

Syntax

```
svctask -- stopcluster [-force] [-node node_name | node_id]
```

Parameters

-force

(Optional) Specifies that the node that is being shut down is the last online node in a given I/O group. The **force** parameter also overrides the checks that this command runs. The parameter overrides the following two checks:

- If the command results in VDIsks going offline, the command fails unless the **force** parameter is used.
- If the node being shut down is the last online node in the I/O group, the command fails unless the **force** parameter is used.

If you use the **force** parameter as a result of an error about VDIsks going offline, you force the node shutdown even if it is the last online node in the I/O group. The **force** parameter should always be used with caution.

-node node_name | node_id

(Optional) Specifies the node that you want to shut down. You can specify one of the following values:

- The node name, or label that you assigned when you added the node to the cluster
- The node ID that is assigned to the node (not the worldwide node name).

If you specify `-node node_name | node_id`, only the specified node is shut down; otherwise, the entire cluster is shut down.

Description

When you enter this command with no parameters, the entire cluster is shutdown. All data is flushed to disk before the power is removed.

Attention: Ensure that you have stopped all FlashCopy, Metro or Global Mirror, or data migration operations before you attempt to shutdown a node or cluster. You must also ensure that all asynchronous deletion operations have completed prior to a shutdown operation.

When you enter this command with either a node ID or node name, the node in question is shut down. After the command completes, the remaining node in the I/O group enters write-through mode until the power to the node is returned and the node rejoins the cluster.

Attention: If you have to remove all cluster input power for more than a few minutes, you must shut down the cluster before you remove the power. If the input power is removed from the uninterruptible power supply units without first shutting down the cluster and the uninterruptible power supply units, the uninterruptible power supply units remain operational and eventually become drained of power.

When input power is restored to the uninterruptible power supply units, they begin to recharge but the nodes do not permit I/O activity on the virtual disks until the uninterruptible power supply is charged enough to enable all the data on the nodes to be saved in the event of an unexpected power loss. This might take as long as two hours. Shutting down the cluster prior to removing input power to the uninterruptible power supply units prevents the draining of battery power and makes it possible for I/O activity to resume as soon as input power is restored.

Attention: Before shutting down a node or the cluster, quiesce all I/O operations that are destined for this node or cluster. Failure to quiesce can result in failed I/O operations being reported to your host operating systems.

Begin the process of quiescing all I/O to the cluster by stopping the applications on your hosts that are using the VDisks that are provided by the cluster. Perform the following steps to quiesce all I/O to the cluster:

1. If you are unsure which hosts are using the VDisks that are provided by the cluster, determine the hosts that the VDisks are mapped to, as follows:
 - a. List the hosts that this VDisk is mapped to by issuing the following command, substituting the `vdiskname | id` variable with the name or ID of the VDisk:

```
svcinfo lsvdiskhostmap vdiskname | id
```
 - b. In the command output, located the host name or ID to determine which host the specified VDisk is mapped to. If no data is displayed, the VDisk is not mapped to any hosts.
2. Repeat the previous procedure for all VDisks.

Attention: If you are shutting down the entire cluster, you lose access to all VDisks that are being provided by this cluster.

When all I/O has been stopped, issue the **svctask stopcluster** to shut down a single node or the entire cluster in a controller manner. If you specify the node ID or node name, you can shut down a single node. After the command completes, the other node in the I/O group goes into write-through mode until the power to the node is returned and the node rejoins the cluster.

Attention: If this is the last node in an I/O group, you must specify the **-force** parameter. You will lose all access to the virtual disks in the I/O group. Before you enter this command, ensure that this is what you want to do.

If a shutdown command has been sent to the cluster and both cluster and uninterruptible power supply units have powered off, when input power is restored, press the power button on the uninterruptible power supply front panel to restart the uninterruptible power supply units.

Ensure that you have stopped all FlashCopy mappings and Metro or Global Mirror relationships. In addition, ensure that all data migration operations and forced deletions have completed before continuing. Entering **y** to the confirmation message processes the command. No feedback is then displayed. Entering anything other than **y** or **Y** results in the command not processing. No feedback is displayed.

Attention: If you are shutting down a single node and the other node in the I/O group is online, the cache on the partner node goes into the write-through mode and that you are exposed to a single point of failure if the partner node fails while the node is shut down. You also lose access to all VDisks being served by this I/O group.

An invocation example

```
svctask stopcluster
```

The resulting output You will be presented with the following warning:

```
Are you sure that you want to continue with the shut down?
```

stopstats

| Use the stopstats command to stop the collection of per-node statistics for virtual
| disks (VDisks), managed disks (MDisks), and nodes.

Syntax

```
▶— svctask — — stopstats —————▶
```

Description

This command turns off the generation of statistics, until you start them again (with the **svctask startstats**).

An invocation example

```
svctask stopstats
```

The resulting output

```
No feedback
```

Chapter 9. Backup and restore commands

The following commands are used for backing up and restoring configuration information with the SAN Volume Controller.

backup

Use the **backup** command to back up your configuration. You can enter this command any time after a cluster has been created.

Syntax

```
svcconfig -- backup [-quiet] [-v on | off]
```

Parameters

-quiet

Suppresses standard output (STDOUT) messages from the console.

-v on | off

On means verbose messages are displayed. Off means normal messages (the default) are displayed.

Description

The **backup** command extracts configuration information from the cluster, allowing you to restore your configuration whenever necessary. The **backup** command produces **.xml**, **.sh**, and **.log** files and saves them in the **/tmp** directory. The **.xml** file contains the extracted configuration information. The **.log** file contains details about command usage.

Note: If a previous **svc.config.backup.xml** file exists in **/tmp**, it is archived as **svc.config.backup.bak**; only one archive file is stored in the **/tmp** directory. Immediately archive the **.xml** file and the related **.key** files, and use the **clear** command to erase those files from the **/tmp** directory. Change all objects with default names to nondefault names; you cannot restore objects with default names.

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 **svconfig** commands. You can enter this command any time after a cluster has been created.

Syntax



Parameters

-all

Erases all configuration files.

-q | quiet

Suppresses console output (STDOUT).

-v on | off

Produces verbose output (on); the default is regular output (off).

Description

This command erases configuration files on the current config node.

You can use the **svconfig clear** command without the **-all** parameter to erase files of the form:

```
/tmp/svc.config*.sh  
/tmp/svc.config*.log
```

You can use the **svconfig clear** command with the **-all** parameter to erase files of the form:

```
/tmp/svc.config*.sh  
/tmp/svc.config*.log  
/tmp/svc.config*.xml  
/tmp/svc.config*.bak
```

An invocation example

```
svconfig clear -all
```


The resulting output

No feedback

help

Use the **help** command to obtain summary information about the syntax of the **svconfig** command. You can enter this command any time after a cluster has been created.

Syntax



Parameters

-ver

Returns the version number for the **svcconfig** command.

(action) -h | -?

Provides command help: the possible values for (action) are backup, clear, and restore.

-h | -?

Provides general help.

Description

This command provides syntax help for **svcconfig**.

An invocation example

```
svcconfig -ver
svcconfig -?
svcconfig backup -h
```

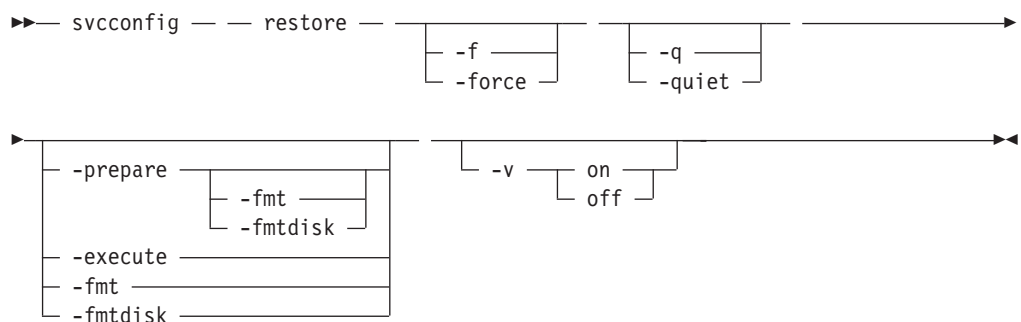
The resulting output

The help text displays.

restore

The **restore** command uses the configuration files in the **/tmp** directory to restore the cluster to its previous configuration.

Syntax



Parameters

-f | force

Forces continued processing where possible.

-q | quiet

Suppresses console output (STDOUT).

-prepare

Checks the current configuration against the information in **svc.config.backup.xml** on the configuration to be restored. Prepares commands for processing in **svc.config.restore.sh**, and produces a log of events in **svc.config.restore.prepare.log**.

-fmt | fmtdisk

Includes the **-fmtdisk** option on all **mkvdisk** commands to be issued.

-execute

Runs the command script **svc.config.restore.sh**. Produces a log of events in **svc.config.restore.execute.log**.

-v on | off

Produces verbose output (on); the default is regular output (off).

Description

The restore command restores the target cluster configuration from the **svc.config.backup.xml** file in the configuration files directory. If neither the **-prepare** nor **-execute** option is specified, only a single event log **svc.config.restore.log** is produced.

The command pauses for 5 minutes if any nodes are added during this process. You are informed of this at run-time.

After restoration, a VDisk consists of a specific list of MDisks. If the relevant MDisk group comprises a larger list, either now or in the future, the restored VDisk cannot manage any MDisks that are not presently in its own list.

The configuration files directory is **/tmp**.

An invocation example

```
svconfig restore -prepare  
svconfig restore -execute
```

The resulting output

No feedback

Chapter 10. Cluster diagnostic and service-aid commands

Cluster diagnostic and service-aid commands are designed to diagnose and find cluster problems.

The SAN Volume Controller enables you to perform service activity, such as problem determination and repair activities, with a limited set of command-line tools. When you are logged in under the administrator role, all command-line activities are permitted. When you are logged in under the service role, only those commands that are required for service are enabled. All of these commands apply under the service role.

applysoftware

The applysoftware command upgrades the cluster to a new level of software.

Syntax

```
svctask -- applysoftware [ -force ] -file filename [ -abort ]
```

Parameters

-force

(Optional) Forces continuation of the command even though any node-dependent VDIsks will be taken offline. Use the **force** parameter with caution; access to data on node-dependent VDIsks will be lost. If you use this parameter, you are prompted to confirm that you want to continue.

In addition, the **force** parameter must be used if one or more nodes are offline when you stop the upgrade with the **abort** parameter. The **force** parameter is not required when applying an upgrade to a cluster with errors in the error log.

-file *filename*

(Required) Specifies the file name of the new software package.

-abort

(Optional) Stops the upgrade and backs out to where the application was before you started the upgrade.

Note: The **force** parameter can be used with the **abort** parameter. If one or more nodes are offline, you must use the **force** parameter with the **abort** parameter.

Description

This command starts the upgrade process of the cluster to a new level of SAN Volume Controller software. The applysoftware command applies a level of software to the node in both service and nonservice modes. In service mode, the

applysoftware command is applied to the specific node. In nonservice mode, the applysoftware command is applied to the entire cluster.

The software package as specified by the file name must first be copied onto the current configuration node in the /home/admin/upgrade directory. You can use the PuTTY secure copy (scp) application to copy the file.

The command completes as soon as the upgrade process has successfully begun. The command fails and the upgrade package is deleted if:

- The given package fails an integrity check due to corruption.
- Any node in the cluster has a hardware type not supported by the new software.
- The new software level does not support upgrades from the currently installed software.
- The software level of a remote cluster is incompatible with the new software.
- There are any VDIsks that are dependent on the status of a node.

Note: The **force** parameter can be used to override this if you are prepared to lose access to data during the upgrade. Before proceeding, use the svcinfo lsnodeindependentvdisks command to list the node-dependent VDIsks at the time the command is run. If the command returns an error, you must move your quorum disks to MDIsks that are accessible through all nodes. Rerun the command until no errors are returned.

The actual upgrade completes asynchronously.

The svcinfo lssoftwareinfo command allows you to view the contents of the /home/admin/upgrade directory.

An invocation example

```
svctask applysoftware -file softwareupdate
```

The resulting output

No feedback

cherrstate

The **cherrstate** command marks an unfixed error as fixed. You can also use it to mark a fixed error as unfixed.

Syntax

►► svctask — — cherrstate — — -sequencenumber — *sequence_number* — —►

└─┬─┬─
 └─┬─
 - unfix

Parameters

-sequencenumber *sequence_number*
(Required) Specifies the error log sequence numbers to mark as fixed or as unfixed.

-unfix

(Optional) Specifies that the sequence numbers be marked as unfixed. Use this parameter when you have marked the wrong sequence number as fixed.

Description

The error log entries that the sequence numbers that you entered are marked as fixed. Use this command as a manual confirmation step that you have performed a maintenance procedure on the cluster, the fabric, or the subsystems.

This step is performed as part of the directed maintenance procedures (DMPs).

Optionally, if you have wrongly marked a sequence number as fixed, you can use the **-unfix** parameter to change the entry to unfixed.

An invocation example

```
svctask cherrstate -sequencenumber 2019
```

The resulting output

No feedback

clearerrlog

The **clearerrlog** command clears all entries from the error log including status events and any unfixed errors.

Syntax

```
svctask — — clearerrlog — [ -force ]
```

Parameters

-force

(Optional) Specifies that the **clearerrlog** command be processed without confirmation requests. If the **-force** parameter is not supplied, you are prompted to confirm that you want to clear the log.

Description

This command clears all entries from the error log. The entries are cleared even if there are unfixed errors in the log. It also clears any status events that are in the log.

Attention: This command is destructive. Use it only use when you have either rebuilt the cluster or have fixed a major problem that has caused entries in the error log that you do not want to manually fix.

An invocation example

```
svctask clearerrlog -force
```

The resulting output

No feedback

dumperrlog

The dumperrlog command dumps the contents of the error log to a text file.

Syntax

```
▶▶ svctask — — dumperrlog — [ -prefix — filename_prefix ] ▶▶
```

Parameters

-prefix *filename_prefix*

(Optional) A file name is created from the prefix and a time stamp, and has the following format:

```
prefix_NNNNNN_YYMMDD_HHMMSS
```

where *NNNNNN* is the node front panel name.

Note: If the **-prefix** parameter is not supplied, the dump is directed to a file with a system-defined prefix of **errlog**.

Description

When run with no parameters, this command dumps the cluster error log to a file using a system-supplied prefix of **errlog**, which includes the node ID and time stamp. When a file name prefix is provided, the same operation is performed but the details are stored in the **dumps** directory within a file with a name that starts with the specified prefix.

A maximum of ten error-log dump files are kept on the cluster. When the 11th dump is made, the oldest existing dump file is overwritten.

Error log dump files are written to **/dumps/elogs**. The contents of this directory can be viewed using the **svcinfolerrlogdumps** command.

Files are not deleted from other nodes until you issue the **cleardumps** command.

An invocation example

```
svctask dumperrlog -prefix testerrorlog
```

The resulting output

No feedback

finderr

The **finderr** command analyzes the error log for the highest severity unfixed error.

Syntax

```
▶▶ svctask — — finderr —▶▶
```

Description

The command scans the error log for any unfixed errors. Given a priority ordering within the code, the highest priority unfixed error is returned to standard output.

You can use this command to determine the order in which to fix the logged errors.

The Web-based directed maintenance procedures (DMPs) also use this command.

An invocation example

```
svctask finderr
```

The resulting output

```
Highest priority unfixed error 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 `svcinfolssnmpserver`.

setlocale

The setlocale command changes the locale setting for the cluster. It also changes command output to the chosen language.

Syntax

```
svctask -- setlocale -- -locale -- locale_id
```

Parameters

-locale *locale_id*
Specifies the locale ID.

Description

This command changes the language in which error messages are displayed as output from the command-line interface. Subsequently, all error messages from the command-line tools are generated in the chosen language. This command is run when you request a change of language (locale) and is generally run from the Web page. Issue the setlocale command to change the locale setting for the cluster; all interface output is changed to the chosen language. For example, to change the language to Japanese, type the following:

```
svctask setlocale -locale 3
```

where **3** is the value for Japanese. The following values are supported:

- **0** US English (default)
- **3** Japanese

Note: This command does not change the front panel display panel settings.

An invocation example

```
svctask setlocale -locale 3
```

The resulting output

No feedback

svqueryclock

The **svqueryclock** command returns the date, time, and current time-zone of the cluster.

Syntax

```
svqueryclock
```

Description

This command returns the date, time and current time-zone of the cluster.

An invocation example

```
svqueryclock
```

The resulting output

```
Mon Nov 25 14:59:28 GMT 2002
```

writesernum

Use the **writesernum** command to write the node serial number into the planar NVRAM.

Syntax

```
svctask -- writesernum -- -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 **svcinfo lsnodevpd** command to verify that the number is correct.
| The `system_serial_number` field contains the serial number.

| **An invocation example**

```
svctask writesernum -sernum 1300027 node1
```

| **The resulting output**

No feedback

Chapter 11. Host commands

The following commands enable you to work with host options with the SAN Volume Controller.

addhostiogrp

The **addhostiogrp** command enables you to map I/O groups to an existing host object.

Syntax

```
svctask -- addhostiogrp -- [-iogrp -- iogrp_list] --  
                             [-iogrpall] --  
-- [host_name | host_id] --
```

Parameters

-iogrp *iogrp_list*

(Required if you do not use **-iogrpall**) Specifies a colon-separated list of one or more I/O groups that must be mapped to the host. You cannot use this parameter with the **-iogrpall** parameter.

-iogrpall

(Required if you do not use **-iogrp**) Specifies that all the I/O groups must be mapped to the specified host. You cannot use this parameter with the **-iogrp** parameter.

host_id | *host_name*

(Required) Specifies the host to which the I/O groups must be mapped, either by ID or by name.

Description

This command allows you to map the list of I/O groups to the specified host object.

An invocation example

```
svctask addhostiogrp -iogrpall testhost
```


The resulting output

No feedback

addhostport

The **addhostport** command adds worldwide port names (WWPNs) or iSCSI names to an existing host object.

Syntax



Parameters

-hbawwpn *wwpn_list*

(Required if you do not use **iscsiname**) Specifies the list of fibre-channel host ports to add to the host. You cannot use this parameter with the **iscsiname** parameter.

-iscsiname *iscsi_name_list*

(Required if you do not use **hbawwpn**) Specifies the comma-separated list of iSCSI names to add to the host. Up to 16 names can be specified, and each name can contain up to 255 characters. Each name should comply with the iSCSI standard, RFC 3720. At least one WWPN or iSCSI name must be specified. You cannot use this parameter with the **hbawwpn** parameter.

-force

(Optional) Specifies that the list of ports be added to the host without the validation of any WWPNs or iSCSI names.

host_id | *host_name*

(Required) Specifies the host object to add ports to, either by ID or by name.

Description

This command adds a list of HBA WWPNs or iSCSI names to the specified host object. Any virtual disks that are mapped to this host object automatically map to the new ports.

Only WWPNs that are logged-in unconfigured can be added. For a list of candidate WWPNs, use the `svctask svcinfo lshbaportcandidate` command.

Some HBA device drivers do not log in to the fabric until they can recognize target LUNs. Because they do not log in, their WWPNs are not recognized as candidate ports. You can specify the **force** parameter with the `svctask addhostport` command to stop the validation of the WWPN list.

Note: When all I/O groups are removed from an iSCSI host, you cannot add a port to the iSCSI host until you have mapped the iSCSI host to at least one I/O group. After mapping the iSCSI host to at least one I/O group, resubmit the `svctask addhostport` command. After adding the port to the host, you must create a host authentication entry using the `svctask chhost` command.

An invocation example

```
svctask addhostport -hbawwpn 210100E08B251DD4 host_one
```

The resulting output

No feedback

An invocation example

```
svctask addhostport -iscsiname iqn.localhost.hostid.7f000001 mchost13
```


The resulting output

No feedback

chhost

The `chhost` command changes the name or type of a host object. This does not affect any existing virtual disk-to-host mappings.

Syntax



Parameters

-type *hpux* | *tpgs* | *generic*

(Optional) Specifies the type of host: **hpux**, **tpgs**, or **generic**. The default is **generic**. The **tpgs** parameter enables extra target port unit attentions. Refer to SAN Volume Controller host attachment documentation for more information on the hosts that require the **type** parameter.

-name *new_name_arg*

(Optional) Specifies the new name that you want to assign to the host object.

-mask *port_login_mask*

(Optional) Specifies which node target ports that a host can access. The port mask is four binary bits and is made of a combination of 0's and 1's, where 0 indicates that the corresponding target port cannot be used and 1 indicates that it can be used. The right-most bit in the mask corresponds to the lowest numbered target port (1 not 4) on a node. Valid mask values range from **0000** (no ports enabled) to **1111** (all ports enabled). For example, a mask of **0011** enables port 1 and port 2. The default value is **1111** (all ports enabled).

-chapsecret *chap_secret*

(Optional) Sets the Challenge Handshake Authentication Protocol (CHAP) secret used to authenticate the host for iSCSI I/O. This secret is shared between the host and the cluster. The CHAP secret for each host can be listed using the `svcinfo lsiscsiauth` command.

-nochapsecret

(Optional) Clears any previously set CHAP secret for this host.

host_name | *host_id*

(Required) Specifies the host object to modify, either by ID or by current name.

Description

This command can change the name of the specified host to a new name, or it can change the type of host. This command does not affect any of the current virtual disk-to-host mappings.

The port mask applies to logins from the host initiator port that are associated with the host object. For each login between a host HBA port and node port, the node examines the port mask that is associated with the host object for which the host HBA is a member and determines if access is allowed or denied. If access is denied, the node responds to SCSI commands as if the HBA port is unknown.

Note: When all I/O groups are removed from an iSCSI host, the `svctask lsiscsiauth` command does not display the authentication entry for that host. Use the `svctask addhostiogr` command to map the iSCSI host to at least one I/O group, and then use the `svctask addhostport` command to add the iSCSI port into it. You must also add authentication for that host using the `svctask chhost` command with either the `chapsecret` or `nochapsecret` parameter.

An invocation example

```
svctask chhost -name testhostlode -mask 0011 hostone
```

The resulting output

No feedback

mkhost

The `mkhost` command creates a logical host object.

Syntax

```
svctask -- mkhost -- [-name -- new_name] [-hbawwpn -- wwpn_list] [-iscsiname -- iscsi_name_list] [-iogrp -- iogrp_list] [-mask -- port_login_mask] [-force] [-type { hpux | tpgs | generic }]
```

Parameters

-name *new_name*

(Optional) Specifies a name or label for the new host object.

-hbawwpn *wwpn_list*

(Required if you do not use **iscsiname**) Specifies a list of host bus adapter (HBA) worldwide port names (WWPNs) to add to the specified host object. You cannot use this parameter with the **iscsiname** parameter.

-iscsiname *iscsi_name_list*

(Required if you do not use **hbawwpn**) Specifies one or more iSCSI names of this host. Up to 16 names can be specified, and each name can contain up to 255 characters. Each name should comply with the iSCSI standard, RFD 3720.

At least one WWPN or iSCSI name must be specified. You cannot use this parameter with the **hbawwpn** parameter.

-iogrp *iogrp_list*

(Optional) Specifies a set of one or more I/O groups that the host can access the VDisks from. I/O groups are specified using their names or IDs, separated by a colon. Names and IDs can be mixed in the list. If this parameter is not specified, the host is associated with all I/O groups.

-mask *port_login_mask*

(Optional) Specifies which node target ports that a host can access. The port mask is four binary bits and is made up of a combination of 0's and 1's, where 0 indicates that the corresponding target port cannot be used and 1 indicates that it can be used. The right-most bit in the mask corresponds to the lowest numbered target port (1 not 4) on a node. Valid mask values range from **0000** (no ports enabled) to **1111** (all ports enabled). For example, a mask of **0011** enables port 1 and port 2. The default value is **1111** (all ports enabled).

-force

(Optional) Specifies that a logical host object be created without validation of the WWPNs.

-type *hpux | tpgs | generic*

(Optional) Specifies the type of host: **hpux**, **tpgs**, or **generic**. The default is **generic**. The **tpgs** parameter enables extra target port unit attentions. Refer to SAN Volume Controller host attachment documentation for more information on the hosts that require the **type** parameter.

Description

The `mkhost` command associates one or more HBA WWPNs or iSCSI names with a logical host object. This command creates a new host. The ID is displayed when the command completes. You can subsequently use this object when you map virtual disks to hosts by using the `mkvdiskhostmap` command.

Issue the `mkhost` command only once. The cluster scans the fabric for WWPNs in the host zone. The cluster itself cannot filter into the hosts to determine which WWPNs are in which hosts. Therefore, you must use the `svctask mkhost` command to identify the hosts.

After you identify the hosts, mappings are created between hosts and virtual disks. These mappings effectively present the virtual disks to the hosts to which they are mapped. All WWPNs in the host object are mapped to the virtual disks.

Some HBA device drivers do not log in to the fabric until they can see target logical unit numbers (LUNs). Because they do not log in, their WWPNs are not recognized as candidate ports. You can specify the **force** parameter with this command to stop the validation of the WWPN list.

This command fails if you add the host to an I/O group that is associated with more host ports or host objects than is allowed by the limits within the cluster.

For additional information, see the `svctask mkvdiskhostmap` and `svcinfo lshbaportcandidate` commands.

An invocation example

```
svctask mkhost -name hostone -hbawwpn 210100E08B251DD4 -force -mask 1001
```

The resulting output

Host id [1] successfully created

An invocation example

```
svctask mkhost -iscsiname iqn.localhost.hostid.7f000001 -name newhost
```

The resulting output

Host, id [10], successfully created

rmhost

The **rmhost** command deletes a host object.

Syntax

```
svctask -- rmhost -- [-force] -- [host_name | host_id]
```

Parameters

-force

(Optional) Specifies that you want the system to delete the host object even if mappings still exist between this host and virtual disks (VDisks). When the **-force** parameter is specified, the mappings are deleted before the host object is deleted.

host_name | *host_id*

(Required) Specifies the host object to delete, either by ID or by name.

Description

The **rmhost** command deletes the logical host object. The WWPNs that were contained by this host object (if it is still connected and logged in to the fabric) are returned to the unconfigured state. When you issue the **svcinfo lshbaportcandidate** command, the host objects are listed as candidate ports.

If any mappings still exist between this host and virtual disks, the command fails unless you specify the **-force** parameter. When the **-force** parameter is specified, the **rmhost** command deletes the mappings before the host object is deleted.

An invocation example

```
svctask rmhost host_one
```

The resulting output

No feedback

rmhostiogr

The **rmhostiogr** command enables you to delete mappings between one or more I/O groups and a specified host object.

Syntax

```
svctask -- rmhostiogr -- [-iogrp | -iogrpall] -- iogrp_list
```



Parameters

-iogrp *iogrp_list*

(Required) Specifies a set of one or more I/O group mappings that will be deleted from the host. You cannot use this parameter with the **-iogrpall** parameter.

-iogrpall

(Optional) Specifies that all the I/O group mappings that are associated with the specified host must be deleted from the host. You cannot use this parameter with the **-iogrp** parameter.

-force

(Optional) Specifies that you want the system to remove the specified I/O group mappings on the host even if the removal of a host to I/O group mapping results in the loss of VDisk-to-host mappings.

host_id | *host_name*

(Required) Specifies the identity of the host either by ID or name from which the I/O group mappings must be deleted.

Description

The `rmhostiogrps` command deletes the mappings between the list of I/O groups and the specified host object.

This command fails if any of the I/O groups that are deleted from the host contain VDisks that have host mappings to the host. To resolve this problem, do one of the following:

- Delete the VDisks-to-host mappings that are causing the error.
- Delete the VDisks or the host.
- Issue the `rmhostiogrps` command with the **-force** parameter.

Note: When all I/O groups are removed from an iSCSI host, and you want to add an iSCSI port to the host, refer to the `svctask addhostport` and `svctask chhost` commands.

An invocation example

```
svctask rmhostiogrps -iogrp 1:2 host0
```


The resulting output

No feedback

rmhostport

The `rmhostport` command deletes worldwide port names (WWPNs) or iSCSI qualified names (IQNs) from an existing host object.

Syntax



Parameters

-hbawwpn *wwpn_list*

(Required) Specifies the list of fibre-channel host ports to delete from the host.

-iscsiname *iscsi_name_list*

(Optional) Specifies the comma-separated list of IQNs to delete from the host. Up to 16 names can be specified. At least one WWPN or iSCSI name must be specified.

-force

(Optional) Forces the deletion of the specified ports. This overrides the check that all of the WWPNs or IQNs in the list are mapped to the host specified.

host_name | host_id

(Required) Specifies the host name or the host ID.

Description

This command deletes the list of HBA WWPNs or IQNs from the specified host object. If the WWPN ports are still logged in to the fabric, they become unconfigured and are listed as candidate WWPNs. See also the `svcinfo lshbaportcandidate` command.

Any virtual disks that are mapped to this host object are automatically unmapped from the ports.

Replacing an HBA in a host: List the candidate HBA ports by issuing the `svcinfo lshbaportcandidate` command. A list of the HBA ports that are available to be added to host objects is displayed. One or more of these ports corresponds with one or more WWPNs that belong to the new HBA. Locate the host object that corresponds to the host in which you have replaced the HBA. The following command lists all the defined host objects:

```
svcinfo lshost
```

To list the WWPNs that are currently assigned to the host, issue the following:

```
svcinfo 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:

```
svctask addhostport -hbawwpn one or more existing WWPNs  
separated by : hostobjectname/ID
```

where *one or more existing WWPNs separated by :* and *hostobjectname/id* correspond to those values listed in the previous steps.

Remove the old ports from the host object by issuing the following command:

```
svctask rmhostport -hbawwpn one or more existing WWPNs  
separated by : hostobjectname/ID
```


where *one or more existing WWPNs separated by :* corresponds with those WWPNs that are listed in the previous step that belong to the old HBA that has been replaced. Any mappings that exist between the host object and VDIs are automatically applied to the new WWPNs. Therefore, the host recognizes that the VDIs are the same SCSI LUNs as before. See the *Multipath Subsystem Device Driver: User's Guide* for additional information about dynamic reconfiguration.

An invocation example

```
svctask rmhostport -hbawwpn 210100E08B251DD4 host_one
```

The resulting output

No feedback

|
|
|
|

An invocation example

```
svctask rmhostport -iscsiname iqn.localhost.hostid.7f000001 mchost13
```

The resulting output

No feedback

Chapter 12. Virtual disk commands

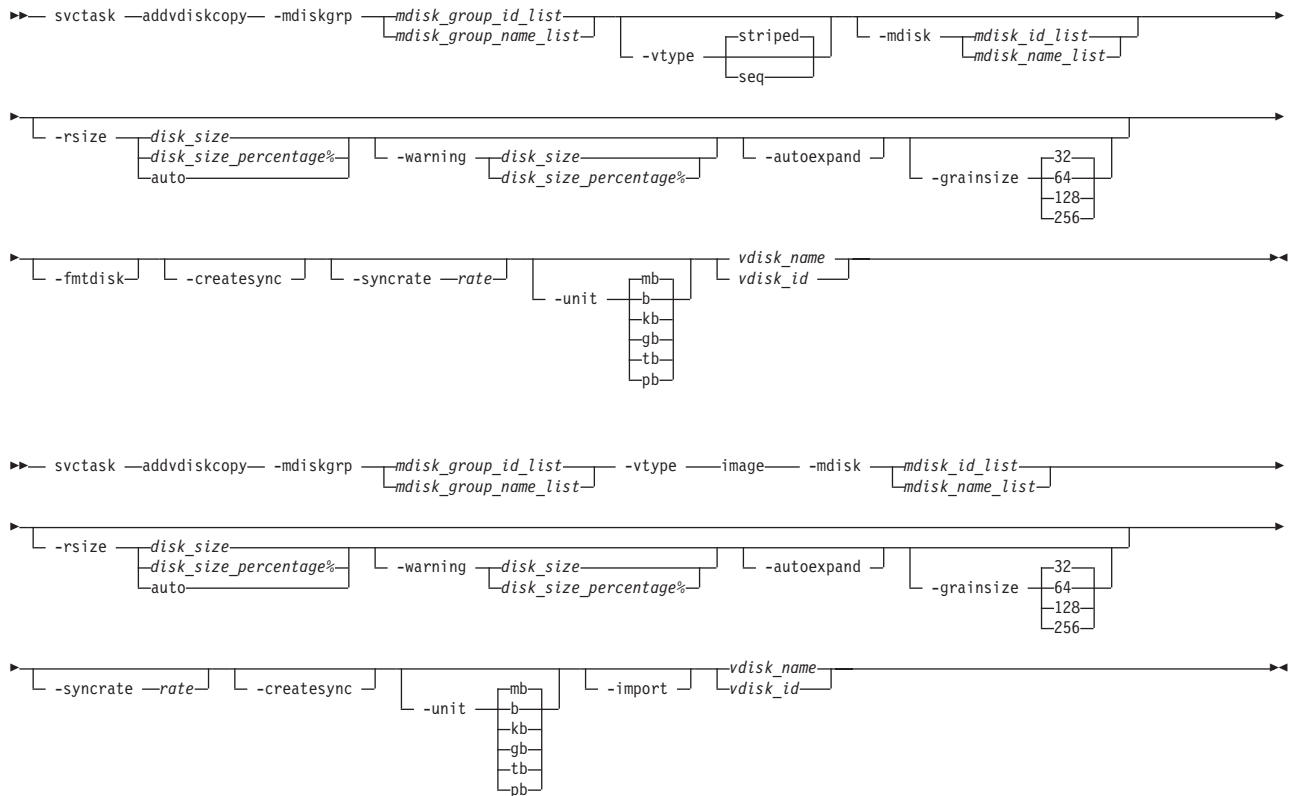
The following commands enable you to work with virtual disk options with the SAN Volume Controller.

addvdiskcopy

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

Note: The first syntax diagram depicts the addition of a sequential or striped mode virtual disk. The second syntax diagram depicts the addition of an image mode virtual disk.

Syntax



Parameters

-mdiskgrp *mdisk_group_id_list* | *mdisk_group_name_list*

(Required) Specifies the managed disk groups to use to create copies for the virtual disk. You must specify a group for each copy that is being added.

-vtype *seq* | *striped* | *image*

(Optional) Specifies the virtualization type for the copy: sequential, striped, or image. The type can be different than the virtualization types for other copies on the VDisk. The default virtualization type is **striped**.

-mdisk *mdisk_id_list* | *mdisk_name_list*

(Optional) Specifies one or more managed disks (MDisks). For sequential and image mode copies, you must specify a single MDisk that has sufficient free extents. For image mode copies, the MDisk must be in unmanaged mode. For sequential mode copies the MDisk must be in the managed mode.

-syncrate *rate*

(Optional) Specifies the copy synchronization rate. A value of zero (0) prevents synchronization. The default value is 50. For the supported **-syncrate** values and their corresponding rates, see Table 6 on page 100.

-createsync

(Optional) Suppresses the synchronization of the new VDisk copy with the primary copy. Using this parameter can cause data corruption if the primary copy fails and leaves an unsynchronized secondary copy to provide data. Using this parameter can cause loss of read stability in unwritten areas if the primary copy fails, data is read from the primary copy, and then different data is read from the secondary copy. To avoid data loss or read stability loss, use this parameter only for a primary copy that has been formatted and not written to, and with the **-fmtdisk** parameter.

-fmtdisk

(Optional) Formats a sequential or striped mode copy. You must also specify the **-createsync** parameter, which labels the formatted copy as identical to the primary copy. The **-fmtdisk** parameter causes the VDisk to go offline until new VDisk copy formatting completes. To query the formatting progress, use the **lsvdiskprogress** command.

-rsize *disk_size* | *disk_size_percentage%* | **auto**

(Optional) Makes the copy space-efficient and specifies the real size of the copy. Specify the *disk_size* | *disk_size_percentage* value using an integer, or an integer immediately followed by the percent character (%). The default units for *disk_size* are megabytes (MB); to specify different units, use the **-unit** parameter. The **auto** option creates a VDisk copy that uses the entire size of the MDisk; if you specify the **-rsize auto** option, you must also specify the **-vtype image** option.

-warning *disk_size* | *disk_size_percentage%*

(Optional) Requires that the **-rsize** parameter also be specified. Generates a warning when the used disk capacity on the space-efficient copy first exceeds the specified threshold. You can specify a *disk_size* integer, which defaults to megabytes (MB) unless the **-unit** parameter is specified; or you can specify a *disk_size%*, which is a percentage of the virtual disk size. If **-autoexpand** is enabled, the default value for **-warning** is 80% of the virtual disk capacity. If **-autoexpand** is not enabled, the default value for warning is 80% of the real capacity. To disable warnings, specify 0.

-autoexpand

(Optional) Requires that the **-rsize** parameter also be specified. Specifies that space-efficient copies automatically expand their real capacities by allocating new extents from their managed disk group. If the **-autoexpand** parameter is specified, the **-rsize** parameter specifies a capacity that is reserved by the copy. This protects the copy from going offline when its managed disk group runs out of space by allowing it to consume this reserved space first.

-grainsize 32 | 64 | 128 | 256

(Optional) Requires that the **-rsize** parameter also be specified. Sets the grain size (KB) for a space-efficient VDisk. The default is 32 KB.

-unit b | kb | mb | gb | tb | pb

(Optional) Specifies the data units for the **-rsize** and **-warning** parameters.

-import

(Optional) Imports an image mode disk that contains a space-efficient volume into the cluster. Requires that the **-rsize** and **-vtype image** parameters also be specified.

vdisk_name | vdisk_id

(Required) Specifies the virtual disk to add the VDisk copy to, either by ID or by name.

Description

The **addvdiskcopy** command adds a copy to an existing VDisk, which changes a nonmirrored VDisk into a mirrored VDisk. Use the **-mdiskgrp** parameter to specify the managed disk group that will provide storage for the copy; the **svcinfo lsmdiskgrp** command lists the available managed disk groups and the amount of available storage in each group.

The virtualization types are defined as follows:

sequential (seq)

This policy requires the **-mdisk** parameter with a single managed disk as its argument. This MDisk must be in the managed mode.

It creates the virtual disk using extents from the given managed disk (assuming there are enough free extents on the managed disk).

striped

This is the default policy. If the **-vtype** parameter is not specified, this policy is used in its default form. That is, all managed disks in the managed disk group are used to create the virtual disk. The striping is at an extent level; one extent from each managed disk in the group is used. For example, a managed disk group with 10 managed disks uses one extent from each managed disk, then it uses the 11th extent from the first managed disk, and so on.

If the **-mdisk** parameter is also specified, you can supply a list of managed disks to use as the stripe set. This can be two or more managed disks from the same managed disk group. The same circular algorithm is used across the striped set. However, a single managed disk can be specified more than once in the list. For example, if you enter **-m 0:1:2:1**, the extents are from the following **managed** disks: 0, 1, 2, 1, 0, 1, 2, and so forth. All MDisks that are specified in the **-mdisk** parameter must be in **managed** mode.

image This policy allows image mode virtual disks to be created when a managed disk already has data on it, perhaps from a previrtualized subsystem. When an image mode virtual disk is created, it directly corresponds to the (previously unmanaged) managed disk that it was created from; therefore, virtual disk logical block address (LBA) *x* equals managed disk LBA *x*. You can use this command to bring a nonvirtualized disk under the control of the cluster. After it is under the control of the cluster, you can migrate the virtual disk from the single managed disk. When it is migrated, the virtual disk is no longer an image mode virtual disk.

You can add image mode VDIs to an already populated MDisk group with other types of VDIs, such as a striped or sequential.

Note: An image mode copy must be at least as large as the VDisk that it is being added to, but any capacity beyond the size of the VDisk is not accessible.

The command returns the ID of the newly created VDisk copy.

Table 6 provides the relationship of the *rate* value to the data copied per second.

Table 6. Relationship between the rate value and the data copied per second

User-specified <i>rate</i> attribute value	Data copied/sec
1 - 10	128 KB
11 - 20	256 KB
21 - 30	512 KB
31 - 40	1 MB
41 - 50	2 MB
51 - 60	4 MB
61 - 70	8 MB
71 - 80	16 MB
81 - 90	32 MB
91 - 100	64 MB

An invocation example

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

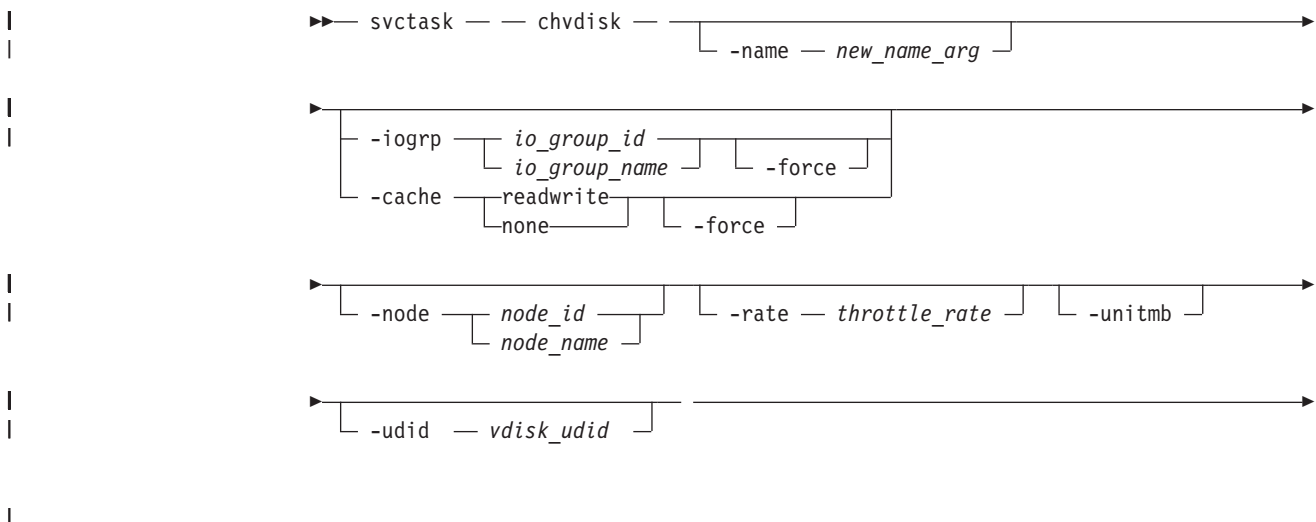
The resulting output

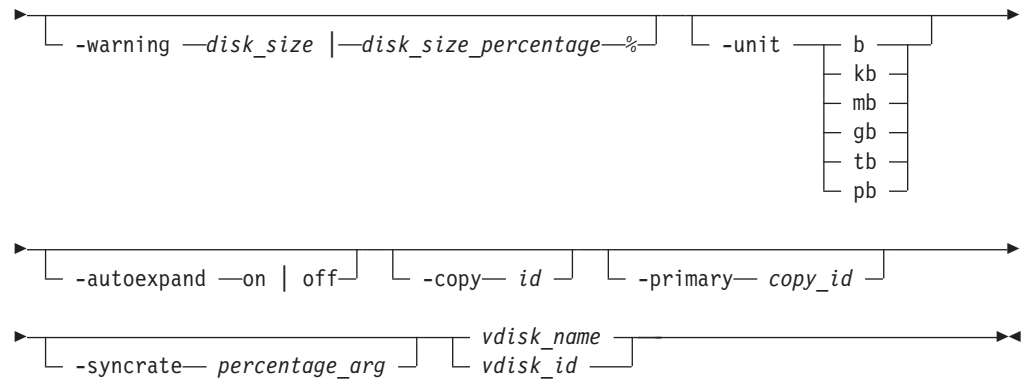
```
Vdisk [8] copy [1] successfully created
```

chvdisk

The chvdisk command modifies the properties of a virtual disk, such as the disk name, I/O group, I/O governing rate, or unit number.

Syntax





Parameters

-name *new_name_arg*

(Optional) Specifies a new name to assign to the virtual disk. You cannot use this parameter with the **-iogrp**, **-rate**, **-node**, or **-udid** parameters. This parameter is required if you do not use the **-iogrp**, **-rate**, or **-udid** parameter.

-iogrp *io_group_id* | *io_group_name*

(Optional) Specifies a new I/O group to move the virtual disk to, by I/O group ID or I/O group name. You can use the **-node** parameter with the **-iogrp** parameter to specify a preferred node for the specified VDisk.

Notes:

1. If the VDisk has a mapping to any hosts, it is not possible to move the VDisk to an I/O group, unless all of those hosts are associated with the new I/O group.
2. This parameter can fail if there is not enough space to allocate bitmaps for a mirrored VDisk in the target I/O group.
3. This parameter can fail if any copy is not synchronized. The **-force** parameter can be used to force the move, but this resynchronizes the VDisk.
4. If the VDisk is offline, use one of the `recovervdisk` commands to recover the VDisk and bring it back online. Beginning with SAN Volume Controller version 4.3.1, use of the recovery I/O group is not required.

-cache readwrite | **none**

(Optional) Specifies the caching options for the VDisk. Valid entries are **readwrite**, to enable the cache for the VDisk, or **none**, to disable the cache mode for the VDisk.

-force

(Optional) The **force** parameter can only be used for changing the I/O group of a VDisk or the caching mode. Use the **force** parameter with the **iogrp** parameter to force the VDisk to be removed from an I/O group. Use the **force** parameter with the **cache** parameter to specify that you want the system to change the cache mode of the VDisk even if the I/O group is offline. This option overrides the cache flush mechanism.

Attention:

1. If the **force** parameter is used for changing the caching mode or I/O group of a VDisk, the contents of the cache are discarded and the VDisk might be corrupted by the loss of the cached data. This could occur if the cluster is able to destage all write data from the cache or not. The **force** parameter should be used with caution.
2. If the **force** parameter is used to move a VDisk that has out-of-sync copies, a full resynchronization is required.

-rate *throttle_rate* [-**unitmb**]

(Optional) Specifies the I/O governing rate for the VDisk, which caps the amount of I/O that is accepted. The default *throttle_rate* units are I/Os. To change the *throttle_rate* units to megabytes per second (MBps), specify the **-unitmb** parameter. The governing rate for a virtual disk can be specified by I/Os or by MBps, but not both. However, you can set the rate to I/Os for some virtual disks and to MBps for others.

You cannot use this parameter with the **-name**, **-iogrp**, **-node**, or **-udid** parameters.

-udid *vdisk_udid*

(Optional) Specifies the unit number (**udid**) for the disk. The *vdisk_udid* is an identifier that is required to support OpenVMS hosts; no other systems use this parameter. Valid options are a decimal number from 0 to 32 767 or a hexadecimal number from 0 to 0x7FFF. A hexadecimal number must be preceded by 0x (for example, 0x1234). If you do not use the **-udid** parameter, the default **udid** is 0.

You cannot use this parameter with the **-name**, **-iogrp**, **-node**, or **-rate** parameters.

-warning *disk_size* | *disk_size_percentage*%

(Optional) Generates a warning when the used disk capacity on the space-efficient copy first exceeds the specified threshold. You can specify a *disk_size* integer, which defaults to MBs unless the **-unit** parameter is specified; or you can specify a *disk_size*%, which is a percentage of the virtual disk size. To disable warnings, specify 0 or 0%.

-unit **b** | **kb** | **mb** | **gb** | **tb** | **pb**

(Optional) Specifies the data units to use for the **-warning** *disk_size* parameter.

-autoexpand **on** | **off**

(Optional) Specifies whether space-efficient VDisk copies automatically expand their real capacities by allocating new extents from their managed disk group. To use this parameter, the VDisk must be space-efficient.

-copy *id*

(Optional) Specifies the copy to apply the changes to. You must specify this parameter with the **-autoexpand** or **-warning** parameter. The **-copy** parameter is required if the specified VDisk is mirrored and only one VDisk copy is space-efficient. If both copies are space-efficient and the **-copy** parameter is not specified, the specified **-autoexpand** or **-warning** parameter is set on both copies.

-primary *copy_id*

(Optional) Specifies the primary copy. Changing the primary copy only takes effect when the new primary copy is online and synchronized. If the new primary is online and synchronized when the command is issued, the change takes effect immediately.

-syncrate *percentage*

(Optional) Specifies the copy synchronization rate, as a percentage of the peak synchronization rate. A value of zero (0) prevents synchronization.

-node *node_id* | *node_name*

(Optional) Specifies a preferred node for the specified VDisk. When using this parameter, you must also specify the **-iogrp** parameter. You cannot use this parameter with the **-name**, **-rate**, or **-udid** parameters.

vdisk_name | *vdisk_id*

(Required) Specifies the virtual disk to modify, either by ID or by name.

Description

The `chvdisk` command modifies a single property of a virtual disk (VDisk). To change the VDisk name and modify the I/O group, for example, you must issue the command twice.

You can specify a new name or label. You can use the new name subsequently to refer to the virtual disk. To specify a preferred node for the VDisk, use the **-node** *node_id* | *node_name* parameter.

You can change the I/O group with which this virtual disk is associated. However, to change the I/O group, you must first flush the cache within the nodes in the current I/O group to ensure that all data is written to disk. Ensure that you suspend I/O operations at the host level before you perform this operation.

Attention:

1. Do not move a VDisk to an offline I/O group under any circumstance. To avoid any data loss, you must ensure that the I/O group is online before you move the VDIs.
2. Do not move an offline VDisk to the recovery I/O group. Beginning with SAN Volume Controller version 4.3.1, use of the recovery I/O group is not required. Instead, use one of the `recovervdisk` commands to recover the VDisk and bring it back online.

You can set a limit on the amount of I/O transactions that is accepted for this virtual disk. It is set in terms of I/Os per second or MBs per second. By default, no I/O governing rate is set when a virtual disk is created.

Attention: All capacities, including changes, must be in multiples of 512 bytes. An error occurs if you specify a capacity that is not a multiple of 512, which can only happen when byte units (**-b**) are used. The default capacity is in MB.

When the virtual disk is created, there is no throttling applied to it. Using the **-rate** parameter can change this. To change the virtual disk back to an unthrottled state, specify 0 (zero) with the **-rate** parameter.

You can migrate a VDisk to a new I/O group to manually balance the workload across the nodes in the cluster. You might end up with a pair of nodes that are overworked and another pair that are underworked. Use the following procedure to migrate a single VDisk to a new I/O group. Repeat for other VDIs as required.

Attention: This is a disruptive procedure. Access to the VDisk is lost while you follow this procedure.

Ensure that when you migrate a VDisk to a new I/O group, you quiesce all I/O operations for the VDisk. Determine the hosts that are using this VDisk. Stop and delete any FlashCopy mappings or Metro or Global Mirror relationships that use this VDisk. To check if the VDisk is part of a relationship or mapping, issue the `svcinfolsvdisk vdiskname | id` command, where `vdiskname | id` is the name or ID of the VDisk.

Look for the `FC_id` and `RC_id` fields. If these are not blank, the VDisk is part of a mapping or relationship. See the FlashCopy commands or Metro Mirror and Global Mirror commands for details on how to stop or delete the mapping or relationship. Issue the following command to migrate the VDisk:

```
svctask chvdisk -iogrp newiogrpname|id vdiskname|id
```

Follow the procedure to discover the new vpaths and to check that each vpath is presenting the correct number of paths. Refer to Multipath Subsystem Device Driver (SDD) documentation for details on how to dynamically reconfigure SDD for the given host operating system.

Note: The command fails if you attempt to change the primary copy of a mirrored VDisk while the `repairvdiskcopy -resync` command is running.

An invocation example

```
svctask chvdisk -rate 2040 -unitmb 6
```

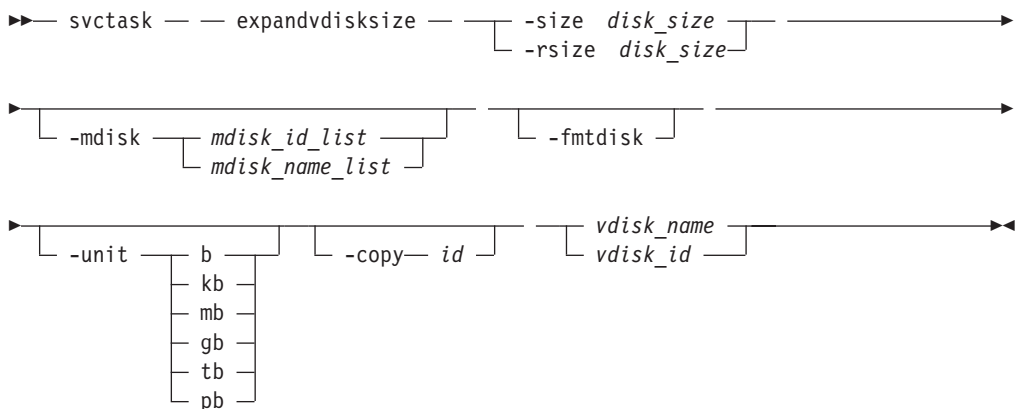
The resulting output

No feedback

expandvdisksize

The `expandvdisksize` command expands the size of a VDisk by a given capacity.

Syntax



Parameters

-size *disk_size*

(Optional) Specifies the capacity by which the virtual disk is expanded. Disk size is used with the value of the unit. All capacities, including changes must be in multiples of 512 bytes. An error occurs if you specify a capacity that is not a multiple of 512, which can only occur when byte units (**-unit b**) are used.

However, an entire extent is reserved even if it is only partially used. The default *disk_size* unit is megabytes (MB). You cannot specify the **-size** parameter with the **-rsize** parameter. You must specify either **-size** or **-rsize**. If the VDisk is space-efficient, MDisks cannot be specified.

-rsize *disk_size*

(Optional) Specifies the capacity by which to increase the real size of a space-efficient VDisk. Specify the *disk_size* value using an integer. Specify the unit for a *disk_size* integer using the **-unit** parameter; the default unit is megabytes (MB). The **-rsize** value can be greater than, equal to, or less than the size of the VDisk. You cannot specify the **-rsize** parameter with the **-size** parameter. You must specify either **-size** or **-rsize**.

-copy *id*

(Optional) Specifies the copy to change the real capacity for. You must also specify the **-rsize** parameter; you can only modify the real capacity of a VDisk copy. The **-copy** parameter is required if the specified VDisk is mirrored and only one copy is space-efficient. If the VDisk is mirrored, both copies are space-efficient and **-copy** is not specified, both copies are modified by the same amount.

-mdisk *mdisk_id_list* | *mdisk_name_list*

(Optional) Specifies the list of one or more MDisks to be used as the stripe set. The extents that expand the VDisk come from the specified list of MDisks. All MDisks in the list must be part of the same MDisk group. The **-mdisk** parameter cannot be used if the specified VDisk is mirrored.

-fmtdisk

(Optional) Specifies that the VDisk be formatted before use. This parameter formats the new extents that have been added to the VDisk as a result of the **expandvdisksize** command. The **expandvdisksize** command completes asynchronously if you use this parameter.

-unit **b** | **kb** | **mb** | **gb** | **tb** | **pb**

(Optional) Specifies the *disk_size* unit for the **-size** or **-rsize** parameter. The default value is megabytes (MB).

vdisk_name | *vdisk_id*

(Required) Specifies the virtual disk to modify, either by ID or by name.

Description

The **expandvdisksize** command can be used to expand the physical capacity that is allocated to a particular VDisk by the specified amount. The command can also be used to expand the virtual capacity of a space-efficient VDisk without altering the physical capacity that is assigned to the VDisk. To change the capacity of a non-space-efficient VDisk, or the virtual capacity of a space-efficient VDisk, use the **-size** parameter. To change the real capacity of a space-efficient VDisk, use the **-rsize** parameter.

The default capacity units are MB.

When a VDisk is expanded, the virtualization policy can change. Its mode becomes striped even if it was previously sequential. See the **svctask mkvdisk** command for details of the virtualization policies.

To run the **expandvdisksize** command on a mirrored VDisk, all copies of the VDisk must be synchronized. The command formats all copies of a mirrored VDisk automatically.

An invocation example

To increase the capacity of VDisk1 by 2048 bytes by using extents from two MDisks and to format the new part of the VDisk, enter:

```
svctask expandvdisksize -size 2048 -unit b -mdisk  
mdisk0:mdisk1 -fmtdisk vdisk1
```

The resulting output

No feedback

An invocation example

To increase the capacity of VDisk1 by 100 MB using extents from two MDisks, and to format the new part of the VDisk, enter:

```
svctask expandvdisksize -size 100 -unit mb -mdisk mdisk0:mdisk1 -fmtdisk vdisk1
```

The resulting output

No feedback

An invocation example

To increase the real capacity of space-efficient VDisk **vdisk2** by 100 MB without changing the virtual capacity, and to spread the extents across all MDisks in the MDisk group, enter:

```
svctask expandvdisksize -rsize 100 -unit mb vdisk2
```

The resulting output

No feedback

An invocation example

To increase the real capacity of space-efficient VDisk copy id 1 of mirrored VDisk **vdisk3** by 100 MB, enter:

```
svctask expandvdisksize -rsize 100 -unit mb -copy 1 vdisk3
```

The resulting output

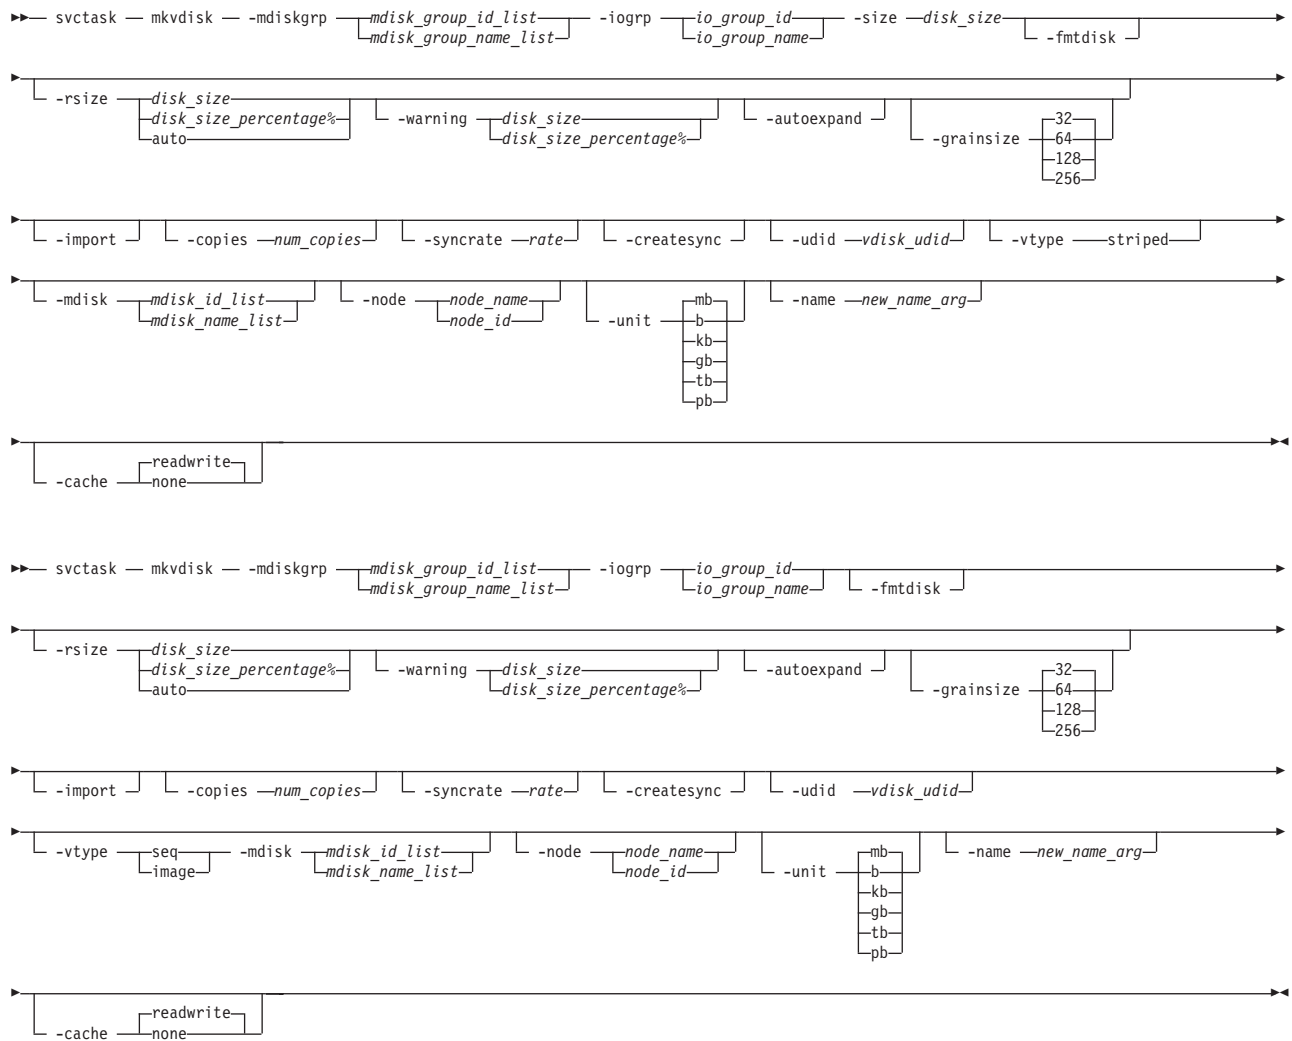
No feedback

mkvdisk

The **mkvdisk** command creates sequential, striped, or image mode virtual disk objects. When they are mapped to a host object, these objects are seen as disk drives with which the host can perform I/O operations.

Note: The first syntax diagram depicts the creation of a **striped** mode virtual disk. The second syntax diagram depicts the creation of a **sequential** or **image** mode virtual disk.

Syntax



Parameters

-mdiskgrp *mdisk_group_id_list* | *mdisk_group_name_list*

(Required) Specifies one or more managed disk groups to use when you are creating this virtual disk. If you are creating multiple copies, you must specify one managed disk group per copy. The primary copy is allocated from the first managed disk group in the list.

-iogrp *io_group_id* | *io_group_name*

(Required) Specifies the I/O group (node pair) with which to associate this virtual disk.

-udid *vdisk_udid*

(Optional) Specifies the unit number (**udid**) for the disk. The **udid** is an identifier that is required to support OpenVMS hosts; no other systems use this parameter. Valid options are a decimal number 0 - 32 767, or a hexadecimal number 0 - 0x7FFF. A hexadecimal number must be preceded by **0x** (for example, **0x1234**).

-size *disk_size*

(Required for sequential [**seq**] or **striped** VDisk creation) (Optional for **image** VDisk creation) Specifies the capacity of the virtual disk, which is used with the value of the unit. All capacities, including changes, must be in multiples of 512 bytes. An error occurs if you specify a capacity that is not a multiple of

512, which can only happen when byte units (-b) are used. However, an entire extent is reserved even if it is only partially used. The default capacity is in MB. You can specify a capacity of 0. Specify the size in bytes in multiples of logical block address (LBA) sizes.

Note: If you do not specify the **-size** parameter when you create an image mode disk, the entire MDisk capacity is used.

-rsize *disk_size* | *disk_size_percentage%* | **auto**

(Optional) Makes the VDisk space-efficient; otherwise, the VDisk is fully allocated. Specify the *disk_size* | *disk_size_percentage* value using an integer, or an integer immediately followed by the percent character (%). Specify the units for a *disk_size* integer using the **-unit** parameter; the default is MB. The **-rsize** value can be greater than, equal to, or less than the size of the VDisk. The **auto** option creates a VDisk copy that uses the entire size of the MDisk; if you specify the **-rsize auto** option, you must also specify the **-vtype image** option.

-warning *disk_size* | *disk_size_percentage%*

(Optional) Requires that the **-rsize** parameter also be specified. Specifies a threshold at which a warning error log is generated for VDisk copies. A warning is generated when the used disk capacity on the space-efficient copy first exceeds the specified threshold. You can specify a *disk_size* integer, which defaults to MBs unless the **-unit** parameter is specified; or you can specify a *disk_size%*, which is a percentage of the virtual disk size. If **-autoexpand** is enabled, the default value for **-warning** is 80% of the virtual disk capacity. If **-autoexpand** is not enabled, the default value for warning is 80% of the real capacity. To disable warnings, specify 0.

-autoexpand

(Optional) Specifies that space-efficient copies automatically expand their real capacities by allocating new extents from their managed disk group. Requires that the **-rsize** parameter also be specified. If the **-autoexpand** parameter is specified, the **-rsize** parameter specifies a capacity that is reserved by the copy. This protects the copy from going offline when its managed disk group runs out of space by having the managed disk group to consume this reserved space first.

The parameter has no immediate effect on **image** mode copies. However, if the image mode copy is subsequently migrated to managed mode, the copy is then automatically expanded.

-grainsize 32 | 64 | 128 | 256

(Optional) Sets the grain size (KB) for a space-efficient VDisk. This parameter also requires that the **-rsize** parameter be specified. The default is 32 KB. If you are using the space-efficient VDisk in a FlashCopy map, use the same grain size as the map grain size for best performance. If you are using the space-efficient VDisk directly with a host system, use a small grain size.

-import

(Optional) Imports a space-efficient VDisk from the Mdisk. This parameter also requires that the **-rsize** parameter be specified.

-copies *num_copies*

(Optional) Specifies the number of copies to create. The *num_copies* value can be 1 or 2. Setting the value to 2 creates a mirrored VDisk. The default value is 1.

-syncrate *rate*

(Optional) Specifies the copy synchronization rate. A value of zero (0) prevents

synchronization. The default value is 50. For the supported **-syncrate** values and their corresponding rates, see Table 7 on page 111.

-createsync

(Optional) Creates copies in sync. Use this parameter if you have already formatted the MDisks, or when read stability to unwritten areas of the VDisk is not required.

-fmtdisk

(Optional) Specifies that the virtual disk be formatted before it can be used. The **-fmtdisk** parameter formats (sets to all zeros) the extents that make up this VDisk after it is created. If this parameter is used, the command completes asynchronously; you can query the status using the **svcinfo lsvdiskprogress** command.

The **-fmtdisk** parameter is not required when creating space-efficient virtual disks. Space-efficient VDIsks return zeros for extents that have not been written to.

The **-fmtdisk** parameter synchronizes mirrored copies by default.

Note: You cannot specify this parameter with the **-vtype image** parameter.

-vtype seq | striped | image

(Optional) Specifies the virtualization type. When creating sequential or image mode VDIsks, you must also specify the **-mdisk** parameter. The default virtualization type is striped.

-node node_id | node_name

(Optional) Specifies the preferred node ID or the name for I/O operations to this virtual disk. You can use the **-node** parameter to specify the preferred access node.

Note: This parameter is required for the subsystem device driver (SDD). The cluster chooses a default if you do not supply this parameter.

-unit b | kb | mb | gb | tb | pb

(Optional) Specifies the data units to use in conjunction with the capacity that is specified by the **-size** and **-rsize** parameters.

-mdisk mdisk_id_list | mdisk_name_list

(Optional) Specifies one or more managed disks. For sequential and image mode VDIsks, the number of MDisks must match the number of copies. For sequential mode VDIsks, each MDisk must belong to the specified MDisk group. For striped VDIsks, you cannot specify the **-mdisk** parameter if the **-copies** value is greater than 1. When creating a single copy striped VDisk, you can specify a list of MDisks to stripe across.

-name new_name_arg

(Optional) Specifies a name to assign to the new virtual disk.

-cache readwrite | none

(Optional) Specifies the caching options for the VDisk. Valid entries are **readwrite** or **none**. The default is **readwrite**. If you do not specify the **-cache** parameter, the default value (**readwrite**) is used.

Description

This command creates a new virtual disk object. You can use the command to create a variety of types of virtual disk objects, making it one of the most complex commands.

You must decide which managed disk group or groups provide the storage for the VDisk. Use the **svcinfolsmdiskgrp** command to list the available managed disk groups and the amount of free storage in each group. If you are creating a VDisk with more than one copy, each MDisk group that you specify must have enough space for the size of the VDisk.

Choose an I/O group for the VDisk. This determines which nodes in the cluster process the I/O requests from the host systems. If you have more than one I/O group, ensure that you distribute the VDIs between the I/O groups so that the I/O workload is shared evenly between all nodes. Use the **svcinfolsiogrp** command to show the I/O groups and the number of virtual disks that are assigned to each I/O group.

Note: It is normal for clusters with more than one I/O group to have MDisk groups that have VDIs in different I/O groups. FlashCopy processing can make copies of VDIs whether the source and target VDIs are in the same I/O group. If, however, you plan to use intracluster Metro or Global Mirror operations, ensure that both the master and auxiliary VDisk are in the same I/O group.

Specify the virtualization type using the **-vtype** parameter; the supported types are sequential (**seq**), **striped**, and **image**.

sequential (seq)

This virtualization type creates the virtual disk using sequential extents from the specified MDisk (or MDisks, if creating multiple copies). The command fails if there are not enough sequential extents on the specified MDisk.

striped

This is the default virtualization type. If the **-vtype** parameter is not specified, **striped** is the default; all managed disks in the managed disk group are used to create the virtual disk. The striping is at an extent level; one extent from each managed disk in the group is used. For example, a managed disk group with 10 managed disks uses one extent from each managed disk, then it uses the 11th extent from the first managed disk, and so on.

If the **-mdisk** parameter is also specified, you can supply a list of managed disks to use as the stripe set. This can be two or more managed disks from the same managed disk group. The same circular algorithm is used across the striped set. However, a single managed disk can be specified more than once in the list. For example, if you enter **-mdisk 0:1:2:1**, the extents are from the following managed disks: 0, 1, 2, 1, 0, 1, 2, and so forth. All MDisks that are specified in the **-mdisk** parameter must be in the managed mode.

A capacity of 0 is allowed.

image This virtualization type allows image mode virtual disks to be created when a managed disk already has data on it, perhaps from a previrtualized subsystem. When an image mode virtual disk is created, it directly corresponds to the (previously unmanaged) managed disk that it was created from. Therefore, with the exception of space-efficient image mode VDIs, virtual disk logical block address (LBA) x equals managed disk LBA x . You can use this command to bring a nonvirtualized disk under the control of the cluster. After it is under the control of the cluster,

you can migrate the virtual disk from the single managed disk. When it is migrated, the virtual disk is no longer an image mode virtual disk.

You can add image mode VDIs to an already populated MDisk group with other types of VDIs, such as a striped or sequential.

Note: An image mode VDisk must be 512 bytes or greater. At least one extent is allocated to an image mode VDisk.

You must use the **-mdisk** parameter to specify an MDisk that has a mode of unmanaged. The **-fmtdisk** parameter cannot be used to create an image mode VDisk.

Note: If you create a mirrored VDisk from two image mode MDisks without specifying a **-capacity** value, the capacity of the resulting VDisk is the smaller of the two MDisks, and the remaining space on the larger MDisk is not accessible.

The command returns the IDs of the newly created VDisk.

Attention:

1. Do not create a VDisk in an offline I/O group. You must ensure that the I/O group is online before you create a VDisk to avoid any data loss. This applies in particular to recreating VDIs that are assigned the same object ID.
2. To create an image mode disk, you must already have a quorum disk in the cluster because an image mode disk cannot be used to hold quorum data. See “Creating a quorum disk” in the *IBM System Storage SAN Volume Controller Software Installation and Configuration Guide* for more details.
3. The command fails if either limit of 2048 VDIs per I/O Group or 8192 VDisk copies per cluster is reached.

Table 7 provides the relationship of the *rate* value to the data copied per second.

Table 7. Relationship between the rate value and the data copied per second

User-specified <i>rate</i> attribute value	Data copied/sec
1 - 10	128 KB
11 - 20	256 KB
21 - 30	512 KB
31 - 40	1 MB
41 - 50	2 MB
51 - 60	4 MB
61 - 70	8 MB
71 - 80	16 MB
81 - 90	32 MB
91 - 100	64 MB

An invocation example

```
svctask mkvdisk -mdiskgrp Group0 -size 0  
-iogrp 0 -vtype striped -mdisk mdisk1 -node 1
```

The resulting output

```
Virtual Disk, id [1], successfully created
```

An invocation example for creating an image mode VDisk

```
svctask mkvdisk -mdiskgrp Group0  
-iogrp 0 -vtype image -mdisk mdisk2 -node 1
```

The resulting output

Virtual Disk, id [2], successfully created

An invocation example for creating a new VDisk

```
svctask mkvdisk -mdiskgrp Group0 -size 0 -unit kb  
-iogrp 0 -vtype striped -mdisk mdisk1 -node 1 -udid 1234
```

The resulting output

Virtual Disk id [2], successfully created

An invocation example for creating a space-efficient VDisk

```
svctask mkvdisk -mdiskgrp Group0 -iogrp 0 -vtype striped  
-size 10 -unit gb -rsize 20% -autoexpand -grainsize 32
```

The resulting output

Virtual Disk id [1], successfully created

An invocation example for creating a mirrored image-mode VDisk

```
svctask mkvdisk -mdiskgrp Group0:Group0 -mdisk mdisk2:mdisk3  
-iogrp 0 -vtype image -copies 2
```

The resulting output

Virtual Disk id [1], successfully created

An invocation example for creating a mirrored VDisk

```
svctask mkvdisk -iogrp 0 -mdiskgrp 0:1 -size 500 -copies 2
```

The resulting output

Virtual Disk id [5], successfully created

mkvdiskhostmap

The `mkvdiskhostmap` command creates a new mapping between a virtual disk and a host, which makes the virtual disk accessible for I/O operations to the specified host.

Syntax

```
▶▶ svctask — — mkvdiskhostmap — — [ -force ] ▶▶  
▶ -host [ host_id ] [ host_name ] [ -scsi scsi_num_arg ] ▶▶  
▶ [ vdisk_name ] [ vdisk_id ] ▶▶
```

Parameters

-force

(Optional) Allows multiple VDisk-to-host assignments, which are not normally allowed.

-host *host_id* | *host_name*

(Required) Specifies the host to map the virtual disk to, either by ID or by name.

-scsi *scsi_num_arg*

(Optional) Specifies the SCSI LUN ID to assign to this virtual disk on the given host. The *scsi_num_arg* parameter contains the SCSI LUN ID that is assigned to the VDisk on the given host. You must check your host system for the next available SCSI LUN ID on the given HBA. If you do not specify the **-scsi** parameter, the next available SCSI LUN ID is provided to the host.

vdisk_name | *vdisk_id*

(Required) Specifies the name of the virtual disk that you want to map to the host, either by ID or by name.

Description

This command creates a new mapping between the virtual disk and the specified host. The virtual disk is presented to the host as if the disk is directly attached to the host. It is only after this command is processed, that the host can perform I/O transactions to the virtual disk.

Optionally, you can assign a SCSI LUN ID to the mapping. When the HBA in the host scans for devices that are attached to it, it discovers all virtual disks that are mapped to its fibre-channel ports. When the devices are found, each one is allocated an identifier (SCSI LUN ID). For example, the first disk found is usually SCSI LUN 1, and so on. You can control the order in which the HBA discovers virtual disks by assigning the SCSI LUN ID, as required. If you do not specify a SCSI LUN ID, the cluster automatically assigns the next available SCSI LUN ID, if any mappings already exist with that host. When you issue the `mkvdiskhostmap` command, the assigned SCSI LUN ID number is returned.

Some HBA device drivers will stop when they find a gap in the SCSI LUN IDs. For example:

- Virtual Disk 1 is mapped to Host 1 with SCSI LUN ID 1
- Virtual Disk 2 is mapped to Host 1 with SCSI LUN ID 2
- Virtual Disk 3 is mapped to Host 1 with SCSI LUN ID 4

When the device driver scans the HBA, it must stop after identifying virtual disks 1 and 2, because no SCSI LUN is mapped with ID 3. For optimal performance, ensure that the SCSI LUN ID allocation is contiguous.

You can create multiple VDisk assignments. Normally, multiple VDisk-to-host assignments are not used because corruption is likely to occur if more than one host can access a disk. However, in certain multiple path environments, such as in the IBM SAN File System, a VDisk must be mapped to more than one host. To map to more than one host, you must use the `mkvdiskhostmap` command with the **force** parameter. For example:

```
svctask mkvdiskhostmap -host host1 -force 4
svctask mkvdiskhostmap -host host2 -force 4
```

These commands create two host-to-VDisk mappings for VDisk 4 that map to host1 and host2. Omitting the **force** parameter causes the mapping to fail if that VDisk is already mapped to a host.

The command also fails if the host object (to which this mapping is being made) is not associated with the I/O group containing the VDisk.

An invocation example

```
svctask mkvdiskhostmap -host host1 -scsi 1 5
```

The resulting output

```
Virtual Disk to Host map, id [1], successfully created
```

recovervdisk

The **recovervdisk** command acknowledges VDisk data loss and brings the VDisk back online.

Syntax

```
svctask --recovervdisk [vdisk_name | vdisk_id]
```

Parameters

vdisk_name | *vdisk_id*

(Required) Specifies the virtual disk to recover.

Description

The specified VDisk, and all copies if mirrored, are recovered and brought back online. If the VDisk is space-efficient or has space-efficient copies, this command triggers the space-efficient repair process. If the VDisk is mirrored, the **recovervdisk** command triggers a resynchronization from a synchronized copy. The progress of the resynchronization can be monitored using the **svcinfolsvdisksyncprogress** command. The VDisk remains online during the resynchronization process.

The **recovervdisk** command also starts the repair of any space-efficient copies that have a `fast_write_state` of `corrupt`. The progress of the repair process can be monitored using the **svcinfolrepairsevdiskcopyprogress** command.

A VDisk that is still offline because it is being repaired following the **recovervdisk** command has a `fast_write_state` of `repairing`. The VDisk is brought online when the repair process is complete.

An invocation example

```
svctask recovervdisk vdisk17
```

The resulting output

```
No feedback
```

recovervdiskbycluster

The **recovervdiskbycluster** command acknowledges data loss for all VDIs in the cluster with a `fast_write_state` of corrupt and brings the VDIs back online.

Syntax

```
▶▶—svctask— —recovervdiskbycluster— ▶▶
```

Parameters

There are no parameters.

Description

All VDIs in the cluster that have a `fast_write_state` of corrupt; and all copies, if mirrored, are recovered and brought back online. If any of the VDIs are space-efficient or have space-efficient copies, the **recovervdiskbycluster** command triggers the space-efficient repair process. If VDIs are mirrored, the command triggers a resynchronization from a synchronized copy. The progress of the resynchronization can be monitored by using the **svcinfolsvdisksyncprogress** command. VDIs remain online during the resynchronization process.

If none of the VDIs in the cluster have a `fast_write_state` of corrupt, the **recovervdiskbycluster** command still starts the repair process for any corrupt copies of mirrored VDIs. The progress of the repair process can be monitored using the **svcinfolrepairsevdiskcopyprogress** command. If there are no corrupt VDIs or no repairs to copies are required, no error is returned.

VDIs that are still offline because they are being repaired following the **recovervdiskbycluster** command have a `fast_write_state` of repairing. VDIs are brought online when the repair process is complete.

An invocation example

```
svctask recovervdiskbycluster
```

The resulting output

No feedback

recovervdiskbyiogrp

The **recovervdiskbyiogrp** command acknowledges data loss for all VDIs in the specified I/O group with a `fast_write_state` of corrupt and brings the VDIs back online.

Syntax

```
▶▶—svctask— —recovervdiskbyiogrp— io_group_name  
io_group_id ▶▶
```

Parameters

io_group_name | *io_group_id*
(Required) Specifies the I/O group for virtual disk recovery.

Description

All VDIs in the specified I/O group that have a `fast_write_state` of corrupt; and all copies, if mirrored, are recovered and brought back online. If any of the VDIs are space_efficient or have space_efficient copies, the **recovervdiskbyiogrp** command triggers the space-efficient repair process. If VDIs are mirrored, the command triggers a resynchronization from a synchronized copy. The progress of the resynchronization can be monitored by using the **svcinfolsvdisksyncprogress** command. VDIs remain online during the resynchronization process.

If none of the VDIs 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 VDIs. The progress of the repair process can be monitored using the **svcinfolrepairsevdiskcopyprogress** command. If there are no corrupt VDIs or no repairs to copies are required, no error is returned.

VDIs that are still offline because they are being repaired following the **recovervdiskbyiogrp** command have a `fast_write_state` of repairing. VDIs are brought online when the repair process is complete.

An invocation example

```
svctask recovervdiskbyiogrp iogrp2
```

The resulting output

No feedback

repairsevdiskcopy

The `repairsevdiskcopy` command repairs the metadata on a space-efficient virtual disk.

Syntax

```
svctask repairsevdiskcopy [-copy 0 | 1] vdisk_name | vdisk_id
```

Parameters

-copy 0 | 1

(Optional) Specifies the VDisk copy to repair.

vdisk_name | *vdisk_id*

(Required) Specifies the virtual disk to repair.

Description

The `repairsevdiskcopy` command repairs the metadata on a space-efficient VDisk. Run this command only when you are directed by the Directed Maintenance Procedures or by IBM support.

Running the command automatically detects corrupted metadata. The command holds the VDisk offline during the repair, but does not prevent the disk from being moved between I/O groups.

If a repair operation completes successfully and the volume was previously offline because of corrupted metadata, the command brings the volume back online. The

only limit on the number of concurrent repair operations is the number of virtual disk copies in the configuration. Once started, a repair operation cannot be paused or canceled; the repair can only be ended by deleting the copy.

An invocation example

```
svctask repairsevdiskcopy vdisk8
```


The resulting output

No feedback

repairvdiskcopy

The **repairvdiskcopy** command detects and optionally, corrects any VDisk copies that are not identical.

Syntax



Parameters

-medium

(Optional) Converts sectors that contain different contents into virtual medium errors on the specified VDisk. This parameter cannot be used with the **-validate** and **-resync** parameters; you must enter one of the three parameters.

-resync

(Optional) Corrects sectors that contain different contents by copying contents from the primary VDisk copy to other copies on the specified VDisk. This parameter cannot be used with the **-medium** and **-resync** parameters; you must enter one of the three parameters.

-validate

(Optional) Reports the first difference found on synchronized online copies of the specified VDisk, on or after the specified **-startlba** value. This parameter cannot be used with the **-medium** and **-resync** parameters; you must enter one of the three parameters.

-startlba lba

(Optional) Specifies a starting Logical Block Address (LBA) on which to begin the command.

vdisk_name | vdisk_id

(Required) Specifies the virtual disk to repair. You must specify this parameter last on the command line.

Description

The **repairvdiskcopy** command detects and optionally, corrects any VDisk copies that are not identical. The results are logged to the SAN Volume Controller error log. The **-validate** parameter compares synchronized online copies of the specified

VDisk. The **-medium** parameter changes any sectors that are not identical into virtual medium errors. The **-resync** parameter copies any sectors that are not identical to the other VDisk copies. You must specify only one of the three parameters.

Attention:

1. Before you run the **repairvdiskcopy** command, ensure that all VDisk copies are synchronized.
2. Only one **repairvdiskcopy** command can run on a VDisk at a time. You must wait for the **repairvdiskcopy** command to complete processing before running the command again.
3. Once you start the **repairvdiskcopy** command, you cannot use the command to stop processing.
4. The primary copy of a mirrored VDisk cannot be changed while the **repairvdiskcopy -resync** command is running.

Use the **-startlba** parameter to specify a starting Logical Block Address (LBA). Enter an LBA value from 0 - full disk size minus one. The parameter logs the first error found and then stops the command. By repeating this parameter, you can collect all of the instances where the VDisk copies are not identical.

During **repairvdiskcopy** command operation, the VDisk remains online. The I/O and synchronization operations are allowed while the command is in progress.

The rate for the **repairvdiskcopy** command is controlled by the synchronization rate of the VDisk that is being repaired. To suspend the repair process, set the synchronization rate of the VDisk to **0** using the **chvdisk** command.

An invocation example

```
svctask repairvdiskcopy -resync vdisk8
```

The resulting output

No feedback

rmvdisk

The **rmvdisk** command deletes a virtual disk (VDisk).

Syntax

```
svctask -- rmvdisk -- -force vdisk_id  
vdisk_name
```

Parameters

-force

(Optional) Deletes the specified VDisk, even if mappings still exist between this virtual disk and one or more hosts. This parameter deletes any host-to-VDisk mappings and any FlashCopy mappings that exist for this VDisk. This parameter also deletes any Metro Mirror or Global Mirror relationships that exist for the specified VDisk. Data that is on the virtual disk is lost. Before you issue this command, ensure that the virtual disk and any data that resides on it are no longer required.

vdisk_id | *vdisk_name*

Specifies the name of the virtual disk to delete, either by ID or by name.

Description

This command deletes an existing managed mode virtual disk or an existing image mode virtual disk. The extents that made up this virtual disk are returned to the pool of free extents that are available on the managed disk group, if the VDisk is in managed mode.

Attention: Any data that was on the virtual disk is lost. Before you issue this command, ensure that the virtual disk (and any data that resides on it) is no longer required.

Deleting a managed mode virtual disk

When you use this command to delete a managed mode virtual disk, all the data on the virtual disk is deleted. The extents that make up the virtual disk are returned to the pool of free extents that are available in the managed disk group.

If host mappings exist for the virtual disk, or if any FlashCopy mappings would be affected, the deletion fails. You can use the **-force** parameter to force the deletion. If you use the **-force** parameter, mappings that have the virtual disk as source or target are deleted, other mappings in a cascade might be stopped, and then the virtual disk is deleted. The **-force** parameter also deletes any Metro Mirror or Global Mirror relationships that exist for the specified VDisk.

If the virtual disk is in the process of migrating to an image mode virtual disk (using the **svctask migratetoimage** command), the deletion fails unless you use the **-force** parameter. If you use the **-force** parameter, the migration is halted and then the virtual disk is deleted. Before you issue this command, ensure that the virtual disk (and any data that resides on it) is no longer required.

Deleting an image mode virtual disk

If the VDisk is mirrored and one or both copies is in image mode, you must first wait for all fast-write data to be moved to the controller logical unit. This ensures that the data on the controller is consistent with the data on the image mode virtual disk before the VDisk is deleted. This process can take several minutes to complete, and is indicated by the *fast_write_state* state of the virtual disk being **empty**. If the **-force** parameter is specified, the fast-write data is discarded and the virtual disk is deleted immediately; the data on the controller logical unit is left inconsistent and unusable. If the copies are not synchronized, you must use the **-force** parameter.

If you run the command while data is in the cache, SVC attempts to move the data out of the cache; this process can time out, however.

If there are any virtual medium errors on the virtual disk, the command fails. You can force the deletion by using the **-force** parameter; however, this can cause data integrity problems.

Note: A virtual medium error occurs when you copy data from one disk (the source) to another (the target). Reading the source indicates that there is a medium error. At that moment, you must have two identical copies of data

and you must then simulate a medium error on the target disk. You can simulate a medium error on the target disk by creating a virtual medium error on the target disk.

If FlashCopy mappings or host mappings exist for the virtual disk, the deletion fails unless you use the **-force** parameter. If you use the **-force** parameter, mappings are deleted and the virtual disk is deleted. If there is any data that is not staged in the fast write cache for this virtual disk, the deletion of the virtual disk fails. When the **-force** parameter is specified, any data that is not staged in the fast write cache is deleted. Deleting an image mode virtual disk causes the managed disk that is associated with the virtual disk to be removed from the managed disk group. The mode of the managed disk is returned to “unmanaged.”

An invocation example

```
svctask rmvdisk -force vdisk5
```

The resulting output

No feedback

rmvdiskcopy

The **rmvdiskcopy** command removes a VDisk copy from a VDisk.

Syntax

```
▶▶ svctask — rmvdiskcopy — --copy — copy_id — -force —————▶
```



```
▶ vdisk_name —————▶  
  vdisk_id
```

Parameters

-copy *copy_id*

(Required) Specifies the ID of the copy to delete.

-force

(Optional) Forces the deletion of the last synchronized copy of a VDisk, which deletes the entire VDisk. The parameter also forces the deletion of a nonmirrored VDisk, a copy that is migrating to image mode, or an image-mode copy that has virtual medium errors.

vdisk_name | *vdisk_id*

(Required) Specifies the virtual disk to delete the copy from. You must specify this parameter last on the command line.

Description

The **rmvdiskcopy** command deletes the specified copy from the specified VDisk. The command fails if all other copies of the VDisk are not synchronized; in this case, you must specify the **-force** parameter, delete the VDisk, or wait until the copies are synchronized.

An invocation example

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

The resulting output

No feedback

rmvdiskhostmap

The **rmvdiskhostmap** command deletes an existing virtual disk-to-host mapping; the virtual disk is no longer accessible for I/O transactions on the given host.

Syntax

```
▶▶ svctask — — rmvdiskhostmap — — -host ———— [ host_id ] —————▶▶
                                     [ host_name ]
▶ [ vdisk_id ] —————▶▶
  [ vdisk_name ]
```

Parameters

-host *host_id* | *host_name*

(Required) Specifies the host that you want to remove from the map with the virtual disk, either by ID or by name.

vdisk_id | *vdisk_name*

(Required) Specifies the name of the virtual disk that you want to remove from the host mapping, either by ID or by name.

Description

This command deletes an existing mapping between the specified virtual disk and the host. This effectively stops the virtual disk from being available for I/O transactions on the given host.

This command also deletes a SCSI or persistent reservation that a host has on a VDisk. Once the reservation is removed, a new host is allowed to access the VDisk in the future because the original host no longer has access.

Use caution when you process this command because to the host, it seems as if the virtual disk has been deleted or is offline.

An invocation example

```
svctask rmvdiskhostmap -host host1 vdisk8
```

The resulting output


No feedback

shrinkvdisksize

The **shrinkvdisksize** command reduces the size of a VDisk by the specified capacity.

Syntax

```
▶▶ svctask — — shrinkvdisksize — [ -size disk_size ] —————▶▶
                                     [ -rsize disk_size ]
```



Parameters

-size *disk_size*

(Required) Specifies the size reduction for the designated virtual disk. The **-size** parameter cannot be used with the **-rsize** parameter. You must specify either **-size** or **-rsize**.

-rsize *disk_size*

(Optional) Reduces the real size of a space-efficient VDisk by the specified amount. Specify the *disk_size* value using an integer. Specify the units for a *disk_size* integer using the **-unit** parameter; the default is MB. The **-rsize** value can be greater than, equal to, or less than the size of the VDisk. You must specify either the **-size** parameter or the **-rsize** parameter.

-copy *id*

(Optional) Specifies the copy to change the real capacity for. You must also specify the **-rsize** parameter. If the **-copy** parameter is not specified, all copies of the VDisk are reduced. This parameter is required if the VDisk is mirrored and only one copy is space-efficient.

-unit **b | kb | mb | gb | tb | pb**

(Optional) Specifies the data units to be used in conjunction with the value that is specified by the **-size** parameter.

vdisk_name | *vdisk_id*

(Required) Specifies the virtual disk that you want to modify, either by ID or by name.

Description

The `shrinkvdisksize` command reduces the capacity that is allocated to the particular virtual disk by the amount that you specify. You cannot shrink the real size of a space-efficient volume below its used size. All capacities, including changes, must be in multiples of 512 bytes. An entire extent is reserved even if it is only partially used. The default capacity units are MB.

The command can be used to shrink the physical capacity that is allocated to a particular VDisk by the specified amount. The command can also be used to shrink the virtual capacity of a space-efficient VDisk without altering the physical capacity assigned to the VDisk. To change the capacity of a non-space-efficient disk, use the **-size** parameter. To change the real capacity of a space-efficient disk, use the **-rsize** parameter. To change the virtual capacity of a space-efficient disk, use the **-size** parameter.

VDisks can be reduced in size, if required.

When the virtual size of a space-efficient VDisk is changed, the warning threshold is automatically scaled to match. The new threshold is stored as a percentage.

To run the `shrinkvdisksize` command on a mirrored VDisk, all copies of the VDisk must be synchronized.

Attention: If the VDisk contains data that is being used, do not shrink the VDisk without backing up the data first.

The cluster arbitrarily reduces the capacity of the VDisk by removing a partial, one or more extents from those allocated to the VDisk. You cannot control which extents are removed and so you cannot assume that it is unused space that is removed.

Attention:

1. If the virtual disk contains data, do not shrink the disk.
2. Some operating systems or file systems use what they consider to be the outer edge of the disk for performance reasons. This command can shrink FlashCopy target virtual disks to the same capacity as the source.
3. Before you shrink a VDisk, validate that the VDisk is not mapped to any host objects. If the VDisk is mapped, data is displayed. You can determine the exact capacity of the source or master VDisk by issuing the `svcinflsvdisk -bytes vdiskname` command. Shrink the VDisk by the required amount by issuing the `svctask shrinkvdisksize -size disk_size -unit b | kb | mb | gb | tb | pb vdisk_name | vdisk_id` command.

An invocation example

To decrease the capacity of `vdisk1` by 2 KB, enter:
`svctask shrinkvdisksize -size 2048 -unit b vdisk1`

The resulting output

No feedback

An invocation example

To decrease the capacity of `vdisk2` by 100 MB, enter:
`svctask shrinkvdisksize -size 100 -unit mb vdisk2`

The resulting output

No feedback

An invocation example

To decrease the real capacity of space-efficient VDisk `vdisk3` by 100 MB without changing its virtual capacity, enter:
`svctask shrinkvdisksize -rsize 100 -unit mb vdisk3`

The resulting output

No feedback

An invocation example

To decrease the real capacity of space-efficient VDisk copy ID 1 of mirrored VDisk `vdisk4` by 100 MB, enter:
`svctask shrinkvdisksize -rsize 100 -unit mb -copy 1 vdisk4`

The resulting output

No feedback

An invocation example

To decrease the virtual capacity of space-efficient VDisk vdisk5 by 1 GB without changing its real capacity, enter:

```
svctask shrinkvdisksize -size 1 -unit gb vdisk5
```

The resulting output

No feedback

splitvdiskcopy

The **splitvdiskcopy** command creates a separate VDisk from a synchronized copy of a mirrored VDisk.

Syntax

```
svctask splitvdiskcopy -copy id  
-iogrp io_group_id | io_group_name  
-node node_id | node_name -name new_name  
-cache readwrite | none -udid udid -force  
vdisk_name  
vdisk_id
```

Parameters

-copy *id*

(Required) Specifies the ID of the copy to split.

-iogrp *io_group_id* | *io_group_name*

(Optional) Specifies the I/O group to add the new virtual disk to. The default is the I/O group of the specified VDisk.

-node *node_id* | *node_name*

(Optional) Specifies the preferred node ID or the name for I/O operations to this virtual disk. You can use the **-node** parameter to specify the preferred access node.

-name *new_name*

(Optional) Assigns a name to the new virtual disk.

-cache *readwrite* | *none*

(Optional) Specifies the caching options for the new virtual disk. Enter **readwrite** or **none**; the default is **readwrite**.

-udid *udid*

(Optional) Specifies the *udid* for the new VDisk. The *udid* is a required identifier for OpenVMS hosts; no other hosts use this parameter. Supported

values are a decimal number 0 - 32 767, or a hexadecimal number 0 - 0x7FFF. A hexadecimal number must be preceded by **0x**; for example, **0x1234**. The default *udid* value is **0**.

-force

(Optional) Allows the split to proceed even when the specified copy is not synchronized, or even when the cache flush is likely to fail. The newly created VDisk might not be consistent.

Description

The **splitvdiskcopy** command creates a new VDisk in the specified I/O Group from a copy of the specified VDisk. If the copy that you are splitting is not synchronized, you must use the **-force** parameter. The command fails if you are attempting to remove the only synchronized copy. To avoid this, wait for the copy to synchronize or split the unsynchronized copy from the VDisk by using the **-force** parameter. You can run the command when either VDisk copy is offline.

An invocation example

```
svctask splitvdiskcopy -copy 1 vdisk8
```

The resulting output

```
Virtual Disk, id [1], successfully created.
```

Chapter 13. Managed disk group commands

The following commands enable you to work with managed disk group options with the SAN Volume Controller.

addmdisk

The `addmdisk` command adds one or more managed disks to an existing managed disk group.

Syntax

```
svctask -- addmdisk -- -mdisk [ mdisk_id_list ] [ mdisk_name_list ]
[ mdisk_group_id ] [ mdisk_group_name ]
```

Parameters

-mdisk *mdisk_id_list* | *mdisk_name_list*

(Required) Specifies one or more managed disk IDs or names to add to the group.

mdisk_group_id | *mdisk_group_name*

(Required) Specifies the ID or name of the managed disk group to add the disks to. When an MDisk is added, the warning threshold for the MDisk group is automatically scaled.

Description

This command adds the managed disks that you specify to the group. The disks can be specified in terms of the managed disk ID or the managed disk name.

The managed disks must be in unmanaged mode. Disks that already belong to a group cannot be added to another group until they have been deleted from their current group. You can delete a managed disk from a group under the following circumstances:

- If the managed disk does not contain any extents in use by a virtual disk
- If you can first migrate the extents in use onto other free extents within the group.

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

An invocation example

```
svctask addmdisk -mdisk mdisk13:mdisk14 Group0
```

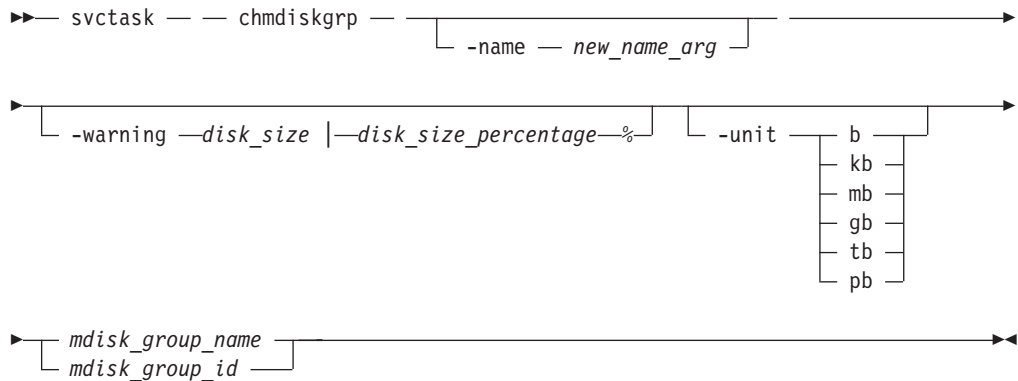
The resulting output

No feedback

chmdiskgrp

Use the **chmdiskgrp** command to modify the name that is assigned to a managed disk (MDisk) group or to set the warning threshold for the MDisk group.

Syntax



Parameters

-name *new_name_arg*

(Optional) Specifies the new name of the managed disk group.

-warning *disk_size* | *disk_size_percentage%*

(Optional) Sets a threshold at which a warning is generated. The warning is generated the first time that the threshold is exceeded by the used-disk capacity in the MDisk group. You can specify a *disk_size* integer, which defaults to megabytes (MB) unless the **-unit** parameter is specified; or you can specify a *disk_size%*, which is a percentage of the MDisk group size. To disable warnings, specify **0** or **0%**.

-unit **b** | **kb** | **mb** | **gb** | **tb** | **pb**

(Optional) Specifies the data units for the **-warning** parameter.

mdisk_group_id | *mdisk_group_name*

(Required) Specifies the ID or name of the managed disk group to modify.

Description

This command modifies the name, or label, assigned to a given managed disk group. Subsequently, you can use the new name to refer to the managed disk group.

The command can also be used to set the warning threshold for the managed disk group. The warning threshold is the threshold at which a warning is generated when it is exceeded by the used-disk capacity in the MDisk group.

An invocation example

```
svctask chmdiskgrp -name testmdiskgrp Group0
```

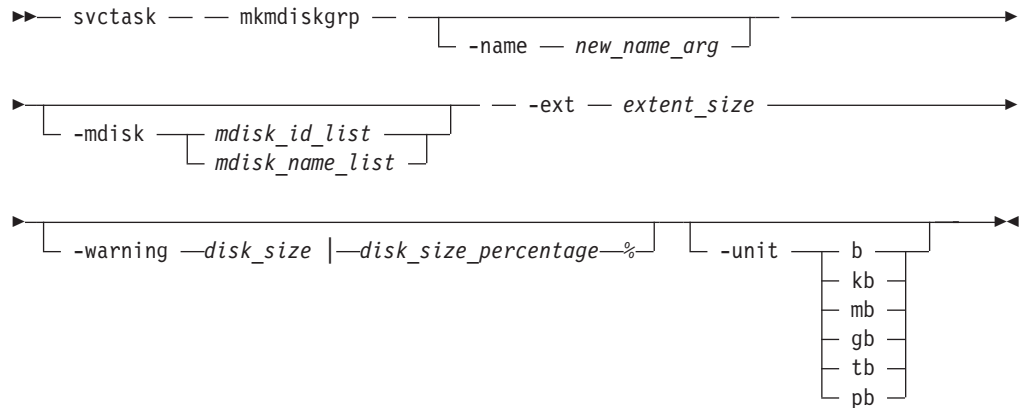
The resulting output

No feedback

mkmdiskgrp

The `mkmdiskgrp` command creates a new managed disk group.

Syntax



Parameters

-name *new_name_arg*

(Optional) Specifies a name to assign to the new group.

-mdisk *mdisk_id_list* | *mdisk_name_list*

(Optional) Specifies a colon-separated list of managed disk IDs or names to add to the group. You can create an empty MDisk group by not specifying the **-mdisk** parameter.

-ext *extent_size*

(Required) Specifies the size of the extents for this group in MB. The *extent_size* parameter must be one of the following values: **16**, **32**, **64**, **128**, **256**, **512**, **1024**, or **2048** (MB).

-warning *disk_size* | *disk_size_percentage%*

(Optional) Generates a warning when the used disk capacity in the MDisk group first exceeds the specified threshold. You can specify a *disk_size* integer, which defaults to megabytes (MB) unless the **-unit** parameter is specified; or you can specify a *disk_size%*, which is a percentage of the MDisk group size. To disable warnings, specify **0** or **0%**. The default value is **0**.

-unit **b** | **kb** | **mb** | **gb** | **tb** | **pb**

(Optional) Specifies the data units for the **-warning** parameter.

Description

The `mkmdiskgrp` command creates a new managed disk group and assigns the group name if specified. The ID of the new group is returned if the command is successful. Managed disk groups are collections of managed disks. Each group is divided into chunks, called extents, which are used to create virtual disks.

Optionally, you can specify a list of managed disks that will be added to this group. These managed disks cannot belong to another group, and they must have a mode of unmanaged. Use the `svcinfo lsmdiskcandidate` command to get a list of suitable candidates.

Each managed disk that is a member of this group is split into extents. The storage that is available on these disks is added to a pool of extents that is available in this group. When a virtual disk is created from this group, free extents from the pool are used, in accordance with the policy used when the virtual disk was first created.

All managed disks subsequently added to this group are split into extents of the same size as the size that is assigned to the group.

When choosing an extent size, take into account the amount of storage you want to virtualize in this group. The system maintains a mapping of extents between virtual disks and managed disks. The cluster can only manage a finite number of extents (4 194 304). One cluster can virtualize the following number of extents:

- 64 TB – if all managed disk groups have extent sizes of 16 MB.
- 2 PB – if all managed disk groups have extent sizes of 512 MB.

Important: The extent size for the MDisk group can also limit VDisk size. Consider the maximum VDisk size you want to use when creating MDisk groups. SAN Volume Controller MDisk groups information provides the maximum VDisk capacity for each extent size.

Note: When an image mode VDisk is created, the MDisk group increases in capacity by the size of the image mode VDisk (not the MDisk capacity), because the image mode VDisk might be smaller than the MDisk itself. If an extent is migrated from the image mode VDisk or MDisk to elsewhere in the group, the VDisk becomes a striped VDisk (no longer image mode). At this point the available capacity might increase, because the extra capacity available on the MDisk (for example, the capacity that was not part of the image mode VDisk) becomes available.

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

An invocation example

```
svctask mkmdiskgrp -mdisk mdisk13 -ext 512
```

The resulting output

```
MDisk Group, id [1], successfully created
```

An invocation example

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

The resulting output


```
MDisk Group, id [0], successfully created
```

rmmdisk

The **rmmdisk** command deletes a managed disk (MDisk) from a managed disk group.

Syntax

```
svctask -- rmmdisk -- -mdisk mdisk_id_list  
mdisk_name_list
```



Parameters

-mdisk *mdisk_id_list* | *mdisk_name_list*

(Required) Specifies one or more managed disk IDs or names to delete from the group.

-force

(Optional) Migrates data on the specified disks to other disks in the group. The command completes asynchronously if **-force** is specified.

mdisk_group_id | *mdisk_group_name*

(Required) Specifies the ID or name of the managed disk group to delete the disks from. The warning threshold for an MDisk group is automatically scaled when MDisks are deleted.

Description

This command attempts to remove the managed disk or disks from the group.

Deleting a managed disk from a group can only be done if the managed disk does not contain any extents in use by a virtual disk. If there are extents in use and you do not supply the force flag, the command fails.

Attention: If this disk being removed has already been powered down, removed, or is experiencing a power outage, the migration is pending and does not complete until the MDisk comes back online. The MDisk is not removed from the list of MDisks that are contained in the group.

If the disk has been deliberately removed, the only method of removing the MDisk is to remove the entire group itself.

Ensure that you do not destroy any controller LUNs until you have deleted them from the MDisk group that they belong to.

The **rmmdisk** command fails if there are insufficient free extents on other disks in the mdisk group for the duration of the command. To avoid this problem, do not issue new commands that use extents until **rmmdisk** processing is completed.

If you do specify the force flag, an attempt will be made to migrate the extents that are in use onto other free extents within the group. If there are not enough free extents in the group, the command will fail even if the force flag is specified.

To delete the disks from the group, you have the following options:

- You can delete the virtual disk that is using the extents specified on the managed disk.
- You can add more managed disks to the group, rerun the command and specify the **-force** parameter.

When data is being migrated from the managed disk, it might take some time for the command to complete. The command itself will return with a success code, notifying you that migration is in progress. An event is logged when the migration is complete and the disk is deleted from the group at this time. You can also check the progress of any active migrations by running the **svcinfo lsmigrate** command.

If the **-force** parameter is used, the **rmmdisk** command fails if the target or source VDisk is offline, or if there is insufficient quorum-disk space to store the metadata. Correct the offline or quorum disk condition and try reissuing the command.

An invocation example

```
svctask rmmdisk -mdisk mdisk12 -force Group3
```

The resulting output

No feedback

rmmdiskgrp

The **rmmdiskgrp** command deletes a managed disk group so that there is no possibility to recover it.

Syntax

```
▶▶ svctask — — rmmdiskgrp — [ -force ] [ mdisk_group_id | mdisk_group_name ] ▶▶
```

Parameters

-force

(Optional) Specifies that all virtual disks and virtual disk-to-host mappings be deleted.

Attention: Use this parameter with extreme caution. When you use this parameter, all managed disks in the group are removed and the group itself is deleted.

mdisk_group_id | mdisk_group_name

(Required) Specifies the ID or name of the managed disk group that is to be deleted.

Description

The **rmmdiskgrp** command deletes the specified managed disk group. The **-force** parameter is required if there are virtual disks that have been created from this group or if there are managed disks in the group. Otherwise, the command fails.

Deleting a managed disk group is essentially the same as deleting a cluster or part of a cluster, because the managed disk group is the central point of control of virtualization. Because virtual disks are created using available extents in the group, mapping between virtual disk extents and managed disk extents is controlled based on the group.

The command deletes all VDisk copies in the specified MDisk group. If the VDisk has no remaining synchronized copies in other MDisk groups, the VDisk is also deleted.

Attention:

1. This command partially completes asynchronously. All virtual disks, host mappings, and Copy Services relationships are deleted before the command completes. The deletion of the managed disk group then completes asynchronously.
2. Before you issue the command, ensure that you want to delete all mapping information; data that is contained on virtual disks cannot be recovered after the managed disk group has been deleted.

In detail, if you specify the **-force** parameter and the virtual disks are still using extents in this group, the following actions are initiated or occur:

- The mappings between that disk and any host objects and the associated Copy Services relationships are deleted.
- If the virtual disk is a part of a FlashCopy mapping, the mapping is deleted.

Note: If the mapping is not in the **idle_or_copied** or **stopped** states, the mapping is force-stopped and then deleted. Force-stopping the mapping might cause other FlashCopy mappings in the cluster to also be stopped. See the description for the **-force** parameter in the **stopfcmap** command for additional information.

- Any virtual disk that is in the process of being migrated into or out of the managed disk group is deleted. This frees up any extents that the virtual disk was using in another managed disk group.
- Virtual disks are deleted without first flushing the cache. Therefore, the storage controller LUNs that underlie any image mode MDisks might not contain the same data as the image mode VDisk prior to the deletion.
- If there are managed disks in the group, all disks are deleted from the group. They are returned to the unmanaged state.
- The group is deleted.

Attention: If you use the **-force** parameter to delete all the managed disk groups in your cluster, you are returned to the processing state where you were after you added nodes to the cluster. All data that is contained on the virtual disks is lost and cannot be recovered.

An invocation example

```
svctask rmmdiskgrp -force Group3
```

The resulting output

```
No feedback
```

Chapter 14. Managed disk commands

The following commands enable you to work with managed disk options with the SAN Volume Controller.

If the cluster detects an MDisk, it automatically adds it to the list of known MDisks. If you subsequently delete the RAID that corresponds to the MDisk, the cluster only deletes the MDisk from the list if the MDisk is offline and it has a mode of unmanaged (it does not belong to an MDisk group).

applymdisksoftware

Use the `applymdisksoftware` command to upgrade the firmware on a specified managed disk (MDisk). Only solid-state drives (SSDs) can be specified.

Syntax

```
svctask -- applymdisksoftware -- -file name ----->
```

```
-type firmware  
fpga -force mdisk_id  
mdisk_name ----->
```

Parameters

-file *name*

(Required) Specifies the name of the firmware upgrade file. The file name must exist in the `/home/admin/upgrade/` directory.

-type *firmware* | *fpga*

(Required) Specifies the type of software that should be applied to the MDisk. The default is `firmware`. With the `fpga` option, the MDisk can remain offline for up to 20 minutes.

Attention: Only use the `fpga` option, which upgrades Field Programmable Gate Array (FPGA) firmware, under the direction of an IBM service representative.

-force

(Optional) Specifies that the upgrade should proceed even if VDisks could go offline.

mdisk *id* | *name*

(Required) Specifies the ID or name of the MDisk to receive the firmware upgrade.

Description

The `applymdisksoftware` command starts an upgrade of the firmware on a specified MDisk. When the command is run, the software image is loaded to the MDisk and the upgrade file is copied to all other online nodes in the cluster.

If the upgrade could cause any VDIs to go offline, the **force** parameter is required. For example, a firmware update to a managed MDisk requires the **force** parameter.

An invocation example

```
svctask applydisksoftware -file SSDsoftware -type firmware mdisk1
```

The resulting output

No feedback

chmdisk

Use the **chmdisk** command to modify the name of a managed disk (MDisk).

Syntax

```
svctask -- chmdisk -- -name -- new_name_arg -- [mdisk_id | mdisk_name]
```

Parameters

-name *new_name_arg*

(Required) Specifies the new name to be applied to the managed disk.

mdisk_id | *mdisk_name*

(Required) Specifies the ID or name of the managed disk to modify.

Description

This command modifies the name, or label, that is assigned to a given managed disk. You can subsequently use the new name to refer to the managed disk.

An invocation example

```
svctask chmdisk -name testmdisk mdisk0
```

The resulting output

No feedback

includemdisk

Use the **includemdisk** command to include a disk that has been excluded by the cluster.

Syntax

```
svctask -- includemdisk -- [mdisk_id | mdisk_name]
```

Parameters

mdisk_id | *mdisk_name*

(Required) Specifies the ID or name of the managed disk to add back into the cluster.

Description

The specified managed disk is included in the cluster.

You might exclude a disk from the cluster because of multiple I/O failures. These failures might be caused by noisy links. Once a fabric-related problem has been fixed, the excluded disk can be added back into the cluster.

Running this command against an MDisk might change its state, whether the state is reported as excluded.

Note: If an MDisk is in the excluded state, is offline, and does not belong to an MDisk group, issuing an include command for this MDisk results in the MDisk record being deleted from the cluster.

An invocation example

```
svctask includemdisk mdisk5
```

The resulting output

No feedback

setquorum

Use the setquorum command to change the managed disks (MDisks) that are assigned as quorum candidate disks.

Syntax

```
svctask -- setquorum -- -quorum { 0 | 1 | 2 } [-active]
      { mdisk_id | mdisk_name }
```

Parameters

-quorum 0 | 1 | 2

(Required) Specifies the quorum index.

-active

(Optional) Makes the specified MDisk the active quorum disk. To identify the quorum disk that is active, use the svcinfo lsquorum command.

mdisk_id | mdisk_name

(Required) Specifies the ID or name of the MDisk to assign as a quorum disk.

Description

The setquorum command sets the MDisk to the specified quorum index. This command is not synchronous, but usually takes only a few seconds to complete. In some situations it can take several minutes.

The cluster uses the quorum disk as a tie breaker when exactly half of the nodes that were previously a member of the cluster is present.

The use of a quorum disk allows the cluster to manage a SAN fault that splits the cluster exactly in half. One half of the cluster continues to operate and the other half stops until SAN connectivity is restored.

There is only one quorum disk; however, the cluster uses three disks as quorum candidate disks. The cluster selects the actual quorum disk from the pool of quorum candidate disks. The quorum candidate disks also hold a copy of important cluster metadata. Just over 256 MB is reserved for this purpose on each quorum candidate disk. The number of extents this reservation requires depends on the extent size for the managed disk group containing the MDisk. Table 8 provides the number of extents reserved for quorum use by extent size.

Table 8. 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

When you issue this command, the MDisk that currently is assigned the quorum index number is set to a nonquorum disk. The cluster automatically assigns quorum indexes.

You can set the active quorum disk with the **active** parameter. This can be useful in a split-site cluster configuration to ensure that the most highly-available quorum disk is used.

An invocation example

```
svctask setquorum -quorum 2 mdisk7
```

The resulting output

```
No feedback
```

triggermdiskdump

Use the `triggermdiskdump` command to collect error data from the specified managed disk (MDisk). Only solid-state drives (SSDs) can be specified.

Syntax

```
▶▶ svctask -- triggermdiskdump [mdisk_id | mdisk_name] ▶▶
```

Parameters

mdisk *id* | *name*

(Required) Specifies the ID or name of the MDisk to have its dump triggered.

Description

The `triggermdiskdump` command exports a copy of the SSD internal log to a disk dump in the `/dumps/mdisk` directory. When the command completes successfully, the node that contains the MDisk has a new file in this directory.

An invocation example

```
svctask triggermdiskdump 0
```

The resulting output

No feedback

Chapter 15. FlashCopy commands

The following commands enable you to work with FlashCopy methods and functions with the SAN Volume Controller.

chfconsistgrp

The **chfconsistgrp** command changes the name of a consistency group or marks the group for auto-deletion.

Syntax

```
svctask -- chfconsistgrp -- [-name new_name_arg] --  
[-autodelete on | off] [fc_consist_group_id | fc_consist_group_name]
```

Parameters

-name *new_name_arg*

(Optional) Specifies the new name to assign to the consistency group.

-autodelete **on** | **off**

(Optional) Deletes the consistency group when the last mapping that it contains is deleted or removed from the consistency group.

fc_consist_group_id | *fc_consist_group_name*

(Required) Specifies the ID or existing name of the consistency group that you want to modify.

Description

The **chfconsistgrp** command changes the name of a consistency group, marks the group for auto-deletion, or both.

An invocation example

```
svctask chfconsistgrp -name testgrp1 fcconsistgrp1
```

The resulting output

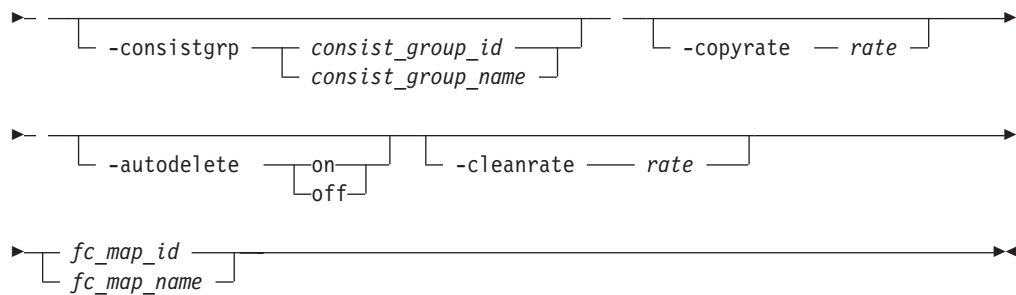
No feedback

chfcmap

The **chfcmap** command modifies attributes of an existing mapping.

Syntax

```
svctask -- chfcmap -- [-name new_name_arg] [-force]
```



Parameters

-name *new_name_arg*

(Optional) Specifies the new name to assign to the mapping. The **-name** parameter cannot be used with any other optional parameters.

-force

(Optional) Specifies that the mapping be modified to a stand-alone mapping (equivalent to creating the mapping without a consistency group ID). You cannot specify the **-force** parameter with the **-consistgrp** parameter.

-consistgrp *consist_group_id* | *consist_group_name*

(Optional) Specifies the consistency group for which you want to modify the mapping. You cannot specify the **-consistgrp** parameter with the **-force** parameter.

Note: The consistency group cannot be modified if the specified consistency group is in the **preparing**, **prepared**, **copying**, **suspended**, or **stopping** state.

-copyrate *rate*

(Optional) Specifies the copy rate. The *rate* value can be 0 - 100. The default value is 50. A value of 0 indicates no background copy process. For the supported **-copyrate** values and their corresponding rates, see Table 9 on page 143.

-autodelete on | **off**

(Optional) Specifies that the autodelete function be turned on or off for the specified mapping. When you specify the **-autodelete on** parameter, you are deleting a mapping after the background copy completes. If the background copy is already complete, the mapping is deleted immediately.

-cleanrate *rate*

(Optional) Sets the cleaning rate for the mapping. The *rate* value can be 0 - 100. The default value is 50.

fc_map_id | *fc_map_name*

(Required) Specifies the ID or name of the mapping to modify. Enter the ID or name last on the command line.

Description

The **svctask chfcmap** command modifies attributes of an existing mapping.

Attention: You must enter the *fc_map_id* | *fc_map_name* last on the command line.

If you have created several FlashCopy mappings for a group of VDisks that contain elements of data for the same application, you can assign these mappings

to a single FlashCopy consistency group. You can then issue a single prepare command and a single start command for the whole group, for example, so that all of the files for a particular database are copied at the same time.

The **copyrate** parameter specifies the copy rate. If **0** is specified, background copy is disabled. The **cleanrate** parameter specifies the rate for cleaning the target VDisk. The cleaning process is only active if the mapping is in the **copying** state and the background copy has completed, the mapping is in the **copying** state and the background copy is disabled, or the mapping is in the **stopping** state. You can disable cleaning when the mapping is in the **copying** state by setting the **cleanrate** parameter to **0**. If the **cleanrate** is set to **0**, the cleaning process runs at the default rate of **50** when the mapping is in the **stopping** state to ensure that the stop operation completes.

Table 9 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 9. Relationship between the rate, data rate and grains per second values

User-specified <i>rate</i> attribute value	Data copied/sec	256 KB grains/sec	64 KB grains/sec
1 - 10	128 KB	0.5	2
11 - 20	256 KB	1	4
21 - 30	512 KB	2	8
31 - 40	1 MB	4	16
41 - 50	2 MB	8	32
51 - 60	4 MB	16	64
61 - 70	8 MB	32	128
71 - 80	16 MB	64	256
81 - 90	32 MB	128	512
91 - 100	64 MB	256	1024

An invocation example

```
svctask chfcmap -name testmap 1
```

The resulting output

No feedback

mkfcconsistgrp

The **mkfcconsistgrp** command creates a new FlashCopy consistency group and identification name.

Syntax

```

▶▶▶ svctask — — mkfcconsistgrp — — [ -name — consist_group_name ]
▶▶▶
▶▶▶ [ -autodelete ]
▶▶▶

```

Parameters

-name *consist_group_name*

(Optional) Specifies a name for the consistency group. If you do not specify a consistency group name, a name is automatically assigned to the consistency group. For example, if the next available consistency group ID is id=2, the consistency group name is fccstgrp2.

-autodelete

(Optional) Deletes the consistency group when the last mapping that it contains is deleted or removed from the consistency group.

Description

This command creates a new consistency group and identification name. The ID of the new group is displayed when the command process completes.

If you have created several FlashCopy mappings for a group of VDisks that contain elements of data for the same application, you might find it convenient to assign these mappings to a single FlashCopy consistency group. You can then issue a single prepare command and a single start command for the whole group, for example, so that all of the files for a particular database are copied at the same time.

An invocation example

```
svctask mkfcconsistgrp
```

The resulting output

```
FlashCopy Consistency Group, id [1], successfully created
```

mkfcmap

The **mkfcmap** command creates a new FlashCopy mapping, which maps a source VDisk to a target VDisk for subsequent copying.

Syntax

```
svctask -- mkfcmap -- -source [ src_vdisk_id ] [ src_vdisk_name ]
-- -target [ target_vdisk_id ] [ target_vdisk_name ] [ -name new_name_arg ]
[ -consistgrp [ consist_group_id ] [ consist_group_name ] ] [ -copyrate rate ]
[ -autodelete ] [ -grainsize 64 | 256 ] [ -incremental ]
[ -cleanrate rate ] [ -iogrp [ iogroup_name ] [ iogroup_id ] ]
```

Parameters

- source** *src_vdisk_id* | *src_vdisk_name*
(Required) Specifies the ID or name of the source VDisk.
- target** *target_vdisk_id* | *target_vdisk_name*
(Required) Specifies the ID or name of the target VDisk.
- name** *new_name_arg*
(Optional) Specifies the name to assign to the new mapping.
- consistgrp** *consist_group_id* | *consist_group_name*
(Optional) Specifies the consistency group to add the new mapping to. If you do not specify a consistency group, the mapping is treated as a stand-alone mapping.
- copyrate** *rate*
(Optional) Specifies the copy rate. The *rate* value can be 0 - 100. The default value is 50. A value of 0 indicates no background copy process. For the supported **-copyrate** values and their corresponding rates, see Table 10 on page 146.
- autodelete**
(Optional) Specifies that a mapping be deleted when the background copy completes. The default, which applies if this parameter is not entered, is that **autodelete** is set to off.
- grainsize** **64** | **256**
(Optional) Specifies the grain size for the mapping. The default value is 256. Once set, this value cannot be changed.
- incremental**
(Optional) Marks the FlashCopy mapping as an incremental copy. The default is nonincremental. Once set, this value cannot be changed.
- cleanrate** *rate*
(Optional) Sets the cleaning rate for the mapping. The *rate* value can be 0 - 100. The default value is 50.
- iogrp** *iogroup_name* | *iogroup_id*
(Optional) Specifies the I/O group for the FlashCopy bitmap. Once set, this value cannot be changed. The default I/O group is either the source VDisk, if a single target map, or the I/O group of the other FlashCopy mapping to which either the source or target VDIsks belong.

Description

This command creates a new FlashCopy mapping. This mapping persists until it is manually deleted, or until it is automatically deleted when the background copy completes and the **autodelete** parameter set to **on**. The source and target VDIsks must be specified on the `mkfcmap` command. The `mkfcmap` command fails if the source and target VDIsks are not identical in size. Issue the `svcinfolsvdisk -bytes` command to find the exact size of the source VDisk for which you want to create a target disk of the same size. The target VDisk that you specify cannot be a target VDisk in an existing FlashCopy mapping. A mapping cannot be created if the resulting set of connected mappings exceeds 256 connected mappings.

The mapping can optionally be given a name and assigned to a consistency group, which is a group of mappings that can be started with a single command. These are groups of mappings that can be processed at the same time. This enables multiple VDIsks to be copied at the same time, which creates a consistent copy of

multiple disks. This consistent copy of multiple disks is required by some database products in which the database and log files reside on different disks.

If the specified source and target VDisks are the target and source VDisks, respectively, of an existing mapping, then the mapping being created and the existing mapping become partners. If one mapping is created as incremental, then its partner is automatically incremental. A mapping can have only one partner.

The **copyrate** parameter specifies the copy rate. If 0 is specified, background copy is disabled. The **cleanrate** parameter specifies the rate for cleaning the target VDisk. The cleaning process is only active if the mapping is in the copying state and the background copy has completed, the mapping is in the copying state and the background copy is disabled, or the mapping is in the stopping state. You can disable cleaning when the mapping is in the copying state by setting the **cleanrate** parameter to 0. If the **cleanrate** is set to 0, the cleaning process runs at the default rate of 50 when the mapping is in the stopping state to ensure that the stop operation completes.

Table 10 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 10. Relationship between the rate, data rate and grains per second values

User-specified <i>rate</i> attribute value	Data copied/sec	256 KB grains/sec	64 KB grains/sec
1 - 10	128 KB	0.5	2
11 - 20	256 KB	1	4
21 - 30	512 KB	2	8
31 - 40	1 MB	4	16
41 - 50	2 MB	8	32
51 - 60	4 MB	16	64
61 - 70	8 MB	32	128
71 - 80	16 MB	64	256
81 - 90	32 MB	128	512
91 - 100	64 MB	256	1024

An invocation example

```
svctask mkfcmap -source 0 -target 2 -name mapone
```

The resulting output

```
FlashCopy Mapping, id [1], successfully created
```

prestartfcconsistgrp

The `prestartfcconsistgrp` command prepares a consistency group (a group of FlashCopy mappings) so that the consistency group can be started. This command flushes the cache of any data that is destined for the source VDisk and forces the cache into the write-through mode until the consistency group is started.

Syntax

```
svctask -- prestartfcconsistgrp -- [-restore] <fc_consist_group_id | fc_consist_group_name>
```

Parameters

-restore

(Optional) Specifies the restore flag. This forces the consistency group to be prepared even if the target VDisk of one of the mappings in the consistency group is being used as a source VDisk of another active mapping. An active mapping is in the copying, suspended, or stopping state.

fc_consist_group_id | *fc_consist_group_name*

(Required) Specifies the name or ID of the consistency group that you want to prepare.

Description

This command prepares a consistency group (a group of FlashCopy mappings) to subsequently start. The preparation step ensures that any data that resides in the cache for the source VDisk is first flushed to disk. This step ensures that the FlashCopy target VDisk is identical to what has been acknowledged to the host operating system as having been written successfully to the source VDisk.

You can use the **restore** parameter to force the consistency group to be prepared even if the target VDisk of one or more mappings in the consistency group is being used as a source VDisk of another active mapping. In this case the mapping restores as shown in the `svcinfo lsfcmmap` view. If the **restore** parameter is specified when preparing a consistency group where none of the target VDIs are the source VDisk of another active mapping, then the parameter is ignored.

You must issue the `svctask prestartfcconsistgrp` command to prepare the FlashCopy consistency group before the copy process can be started. When you have assigned several mappings to a FlashCopy consistency group, you must issue a single prepare command for the whole group to prepare all of the mappings at once.

The consistency group must be in the `idle_or_copied` or `stopped` state before it can be prepared. When you enter the `prestartfcconsistgrp` command, the group enters the preparing state. After the preparation is complete, the consistency group status changes to `prepared`. At this point, you can start the group.

If FlashCopy mappings are assigned to a consistency group, the preparing and the subsequent starting of the mappings in the group must be performed on the consistency group rather than on an individual FlashCopy mapping that is assigned to the group. Only stand-alone mappings, which are mappings that are not assigned to a consistency group, can be prepared and started on their own. A FlashCopy mapping must be prepared before it can be started.

An invocation example

```
svctask prestartfcconsistgrp 1
```

The resulting output

No feedback

prestartfcmap

The `prestartfcmap` command prepares a FlashCopy mapping so that it can be started. This command flushes the cache of any data that is destined for the source VDisk and forces the cache into the write-through mode until the mapping is started.

Syntax

```
svctask -- prestartfcmap -- [-restore] [fc_map_id | fc_map_name]
```

Parameters

-restore

(Optional) Specifies the restore flag. This forces the mapping to be prepared even if the target VDisk is being used as a source VDisk in another active mapping. An active mapping is in the copying, suspended, or stopping state.

fc_map_id | *fc_map_name*

(Required) Specifies the name or ID of the mapping to prepare.

Description

This command prepares a single mapping for subsequent starting. The preparation step ensures that any data that resides in the cache for the source VDisk is first transferred to disk. This step ensures that the copy that is made is consistent with what the operating system expects on the disk.

The **restore** parameter can be used to force the mapping to be prepared even if the target VDisk is being used as a source VDisk of another active mapping. In this case, the mapping is restoring as shown in the `svctask lscmap` view. If the **restore** parameter is specified when preparing a mapping where the target VDisk is not the source VDisk of another active mapping, then the parameter is ignored.

Note: To prepare a FlashCopy mapping that is part of a consistency group, you must use the `prestartfcconsistgrp` command.

The mapping must be in the `idle_or_copied` or `stopped` state before it can be prepared. When the `prestartfcmap` command is processed, the mapping enters the `preparing` state. After the preparation is complete, it changes to the `prepared` state. At this point, the mapping is ready to start.

Attention: This command can take a considerable amount of time to complete.

An invocation example

```
svctask prestartfcmap 1
```

The resulting output

No feedback

rmfcconsistgrp

The **rmfcconsistgrp** command deletes a FlashCopy consistency group.

Syntax

```
svctask -- rmfcconsistgrp -- [-force] --
fc_consist_group_id | fc_consist_group_name
```

Parameters

-force

(Optional) Specifies that all of the mappings that are associated with a consistency group that you want to delete are removed from the group and changed to stand-alone mappings. This parameter is only required if the consistency group that you want to delete contains mappings.

fc_consist_group_id | *fc_consist_group_name*

(Required) Specifies the ID or name of the consistency group that you want to delete.

Description

This command deletes the specified FlashCopy consistency group. If there are mappings that are members of the consistency group, the command fails unless you specify the **-force** parameter. When you specify the **-force** parameter, all of the mappings that are associated with the consistency group are removed from the group and changed to stand-alone mappings.

To delete a single mapping in the consistency group, you must use the **svctask rmfcmap** command.

An invocation example

```
svctask rmfcconsistgrp fcconsistgrp1
```

The resulting output

No feedback

rmfcmap

The **rmfcmap** command deletes an existing mapping.

Syntax

```
svctask -- rmfcmap -- [-force] -- fc_map_id | fc_map_name
```

Parameters

-force

(Optional) Specifies that the target VDisk is brought online. This parameter is required if the FlashCopy mapping is in the stopped state.

fc_map_id | *fc_map_name*

(Required) Specifies the ID or name of the FlashCopy mapping to delete. Enter the ID or name last on the command line.

Description

The `rmfcmap` command deletes the specified mapping if the mapping is in the `idle_or_copied` or `stopped` state. If it is in the `stopped` state, the `-force` parameter is required. If the mapping is in any other state, you must stop the mapping before you can delete it.

Deleting a mapping only deletes the logical relationship between the two virtual disks; it does not affect the virtual disks themselves. However, if you force the deletion, the target virtual disk (which might contain inconsistent data) is brought back online.

An invocation example

```
svctask rmfcmap testmap
```

The resulting output

No feedback

startfcconsistgrp

The `startfcconsistgrp` command starts a FlashCopy consistency group of mappings. This command makes a point-in-time copy of the source VDisks at the moment that the command is started.

Syntax

```
svctask -- startfcconsistgrp -- [-prep] [-restore]
fc_consist_group_id | fc_consist_group_name
```

Parameters

-prep

(Optional) Specifies that the designated FlashCopy consistency group be prepared prior to starting the FlashCopy consistency group. A FlashCopy consistency group must be prepared before it can be started. When you use this parameter, the system automatically issues the `prestartfcconsistgrp` command for the group that you specify.

-restore

(Optional) Specifies the restore flag. When combined with the `prep` option, this forces the consistency group to be prepared even if the target VDisk of one of the mappings in the consistency group is being used as a source VDisk in another active mapping. An active mapping is in the `copying`, `suspended`, or `stopping` state.

fc_consist_group_id | *fc_consist_group_name*

(Required) Specifies the ID or name of the consistency group mapping to start.

Description

This command starts a consistency group, which results in a point-in-time copy of the source VDisks of all mappings in the consistency group. You can combine the **restore** parameter with the **prep** parameter to force the consistency group to be prepared prior to starting, even if the target VDisk of one or more mappings in the consistency group is being used as a source VDisk of another active mapping. In this case, the mapping is restoring as shown in the svcinfo lsfcmmap view. If the **restore** parameter is specified when starting a consistency group where none of the target VDisks are the source VDisk of another active mapping, then the parameter is ignored.

If a consistency group is started and the target VDisk of the mapping being started has up to four other incremental FlashCopy mappings using the target, the incremental recording is left on. If there are more than four other incremental FlashCopy mappings using the target VDisk, the incremental recording for all of these mappings is turned off until they are restarted.

Note: The startfcconsistgrp command can take some time to process particularly if you have specified the **prep** parameter. If you use the **prep** parameter, you give additional processing control to the system because the system must prepare the mapping before the mapping is started. If the prepare process takes too long, the system completes the prepare but does not start the consistency group. In this case, error message CMMVC6209E displays. To control the processing times of the prestartfcconsistgrp and startfcconsistgrp commands independently of each other, do not use the **prep** parameter. Instead, first issue the prestartfcconsistgrp command, and then to start the copy, issue the startfcconsistgrp command.

An invocation example

```
svctask startfcconsistgrp -prep 2
```

The resulting output

No feedback

startfcmap

The startfcmap command starts a FlashCopy mapping. This command makes a point-in-time copy of the source VDisk at the moment that the command is started.

Syntax

```
svctask -- startfcmap -- [-prep] [-restore]
fc_map_id
fc_map_name
```

Parameters

-prep

(Optional) Specifies that the designated mapping be prepared prior to starting the mapping. A mapping must be prepared before it can be started. When you use this parameter, the system automatically issues the prestartfcmap command for the group that you specify.

-restore

(Optional) Specifies the restore flag. When combined with the **prep** option, this forces the mapping to be prepared even if the target VDisk is being used as a source VDisk in another active mapping. An active mapping is in the copying, suspended, or stopping state.

fc_map_id | *fc_map_name*

(Required) Specifies the ID or name of the mapping to start.

Description

This command starts a single mapping, which results in a point-in-time copy of the source VDisk. You can combine the **restore** parameter with the **prep** parameter to force the mapping to be prepared prior to starting, even if the target VDisk is being used as a source VDisk of another active mapping. In this case, the mapping is restoring as shown in the svcinfo lsfcmmap view. If the **restore** parameter is specified when starting a mapping where the target VDisk is not the source VDisk of another active mapping, then the parameter is ignored and the mapping is not restoring as shown in the svcinfo lsfcmmap view.

If a mapping is started and the target VDisk of the mapping being started has up to four other incremental FlashCopy mappings using the target, the incremental recording is left on. If there are more than four other incremental FlashCopy mappings using the target VDisk, the incremental recording for all of these mappings is turned off until they are restarted.

Note: The startfcmap command can take some time to start, particularly if you use the **prep** parameter. If you use the **prep** parameter, you give additional starting control to the system. The system must prepare the mapping before the mapping is started. To keep control when the mapping starts, you must issue the prestartfcmap command before you issue the startfcmap command.

An invocation example

svctask startfcmap -prep 2


The resulting output

No feedback

stopfcconsistgrp

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

Syntax



Parameters

-force

(Optional) Specifies that all processing that is associated with the mappings of the designated consistency group be stopped immediately.

Note: When you use this parameter, all FlashCopy mappings that depend on the mappings in this group (as listed by the `lsfcmapdependentmaps` command) are also stopped.

-split

(Optional) Breaks the dependency on the source VDisks of any mappings that are also dependent on the target VDisk. This parameter can only be specified when stopping a consistency group where all maps in the group have progress of 100 as shown by the `svcinfo lsfcmap` command.

fc_consist_group_id | fc_consist_group_name

(Required) Specifies the name or ID of the consistency group that you want to stop.

Description

This command stops a group of mappings in a consistency group. If the copy process is stopped, the target disks become unusable unless they already contain complete images of the source. Disks that contain complete images of the source have a progress of 100, as indicated in the `lsfcmap` command output. The target VDisk is reported as offline if it does not contain a complete image. Before you can access this VDisk, the group of mappings must be prepared and restarted.

If the consistency group is in the `idle_or_copied` state, then the `stopfcconsistgrp` command has no effect and the consistency group stays in the `idle_or_copied` state.

Note: Prior to SVC 4.2.0, the `stopfcconsistgrp` command always caused the consistency group to go to the stopped state, taking the target VDisks offline.

The **split** option can be used when all of the maps in the group have progress of 100. It removes the dependency of any other maps on the source VDisks. It might be used prior to starting another FlashCopy consistency group whose target disks are the source disks of the mappings being stopped. Once the consistency group has been stopped with the **split** option, the other consistency group could then be started without the **restore** option.

An invocation example

```
svctask stopfcconsistgrp testmapone
```

The resulting output

```
No feedback
```

stopfcmap

The `stopfcmap` command stops all processing that is associated with a FlashCopy mapping that is in one of the following processing states: prepared, copying, stopping, or suspended.

Syntax

```
svctask — — stopfcmap — [ -force ] [ -split ] [ fc_map_id | fc_map_name ]
```

Parameters

-force

(Optional) Specifies that all processing that is associated with the designated mapping be stopped immediately.

Note: When you use this parameter, all FlashCopy mappings that depend on this mapping (as listed by the `lsfcmapdependentmaps` command) are also stopped.

-split

(Optional) Breaks the dependency on the source VDisk of any mappings that are also dependent on the target disk. This parameter can only be specified when stopping a map that has progress of 100 as shown by the `svcinfo lsfcmap` command.

fc_map_id | *fc_map_name*

(Required) Specifies the name or ID of the mapping to stop.

Description

This command stops a single mapping. If the copy process is stopped, the target disk becomes unusable unless it already contained a complete image of the source (that is, unless the map had a progress of 100 as shown by the `lsfcmap` command). Before you can use the target disk, the mapping must once again be prepared and then reprocessed (unless the target disk already contained a complete image).

Only stand-alone mappings can be stopped using the `stopfcmap` command. Mappings that belong to a consistency group must be stopped using the `stopfcconsistgrp` command.

If the mapping is in the `idle_or_copied` state, the `stopfcmap` command has no effect and the mapping stays in the `idle_or_copied` state.

Note: Before SAN Volume Controller 4.2.0, the `stopfcmap` command always changed the mapping state to `stopped` and took the target VDisk offline. This change can break scripts that depend on the previous behavior.

The **split** option can be used when the mapping has progress of 100. It removes the dependency of any other mappings on the source VDisk. It might be used prior to starting another FlashCopy mapping whose target disk is the source disk of the mapping being stopped. Once the mapping has been stopped with the **split** option, the other mapping could then be started without the **restore** option.

An invocation example

```
svctask stopfcmap testmapone
```

The resulting output

No feedback

Chapter 16. Metro Mirror and Global Mirror commands

The following Copy Service commands enable you to work with the Metro Mirror and Global Mirror services that the SAN Volume Controller provides.

chpartnership

The **chpartnership** command modifies the bandwidth of the partnership between the local cluster and the remote cluster that is specified in the command. This affects the bandwidth that is available for background copy in a cluster partnership by either Metro Mirror or Global Mirror operations. This command can also be used to disable and re-enable the partnership, to permit the local cluster to be disconnected and then reconnected to the remote cluster.

Syntax

```
svctask -- chpartnership [ -bandwidth bandwidth_in_mbps ]
[ -start | -stop ] remote_cluster_id | remote_cluster_name
```

Parameters

-bandwidth *bandwidth_in_mbps*

(Optional) Specifies the new bandwidth in megabytes per second (MBps). This bandwidth is used to cap the background remote copy progress. Set the bandwidth to the maximum rate that the remote copies should resynchronize at. Write operations from the host add to the use of the cluster link. If this parameter is set to a value that is greater than the intercluster links can sustain, the actual copy rate defaults to what is available on the link.

-start | **-stop**

(Optional) Starts or stops a Metro Mirror or Global Mirror partnership. To start or stop a partnership, run the **svctask chpartnership** command from either cluster.

remote_cluster_id | *remote_cluster_name*

(Required) Specifies the cluster ID or name of the remote cluster. The intracluster bandwidth cannot be modified, so if you enter the local cluster name or ID, an error occurs.

Description

This command modifies the bandwidth of the partnership between the local cluster and the remote cluster that is specified in the command. This affects the bandwidth that is available for a background copy in Metro Mirror or Global Mirror relationships, in the direction from the local to the remote cluster. To modify the background copy bandwidth in the other direction (remote cluster → local cluster), it is necessary to issue the corresponding **chpartnership** command to the remote cluster.

When you stop the cluster partnership, you temporarily disable the partnership and disconnect the local cluster from the remote cluster. The configuration is retained. The cluster partnership must be in either the `partially_configured_stopped` or `fully_configured_stopped` states to be started.

An invocation example

```
svctask chpartnership -bandwidth 20 cluster1
svctask chpartnership -stop cluster1
```

The resulting output

No feedback

chrconsistgrp

The **chrconsistgrp** command modifies the name of an existing Metro Mirror or Global Mirror consistency group.

Syntax

```
svctask -- chrconsistgrp -- -name -- new_name_arg --
└── rc_consist_group_name ────────────────────────────────────▶
    └── rc_consist_group_id ───────────────────────────────────▶
```

Parameters

-name *new_name_arg*

(Required) Specifies the new name to assign to the consistency group.

rc_consist_group_name | *rc_consist_group_id*

(Required) Specifies the ID or existing name of the consistency group that you want to modify.

Description

This command changes the name of the specified consistency group.

An invocation example

Change the name of the consistency group called `rc_testgrp` to `rctestone`.

```
svctask chrconsistgrp -name rctestone rc_testgrp
```


The resulting output

No feedback

chrrelationship

The **chrrelationship** command enables you to modify certain attributes of an existing relationship, such as to add a relationship to a consistency group, to remove a relationship from a consistency group, and to change the name of the relationship. You can only change one attribute at a time per command submittal.

Syntax



Parameters

-name *new_name_arg*

(Optional) Specifies a new label to assign to the relationship.

This parameter is required if you do not specify the **-consistgrp** or **-force** parameter.

-consistgrp *consist_group_id* | *consist_group_name*

(Optional) Specifies a new consistency group to assign the relationship to. Only relationships of the same copy type (Metro Mirror or Global Mirror) can be assigned to the same consistency group. You cannot use this parameter with the **-name**, or **-force** parameters.

This parameter is required if you do not specify the **-name** or **-force** parameter.

-force

(Optional) Specifies that you want the system to remove the relationship from a consistency group making the relationship a stand-alone relationship. You cannot use this parameter with the **-name** or **-consistgrp** parameters.

This parameter is required if you do not specify the **-name** or **-consistgrp** parameter.

rc_rel_name | *rc_rel_id*

(Required) Specifies the ID or name of the relationship.

Description

This command modifies the specified attributes of the supplied relationship, one attribute at a time; you cannot use more than one of the three optional parameters in the same command. In addition to changing the name of a consistency group, this command can be used for the following purposes.

- You can add a stand-alone relationship to a consistency group by specifying the **-consistgrp** parameter and the name or ID of the consistency group. The relationship and consistency group must be connected when the command is issued and must share the following components:
 - Master cluster
 - Auxiliary cluster
 - State (unless the group is empty)
 - Primary (unless the group is empty)
 - Type (unless the group is empty)

When the first relationship is added to an empty group, the group takes on the same state, primary (copy direction), and type (Metro or Global Mirror) as the relationship. Subsequent relationships must have the same state, copy direction, and type as the group in order to be added to it. A relationship can only belong to one consistency group.

- You can remove a relationship from a consistency group by specifying the **-force** parameter and the name or ID of the relationship. Although you do not have to

specify or confirm the name of the consistency group, verify which group the relationship belongs to before you issue this command.

This form of the modify relationship command succeeds in the connected or disconnected states. If the clusters are disconnected the relationship is only removed from the consistency group on the local cluster, at the time the command is issued. When the clusters are reconnected the relationship is automatically removed from the consistency group on the other cluster. Alternatively, you can issue an explicit modify (**chrrelationship**) command to remove the relationship from the group on the other cluster while it is still disconnected.

Note: If you remove all relationships from the group, the relationship type is reset to **empty_group**. When you add a relationship to the empty group, the group again takes on the same type as the relationship.

- To move a relationship between two consistency groups, you must issue the **chrrelationship** command twice. Use the **-force** parameter to remove the relationship from its current group, and then use the **-consistgrp** parameter with the name of the new consistency group.

An invocation example

Change the name of the relationship rccopy1 to testrel
 svctask chrrelationship -name testrel rccopy1

Add relationship rccopy2 to group called newgroup.
 svctask chrrelationship -consistgrp newgroup rccopy2

Remove relationship rccopy3 from whichever consistency group it is a member of.
 svctask chrrelationship -force rccopy3

The resulting output

No feedback

There is no feedback in any of these cases.

mkpartnership

The **mkpartnership** command establishes a one-way Metro Mirror or Global Mirror relationship between the local cluster and a remote cluster.

Syntax

To establish a fully functional Metro Mirror or Global Mirror partnership, you must issue this command to both clusters. This step is a prerequisite to creating Metro Mirror or Global Mirror relationships between VDisks on the clusters.

```

▶▶▶ svctask — — mkpartnership — — -bandwidth — bandwidth_in_mbps — —————▶
▶ remote_cluster_id —————▶▶▶
  └─ remote_cluster_name ─┘
```

|
|
|
|
|

Parameters

-bandwidth *bandwidth_in_mbps*

(Required) Specifies the bandwidth, in megabytes per second (MBps), that is used by the background copy process between the clusters. It adjusts the bandwidth that is used by Metro Mirror or Global Mirror for the initial background copy process. Set the bandwidth to a value that is less than or equal to the bandwidth that can be sustained by the intercluster link. If the **-bandwidth** parameter is set to a higher value than the link can sustain, the background copy process uses the actual available bandwidth.

remote_cluster_id | *remote_cluster_name*

(Required) Specifies the cluster ID or name of the remote cluster. Issue the `svcinfolscandidate` command to list the remote clusters that are available. If two or more remote clusters have the same name and the name is included in this command, the command fails and it requests the ID of the cluster instead of the name.

Description

This command creates a one-way partnership between the local cluster and the remote cluster that you specify in the command. To create a two-way partnership, the equivalent `svctask mkpartnership` command must be issued from the other cluster. The maximum supported number of clusters in a partnership set is four.

Intercluster Mirror relationships can be created between primary VDisks in the local cluster and auxiliary VDisks in the remote cluster. Intracluster relationships can be created between VDisks that reside in a local cluster. The VDisks must belong to the same I/O group within the cluster.

Note: Background copy bandwidth can impact foreground I/O latency. To set the background copy bandwidth optimally, you must consider all three resources: the primary storage, intercluster link bandwidth, and secondary storage. Provision the most restrictive of these three resources between the background copy bandwidth and the peak foreground I/O workload.

An invocation example

```
svctask mkpartnership -bandwidth 20 cluster1
```

The resulting output


No feedback

mkrcconsistgrp

The **mkrcconsistgrp** command creates a new, empty Metro Mirror or Global Mirror consistency group. If the **-cluster** parameter is not specified, the consistency group is created on the local cluster only.

Syntax

```
svctask -- mkrcconsistgrp -- [ -name new_name ]
```



Parameters

-name *new_name*

(Optional) Specifies a name for the new consistency group.

-cluster *cluster_id* | *cluster_name*

(Optional) Specifies the name or ID of the remote cluster. If **-cluster** is not specified, a consistency group is created only on the local cluster.

Description

This command creates a new consistency group. The ID of the new group is displayed after the command processes. The name must be unique across all consistency groups that are known to the clusters within this consistency group. If the consistency group involves two clusters, the clusters must be in communication throughout the create process.

The new consistency group does not contain any relationships and will be in the empty state. You can add Metro Mirror or Global Mirror relationships to the group using the **svctask chrrelationship** command.

An invocation example

```
svctask mkrconsistgrp -name rc_testgrp
```

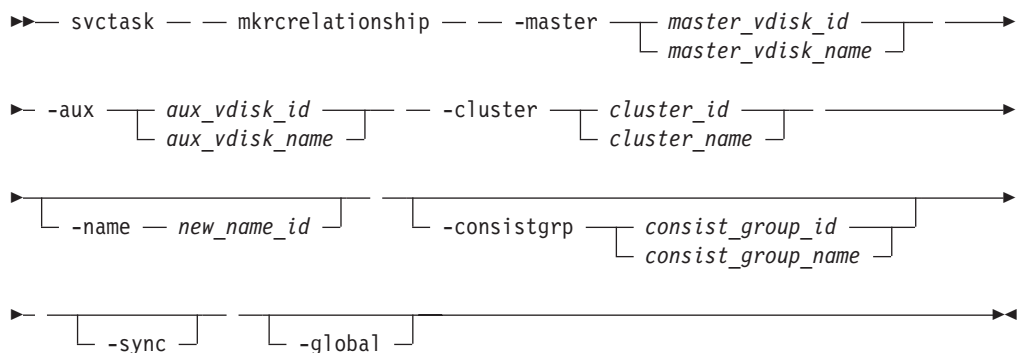
The resulting output

```
RC Consistency Group, id [255], successfully created
```

mkrrelationship

The **mkrrelationship** command creates a new Metro Mirror or Global Mirror relationship with virtual disks (VDisks) in the same cluster (intracluster relationship) or in two different clusters (intercluster relationship).

Syntax



Parameters

-master *master_vdisk_id* | *master_vdisk_name*

(Required) Specifies the ID or name of the master virtual disk.

-aux *aux_vdisk_id* | *aux_vdisk_name*

(Required) Specifies the ID or name of the auxiliary virtual disk.

-cluster *cluster_id* | *cluster_name*

(Required) Specifies the ID or name of the remote cluster.

If you are creating an intracluster relationship, enter the ID of the local cluster. The VDIs in the relationship must belong to the same I/O group within the cluster.

If you are creating an intercluster relationship, enter the ID of the remote cluster. To create a relationship in two different clusters, the clusters must be connected at the time that the **svctask mkrrelationship** command is received.

-name *new_name_id*

(Optional) Specifies a label to assign to the relationship.

-consistgrp *consist_group_id* | *consist_group_name*

(Optional) Specifies a consistency group that this relationship joins. If you do not supply the **-consistgrp** parameter, the relationship is created as a stand-alone relationship that can be started, stopped, and switched on its own.

Note: Metro and Global Mirror relationships cannot belong to the same consistency group. When the first relationship is added to the consistency group, the group takes on the same type as the relationship. Subsequently, only relationships of that type can be added to the consistency group.

-sync

(Optional) Specifies that you want the system to create a synchronized relationship. The **-sync** parameter guarantees that the master and auxiliary disks contain identical data at the point that the relationship is created. You must ensure that the auxiliary disk is created to match the master disk and that no input transactions take place to either disk before you issue the create command. The initial background synchronization is skipped.

-global

(Optional) Specifies that you want the system to create a new Global Mirror relationship. If you do not specify the **-global** parameter, a Metro Mirror relationship is created instead.

Description

This command creates a new Metro Mirror or Global Mirror relationship. A Metro Mirror relationship defines the relationship between two virtual disks (VDIs): a master VDI and an auxiliary VDI. This relationship persists until it is deleted. The auxiliary virtual disk must be identical in size to the master virtual disk or the command fails, and if both VDIs are in the same cluster, they must both be in the same I/O group. The master and auxiliary cannot be in an existing relationship. Neither disk can be the target of a FlashCopy mapping. The command also returns the new relationship ID.

Metro Mirror relationships use one of the following copy types:

- A Metro Mirror copy ensures that updates are committed to both the primary and secondary VDIs before sending confirmation of I/O completion to the host application. This ensures that the secondary VDI is synchronized with the primary VDI in the event that a failover operation is performed.
- A Global Mirror copy allows the host application to receive confirmation of I/O completion before the updates are committed to the secondary VDI. If a

failover operation is performed, the host application must recover and apply any updates that were not committed to the secondary VDisk.

You can optionally give the relationship a name. The name must be a unique relationship name across both clusters.

The relationship can optionally be assigned to a consistency group. A consistency group ensures that a number of relationships are managed so that, in the event of a disconnection of the relationships, the data in all relationships within the group is in a consistent state. This can be important in, for example, a database application where data files and log files are stored on separate VDIs and consequently are managed by separate relationships. In the event of a disaster, the primary and secondary sites might become disconnected. As the disconnection occurs and the relationships stop copying data from the primary to the secondary site, there is no assurance that updates to the two separate secondary VDIs will stop in a consistent manner if the relationships that are associated with the VDIs are not in a consistency group.

For proper database operation, it is important that updates to the log files and the database data are made in a consistent and orderly fashion. It is crucial in this example that the log file VDisk and the data VDisk at the secondary site are in a consistent state. This can be achieved by putting the relationships that are associated with these VDIs into a consistency group. Both Metro Mirror and Global Mirror processing ensure that updates to both VDIs at the secondary site are stopped, leaving a consistent image based on the updates that occurred at the primary site.

If you specify a consistency group, both the group and the relationship must have been created using the same master cluster and the same auxiliary cluster. The relationship must not be a part of another consistency group. If the consistency group is *empty*, it acquires the type of the first relationship that is added to it. Therefore, each subsequent relationship that you add to the consistency group must have the same type.

If the consistency group is *not empty*, the consistency group and the relationship must be in the same state. If the consistency group is *empty*, it acquires the state of the first relationship that is added to it. If the state has an assigned copy direction, the direction of the consistency group and the relationship must match that direction.

If you do not specify a consistency group, a stand-alone relationship is created.

If you specify the **-sync** parameter, the master and auxiliary virtual disks contain identical data at the point when the relationship is created. You must ensure that the auxiliary is created to match the master and that no data movement occurs to either virtual disk before you issue the **svctask mkrrelationship** command.

If you specify the **-global** parameter, a Global Mirror relationship is created. Otherwise, a Metro Mirror relationship is created instead.

An invocation example

```
svctask mkrrelationship -master vdisk1 -aux vdisk2 -name rccopy1  
-cluster 0000020063432AFD
```

The resulting output

```
RC Relationship, id [28], successfully created
```

rmpartnership

| The `rmpartnership` command removes a Metro Mirror or Global Mirror partnership
| on one cluster. Because the partnership exists on both clusters, it is necessary to
| run this command on both clusters to remove both sides of the partnership. If the
| command is run on only one cluster, the partnership enters a partially configured
| state on the other cluster.

Syntax

```
▶▶— svctask — — rmpartnership — — [ remote_cluster_id ] —————▶▶  
| [ remote_cluster_name ] |
```

Parameters

remote_cluster_id | *remote_cluster_name*

(Required) Specifies the cluster ID or the name of the remote cluster.

Description

This command deletes one half of a partnership on a cluster. To remove the entire partnership, you must run the command twice, once on each cluster.

Attention: Before running the `svctask rmpartnership` command, you must remove all relationships and groups that are defined between the two clusters. To display cluster relationships and groups, run the `svcinfo lsrelationship` and `svcinfo lsrconsistgrp` commands. To remove the relationships and groups that are defined between the two clusters, run the `svctask rmrrelationship` and `svctask rmrconsistgrp` commands.

An invocation example

```
svctask rmpartnership cluster1
```

The resulting output

No feedback

rmrconsistgrp

The `rmrconsistgrp` command deletes an existing Metro Mirror or Global Mirror consistency group.

Syntax

```
▶▶— svctask — — rmrconsistgrp — — [ -force ] —————▶▶  
  
▶ [ rc_consist_group_id ] —————▶▶  
| [ rc_consist_group_name ] |
```

Parameters

`-force`

(Optional) Specifies that you want the system to remove all relationships

belonging to a group before the consistency group is deleted. The relationships themselves are not deleted; they become stand-alone relationships.

Note: The **-force** parameter must be used to delete a consistency group when the consistency group has any Metro Mirror or Global Mirror relationships that is associated with it. If you do not use the **-force** parameter, the command fails.

rc_consist_group_id | *rc_consist_group_name*

(Required) Specifies the ID or the name of the consistency group to delete.

Description

This command deletes the specified consistency group. You can issue this command for any existing consistency group. If the consistency group is disconnected at the time that the command is issued, the consistency group is only deleted on the cluster that is connected. When the clusters reconnect, the consistency group is automatically deleted on the other cluster. Alternatively, if the clusters are disconnected, and you still want to remove the consistency group on both clusters, you can issue the **svctask rmrconsistgrp** command separately on both of the clusters.

If the consistency group is not empty, the **-force** parameter is required to delete the group. This removes the relationships from the consistency group before the group is deleted. These relationships become stand-alone relationships. The state of these relationships is not changed by the action of removing them from the consistency group.

An invocation example

```
svctask rmrconsistgrp rctestone
```

The resulting output

No feedback

rmrrelationship

The **rmrrelationship** command deletes an existing Metro Mirror or Global Mirror relationship.

Syntax

```
svctask -- rmrrelationship --  $\left. \begin{array}{l} rc\_rel\_id \\ rc\_rel\_name \end{array} \right\}$ 
```

Parameters

rc_rel_id | *rc_rel_name*

(Required) Specifies the ID or the name of the relationship. A relationship cannot be deleted if it is part of a consistency group.

Description

This command deletes the relationship that is specified.

Deleting a relationship only deletes the logical relationship between the two virtual disks; it does not affect the virtual disks themselves.

If the relationship is disconnected at the time that the command is issued, the relationship is only deleted on the cluster where the command is being run. When the clusters reconnect, the relationship is automatically deleted on the other cluster. Alternatively, if the clusters are disconnected and if you still want to remove the relationship on both clusters, you can issue the `svctask rmrrelationship` command independently on both of the clusters.

If you delete an inconsistent relationship, the secondary virtual disk becomes accessible even though it is still inconsistent. This is the one case in which Metro or Global Mirror does not inhibit access to inconsistent data.

An invocation example

```
svctask rmrrelationship rccopy1
```

The resulting output

No feedback

starttrconsistgrp

The `starttrconsistgrp` command starts the Metro Mirror or Global Mirror consistency group copy process, sets the direction of copy if it is undefined, and optionally marks the secondary VDisks of the consistency group as clean.

Syntax

```
svctask -- starttrconsistgrp -- [-primary master | aux] [-force] [-clean] rc_consist_group_id rc_consist_group_name
```

Parameters

-primary *master* | *aux*

(Optional) Specifies the copy direction by defining whether the master or auxiliary disk becomes the primary (source). This parameter is required when the primary is undefined if, for example, the consistency group is in the **Idling** state.

-force

(Optional) Specifies that you want the system to process the copy operation even if it might lead to a temporary loss of consistency while synchronization occurs. This parameter is required if the consistency group is in the **ConsistentStopped** state, but is not synchronized or is in the **Idling** state, but is not synchronized.

-clean

(Optional) Specifies that the VDisk that is to become a secondary is clean for each of the relationships belonging to the group; any changes made on the secondary VDisk are ignored, and only changes made on the clean primary VDisk are considered during synchronization of the primary and secondary disks. The consistency group must be in an **Idling** (connected) state for this parameter to work.

| **Attention:** This flag should only be used when the primary and secondary
| VDisks contain identical data. Otherwise, relationships that are not consistent
| are reported as consistent. Once this has been done there is no method to
| determine whether these VDisks ever reach a true consistent state until a full
| background copy can be carried out again.

rc_consist_group_id | rc_consist_group_name

(Required) Specifies the ID or name of the consistency group to start.

Description

This command starts a Metro Mirror or Global Mirror stand-alone consistency group.

This command can only be issued to a consistency group that is connected. For a consistency group that is idling, this command assigns a copy direction (primary and secondary roles) and begins the copy process. Otherwise, this command restarts a previous copy process that was stopped either by a stop command or by an I/O error.

If the resumption of the copy process leads to a period of time when the relationship is not consistent, then you must specify the **-force** parameter when you restart the relationship. This situation can arise if the relationship had been stopped and then further input transactions had been performed on the original primary disk of the relationship. When you use the **-force** parameter in this situation, the data on the secondary disk is not usable (because it is inconsistent) in a disaster recovery circumstance.

In the idling state, you must provide the **-primary** parameter. In other connected states, you can provide the **-primary** parameter, but it must match the existing setting.

The **-force** parameter is required if consistency would be lost by starting a copy operation. This can occur if write operations on either primary or secondary VDisks have taken place since the **ConsistentStopped** or **idling** state occurred. If the command is issued without the **-force** parameter in such circumstances, the command fails. In general, the **-force** parameter is required if the group is in one of the following states:

- Consistent_Stopped but not synchronized (sync=out_of_sync)
- Idling but not synchronized

The **-force** parameter is not required if the group is in one of the following states:

- Inconsistent_Stopped
- Inconsistent_Copying
- Consistent_Synchronized

However, the command does not fail if you specify the **-force** parameter.

The **-clean** parameter is used when a Metro Mirror or Global Mirror group is started and the secondary VDisks in this group are assumed to be clean. Clean in this sense, means that any changes that have been made at the secondary are ignored and only changes made at the primary are considered when synchronizing the primary and secondary VDisks. The **-clean** parameter can be used in the following scenario:

1. A consistency group is created with the **-sync** parameter. At this point, it does not matter if the primary and secondary contain the same data, even though the use of the **-sync** parameter implies that this is true.
2. A `stoprconsistgrp` command is issued with the **-access** parameter. This permits access to the secondary disk. Change recording begins at the primary.
3. An image of the primary disk is copied and loaded on to the secondary disk. It is permissible to allow updates to the primary disk during the image copy as this image can be only a fuzzy image of the primary disk.
4. A `starttrconsistgrp` command that specifies the **-primary master**, **-force**, and **-clean** parameters is issued. The auxiliary disk is marked as clean and changes on the master disk that have occurred since the relationship was stopped are copied to the auxiliary disk.
5. Once the background copy has completed, relationships in the group become consistent and synchronized.

An invocation example

```
svctask starttrconsistgrp rccopy1
```


The resulting output

No feedback

startrelationship

The `startrelationship` command starts the Metro Mirror or Global Mirror relationship copy process, sets the direction of copy if undefined, and optionally, marks the secondary VDisk of the relationship as clean. The relationship must be a stand-alone relationship.

Syntax



Parameters

-primary *master* | *aux*

(Optional) Specifies the copy direction by defining whether the master or auxiliary disk becomes the primary (source). This parameter is required when the primary is undefined if, for example, the relationship is in the idling state.

-force

(Optional) Specifies that you want the system to process the copy operation even if it might lead to a temporary loss of consistency while synchronization occurs. This parameter is required if the relationship is in the **Consistentstopped** state, but is not synchronized or in the Idling state, but is not synchronized.

-clean

(Optional) Specifies that the VDisk that is to become a secondary is clean; any changes made on the secondary VDisk are ignored, and only changes made on

the clean primary VDisk are considered when synchronizing the primary and secondary disks. The relationship must be in an Idling (connected) state for this parameter to work.

Attention: This flag should only be used when the primary and secondary VDisks contain identical data. Otherwise, relationships that are not consistent are reported as consistent. Once this has been done there is no method to determine whether these VDisks ever reach a true consistent state until a full background copy can be carried out again.

rc_rel_id | rc_rel_name

(Required) Specifies the ID or name of the relationship that you want to start in a stand-alone relationship.

Description

The `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.

This command can only be issued to a relationship that is connected. For a relationship that is idling, this command assigns a copy direction (primary and secondary roles) and begins the copy process. Otherwise, this command restarts a previous copy process that was stopped either by a stop command or by some I/O error.

If the resumption of the copy process leads to a period of time when the relationship is not consistent, you must specify the **-force** parameter when you restart the relationship. This situation can arise if the relationship had been stopped, and then further input transactions had occurred on the original primary of the relationship. Even though you use of the **-force** parameter in this situation, the data on the secondary is not be useful for disaster recovery purposes because the relationship is in an inconsistent state.

In the idling state, you must provide the **-primary** parameter. In other connected states, you can provide the **-primary** parameter, but it must match the existing setting.

The **-force** parameter is required if consistency would be lost by starting a copy operation. This can occur if input transactions have occurred on either the primary or secondary VDisks since the **ConsistentStopped** or **Idling** state occurred. If the **startrelationship** command is issued without the **-force** parameter in such circumstances, the command fails. In general, the **-force** parameter is required if the relationship is in one of the following states:

- ConsistentStopped but not synchronized
- Idling but not synchronized

The **-force** parameter is not required if the relationship is in one of the following states:

- InconsistentStopped
- InconsistentCopying
- ConsistentSynchronized

However, the command will not fail if you do specify the **-force** parameter.

The **-clean** parameter is used when a Metro Mirror or Global Mirror relationship is started and the secondary VDisk in the relationship is assumed to be clean; any

changes made on the secondary VDisk are ignored, and only changes made on the clean primary VDisk are considered when synchronizing the primary and secondary disks. The **-clean** parameter can be used in the following circumstance:

1. A relationship is created with the **-sync** parameter specified. (At this point it does not matter if the primary and secondary disks contain the same data, even though the use of the **-sync** parameter implies that this is true).
2. A svctask stopprrelationship command is issued with the **-access** parameter specified. This permits access to the secondary disk. Change recording begins at the primary disk.
3. An image of the primary disk is copied and loaded on to the secondary disk. It is permissible to allow updates to the primary disk during the image copy as this image need only be a *fuzzy* image of the primary disk.
4. A svctask startprrelationship command that specifies the **-primary master**, **-force**, and **-clean** parameters is issued. The auxiliary disk is marked as clean and changes on the master disk that have occurred since the relationship was stopped are copied to the auxiliary disk.
5. Once the background copy has completed, the relationship becomes consistent and synchronized.

An invocation example

```
svctask startprrelationship rccopy1
```

The resulting output

```
No feedback
```

stopprconsistgrp

The stopprconsistgrp command stops the copy process for a Metro Mirror or Global Mirror consistency group. This command can also be used to enable write access to the secondary VDIs in the group if the group is in a consistent state.

Syntax

```
svctask -- stopprconsistgrp [-access] rc_consist_group_id | rc_consist_group_name
```

Parameters

-access

(Optional) Allows write access to consistent secondary VDIs 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 VDIs to secondary VDIs.

If the consistency group is in an inconsistent state, all copy operations stop and do not resume until you issue the `svctask startcrconsistgrp` command. For a consistency group in the `consistent_synchronized` state, this command causes a consistency freeze.

When a consistency group is in a consistent state (for example, in the `consistent_stopped`, `consistent_synchronized`, or `consistent_disconnected` state) you can issue the `access` parameter with the `stopcrconsistgrp` command to enable write access to the secondary virtual disks within that group. Table 11 shows consistency group initial and final states:

Table 11. stopcrconsistgrp consistency group states

Initial state	Final state	Notes®
<code>inconsistent_stopped</code>	<code>inconsistent_stopped</code>	If <code>access</code> is specified, the command is rejected.
<code>inconsistent_copying</code>	<code>inconsistent_stopped</code>	If <code>access</code> is specified, the command is rejected with no effect and the relationship remains in the <code>inconsistent_copying</code> state.
<code>consistent_stopped</code>	<code>consistent_stopped</code>	If <code>access</code> is specified, the final state is <code>idling</code> .
<code>consistent_synchronized</code>	<code>consistent_stopped</code>	If <code>access</code> is specified, the final state is <code>idling</code> . If <code>access</code> is not specified, the final state is <code>consistent_stopped</code> .
<code>idling</code>	<code>idling</code>	Remains in <code>idling</code> state whether <code>access</code> is specified or not.
<code>idling_disconnected</code>	<code>unchanged</code>	If specified without <code>access</code> , the relationship/group remains in <code>idling_disconnected</code> state. If the clusters reconnect, the relationship/group is in either <code>inconsistent_stopped</code> or <code>consistent_stopped</code> state.
<code>inconsistent_disconnected</code>	<code>inconsistent_stopped</code>	The command is rejected, with or without the <code>access</code> flag.
<code>consistent_disconnected</code>	<code>consistent_stopped</code>	The command is rejected if specified without <code>access</code> . If specified with <code>access</code> , the relationship/group moves to <code>idling_disconnected</code> .

An invocation example

```
svctask stopcrconsistgrp rccopy1
```

The resulting output

No feedback

stopcrrelationship

The stopcrrelationship command stops the copy process for a Metro Mirror or Global Mirror stand-alone relationship. You can also use this command to enable write access to a consistent secondary VDisk.

Syntax

```
svctask — stopcrrelationship — [-access] [rc_rel_id | rc_rel_name]
```

Parameters

-access

(Optional) Specifies that the system allow write access to a consistent secondary VDisk.

rc_rel_id | *rc_rel_name*

(Required) Specifies the ID or the name of the relationship to stop all processing for.

Description

The stopcrrelationship 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 VDIs.

If the relationship is in an inconsistent state, any copy operation stops and does not resume until you issue a svctask startcrrelationship command. For a relationship in the consistent_synchronized state, this command causes a consistency freeze.

When a relationship is in a consistent state – in the consistent_stopped, consistent_synchronized, or consistent_disconnected state – you can use the **access** parameter to enable write access to the secondary virtual disk. Table 12 provides consistency group initial and final states.

Table 12. stopcrrelationship consistency group states

Initial state	Final state	Notes
inconsistent_stopped	inconsistent_stopped	If access is specified, the command is rejected.
inconsistent_copying	inconsistent_stopped	If access is specified, the command is rejected with no effect and the relationship remains in the inconsistent_copying state.
consistent_stopped	consistent_stopped	If access is specified, the final state is idling.
consistent_synchronized	consistent_stopped	If access is specified, the final state is idling. If access is not specified, the final state is consistent_stopped.

Table 12. 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/group remains in idling_disconnected state. If the clusters reconnect, the relationship/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/group moves to idling_disconnected.

An invocation example

```
svctask stopprrelationship rccopy1
```

The resulting output

No feedback

switchrconsistgrp

The **switchrconsistgrp** command reverses the roles of the primary and secondary virtual disks (VDisks) in a Metro Mirror or Global Mirror consistency group when that consistency group is in a consistent state. All the relationships in the consistency group are affected by this change.

Syntax

```

▶▶ svctask — — switchrconsistgrp — — -primary — master —————▶
                                     └─ aux ─┘

▶ ┌ rc_consist_group_id ─────────────────────────────────────────────────▶
  └ rc_consist_group_name ─────────────────────────────────────────────────▶

```

Parameters

-primary master | aux

(Required) Specifies whether the master or auxiliary side of the relationships in the group will become the primary VDisks.

rc_consist_group_id | *rc_consist_group_name*

(Required) Specifies the ID or name of the consistency group to switch.

Description

This command applies to a consistency group. It is normally issued to reverse the roles of the primary and secondary virtual disks in a consistency group, perhaps as part of a failover process that is associated with a disaster recovery event. Write access to the former primary VDIs is lost and write access to the new primary VDIs is acquired. This command is successful when the consistency group is in a connected, consistent state, and when reversing the direction of the relationships would not lead to a loss of consistency, for example, when the consistency group is consistent and synchronized. The consistency group must be in one of the following states in order for the **switchrconsistgrp** command to process correctly:

- ConsistentSynchronized
- ConsistentStopped and Synchronized
- Idling and Synchronized

The consistency group moves to the ConsistentSynchronized state after the successful completion of this command. If you specify the **-primary** parameter and it is the same as the current primary, the command has no effect.

An invocation example

```
svctask switchrconsistgrp -primary aux rccopy2
```

The resulting output

```
No feedback
```

switchrrelationship

The **switchrrelationship** command reverses the roles of primary and secondary virtual disks in a stand-alone Metro Mirror or Global Mirror relationship when that relationship is in a consistent state.

Syntax

```
svctask -- switchrrelationship -- -primary [ master | aux ]
rc_rel_id | rc_rel_name
```

Parameters

-primary master | aux

(Required) Specifies whether the master disk or the auxiliary disk is to be the primary.

rc_rel_id | rc_rel_name

(Required) Specifies the ID or the name of the relationship to switch.

Description

The **switchrrelationship** command applies to a stand-alone relationship. It is rejected if it is used to try to switch a relationship that is part of a consistency group. It is normally issued to reverse the roles of the primary and secondary virtual disk in a relationship perhaps as part of a failover process during a disaster recovery event. Write access to the old primary disk is lost and write access to the new primary disk is acquired. This command is successful when the relationship is

in a connected, consistent state, and when reversing the direction of the relationship does not lead to a loss of consistency; that is, when the relationship is consistent and synchronized. The relationship must be in one of the following states in order for the **switchrcrelationship** command to process correctly:

- ConsistentSynchronized
- ConsistentStopped and Synchronized
- Idling and Synchronized

The relationship moves to the **ConsistentSynchronized** state after the successful completion of this command. If you specify the **-primary** parameter with the current primary, the command has no effect.

An invocation example

```
svctask switchrcrelationship -primary master rccopy2
```

The resulting output

No feedback

Chapter 17. Migration commands

The following commands enable you to work with migration options with the SAN Volume Controller.

migrateexts

The **migrateexts** command migrates extents from one managed disk to another.

Syntax

```
svctask -- 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 VDisk has more than one copy) Specifies the VDisk copy that the extents belong to.
- vdisk** *vdisk_id* | *vdisk_name*
(Required) Specifies the VDisk that the extents belong to.

Description

This command migrates a given number of extents from the source virtual disk and the managed disk that contains extents that are used to make up the virtual disk. The target is a managed disk within the same managed disk group.

If a large number of extents are being migrated, you can specify 1 - 4 threads. You can issue the **svcinfo 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 VDisk is offline. Correct the offline condition before attempting to migrate the VDisk.

Note: Migration activity on a single managed disk is limited to a maximum of 4 concurrent operations. This limit does not take into account whether the managed disk is the source or the destination target. If more than four migrations are scheduled for a particular managed disk, further migration operations are queued pending the completion of one of the currently running migrations. If a migration operation is stopped for any reason, a queued migration task can be started. However, if a migration is suspended, the current migration continues to use resources and a pending migration is not started. For example, the following setup is a possible initial configuration:

- MDiskGrp 1 has VDisk 1 created in it
- MDiskGrp 2 has VDisk 2 created in it
- MDiskGrp 3 has only one MDisk

With the previous configuration, the following migration operations are started:

- Migration 1 migrates VDisk 1 from MDiskGrp 1 to MDiskGrp 3, running with 4 threads.
- Migration 2 migrates VDisk 2 from MDiskGrp 2 to MDiskGrp 3, running with 4 threads.

Due to the previous limitations, the two migration operations do not always run at the same speed. MDiskGrp 3 has only one MDisk and the two migration operations have a total of 8 threads that are trying to access the one MDisk. Four threads are active. The remaining threads are in standby mode waiting to access the MDisk.

An invocation example

```
svctask migrateexts -vdisk vdisk4 -source mdisk4 -exts  
64 -target mdisk6 -threads 4
```

The resulting output


No feedback

migratetoimage

The **migratetoimage** command migrates data from a VDisk (image mode or managed mode) onto a new image mode VDisk copy. The target disk does not have to be in the same MDisk group as the source disk.

Syntax

```
▶▶ svctask — — migratetoimage — — [ -copy— id ] —————▶  
  
▶ -vdisk — [ source_vdisk_id —————▶  
          [ source_vdisk_name ] —————▶  
                                  [ -threads — number_of_threads ] —▶
```



Parameters

-vdisk *source_vdisk_id | name*
 (Required) Specifies the name or ID of the source VDisk to be migrated.

-copy *id*
 (Required if the specified VDisk has more than one copy) Specifies the VDisk copy to migrate from.

-threads *number_of_threads*
 (Optional) Specifies the number of threads to use during the migration of extents. You can specify 1 - 4 threads. The default number of threads is 4.

-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 MDisk group into which the MDisk must be placed, after the migration has completed.

Description

The **migratetoimage** command migrates the data of the user-specified source virtual disk onto the managed disk that is specified as the target. At completion of the command, the VDisk is classified as an image mode disk.

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 MDisk group.

The **migratetoimage** command fails if the target or source VDisk is offline. Correct the offline condition before attempting to migrate the VDisk.

The following example specifies that the user wants to migrate the data from vdisk1 onto mdisk5 and that the MDisk must be put into the MDisk group mdgrp2.

An invocation example

```
svctask migratetoimage -vdisk vdisk1 -mdisk mdisk5 -mdiskgrp mdgrp2
```

The resulting output

No feedback

migratevdisk

The **migratevdisk** command enables you to migrate an entire virtual disk from one managed disk group to another managed disk group.

Chapter 18. Tracing commands

Tracing commands capture information that can assist you with troubleshooting managed disks and virtual disks.

setdisktrace

Use the **setdisktrace** command to set a list of disks of a given type, to include in a disk trace.

Syntax

```
svctask -- setdisktrace -- -type [ mdisk | vdisk ] [ -set | -reset ]
[ -all | -objectid id_or_name_list ] [ -objectid id_or_name_list ]
```

Parameters

-type *mdisk* | *vdisk*

(Required) Specifies the object type for the disks.

-set

(Optional) Specifies the set argument. You cannot use the **-set** parameter with the **-reset** parameter.

-reset

(Optional) Specifies the reset argument. You cannot use the **-set** parameter with the **-reset** parameter.

-all

(Optional) Traces all disks of the specified type. You cannot use the **-all** parameter with the **-objectid** parameter.

-objectid *id_or_name_list*

(Optional) Specifies a list of one or more disk IDs or names. You cannot use the **-objectid** parameter with the **-all** parameter.

Description

The **setdisktrace** command marks the disks to be included in the next triggered trace.

The command is used with the **svctask settrace** command, which sets the options that result in a trace file and the data that is included in the trace file.

An invocation example

```
svctask setdisktrace -type mdisk -objectid
mdisk1:mdisk3:mdisk11:mdisk10:mdisk9:mdisk5 -reset
```

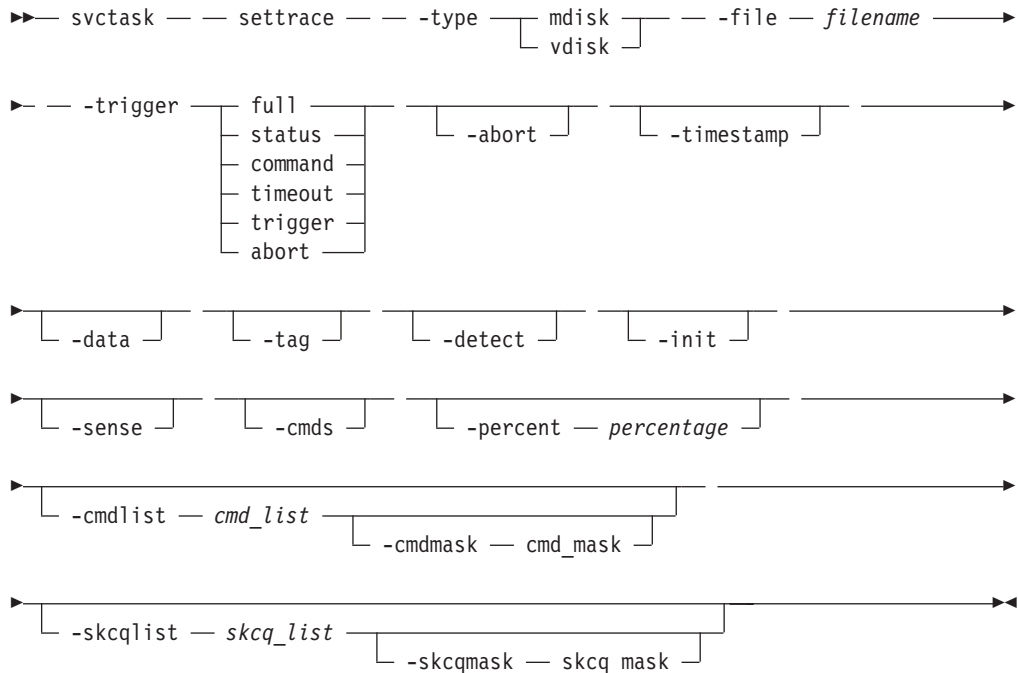
The resulting output

No feedback

settrace

The **settrace** command sets options to trace certain I/O operations through the system.

Syntax



Parameters

-type mdisk | vdisk

(Required) Specifies the type of objects to trace.

-file filename

(Required) Specifies the file name prefix for the trace file.

-trigger full | status | command | timeout | trigger | abort

(Required) Specifies an action for when the trace is started (triggered).

full Specifies to stop the trace when the trace buffer is full, for MDisks and VDIs.

status Sets a trigger for when the specified SCSI status (**-skcqlist**) is reported in sense data, for MDisks and VDIs.

command

Specifies a trigger for when the given SCSI command (**-cmdlist**) is sent, for MDisks and VDIs.

timeout

Sets a trigger for when a timeout occurs, for MDisks only.

trigger

Specifies to keep running until the trigger event, for MDisks only.

abort Sets a trigger for when an abnormal end occurs, for VDIs only.

-abort

(Optional) Adds abnormal ending details to the trace, for VDIs only.

- timestamp**
(Optional) Adds a time-stamp to each entry in the trace. A file name is created from the prefix plus a time-stamp. The file name is in the form *prefix_AAAAAA_YYMMDD_HHMMSS*, where *AAAAAA* is the panel name of the node generating the trace file.
- data**
(Optional) Adds I/O data to the trace.
- tag**
(Optional) Adds CCB tags to the trace, for MDisks only.
- detect**
(Optional) Adds MDisk discovery details to the trace, for MDisks only.
- init**
(Optional) Adds MDisk initialization details to the trace, for MDisks only.
- sense**
(Optional) Adds SCSI sense data to the trace, for VDIs only.
- cmds**
(Optional) Adds commands data to the trace, for VDIs only.
- percent**
(Optional) Specifies the trigger point in the trace file, which determines the amount of data to collect after the trigger point. The default value is **50**, which places the trigger point in the middle of the trace file.
- cmdlist** *cmd_list*
(Optional) Adds the commands in the *cmd_list* to the trace file.
- cmdmask** *cmd_mask*
(Optional) Adds the commands in the *cmd_mask* to the trace file. The **-cmdmask** parameter must be used with the **-cmdlist** parameter.
- skcqlist** *skcq_list*
(Optional) Specifies an SKCQ list, which adds only those SKCQ details to the trace file.
- skcqmask** *skcq_mask*
(Optional) Specifies an SKCQ mask, which adds only those SKCQ details to the trace file. The **-skcqmask** parameter must be used with the **-skcqlist** parameter.

Description

The **settrace** command sets the various I/O tracing options for managed disks or virtual disks. When the relevant disk type trace is subsequently triggered, the options specify the data to be included in the trace file.

The file name specifies a file name prefix to use when you are generating a trace file. The system appends the node panel name and a timestamp to the file name.

A maximum of 10 trace files are kept on the cluster. When the eleventh trace is made, the oldest existing trace file is overwritten.

The directory can also hold files that are retrieved from other nodes. These files are not counted. The cluster deletes the oldest file to maintain the maximum number of files.

An invocation example

```
svctask settrace -type vdisk -file tracedump -trigger abort  
-percent 100 -abort -timestamp
```

The resulting output

No feedback

starttrace

Use the **starttrace** command to begin tracing I/O operations that are based on the option currently set for the specified object type and the list of disks to trace.

Syntax

```
▶▶ svctask — — starttrace — — -type [ mdisk ]  
[ vdisk ]
```

Parameters

-type mdisk | vdisk

Specifies the object type to trigger.

Description

This command starts the collection of I/O tracing information. The trace file is generated according to the options that you specified in the **svctask settrace** command. The disks that are traced are those that are identified in the list that is set by the **svctask setdisktrace** command.

The traces are written to the **/dumps/iotrace** directory. You can view the contents of this directory using the **svcinfo lsiotracedumps** command.

An invocation example

```
svctask starttrace -type vdisk
```

The resulting output

No feedback

stoptrace

Use the **stoptrace** command to stop tracing operations for the specified disk type.

Syntax

```
▶▶ svctask — — stoptrace — — -type [ mdisk ]  
[ vdisk ]
```

Parameters

-type mdisk | vdisk

(Required) Specifies the object type to stop tracing.

Description

This command stops the tracing of I/O operations for the specified object type. A trace file is not generated if the trigger options have not been met.

An invocation example

```
svctask stoptrace -type mdisk
```

The resulting output

No feedback

Chapter 19. Attributes of the `-filtervalue` parameters

The `-filtervalue` parameter filters a view that is based on specific attribute values that relate to each object type. You can combine multiple filters to create specific searches, for example, `-filtervalue name=fred:status=online`. The help (`-filtervalue?`) specifies the attributes that are available for each object type.

The `-filtervalue` parameter must be specified with `attrib=value`. The `-filtervalue?` and `-filtervalue` parameters cannot be specified together.

Note: The qualifier characters left bracket (`<`) and right bracket (`>`) must be enclosed within double quotation marks (`"`). For example, `-filtervalue vdisk_count "<"4 or port_count ">"1`. It is also valid to include the entire expression within double quotation marks. For example, `-filtervalue "vdisk_count<4"`

When an attribute requires the `-unit` parameter, it is specified after the attribute. For example, `-filtervalue capacity=24 -unit mb`. The following input options are valid for the `-unit` parameter:

- **b** (bytes)
- **mb** (Megabytes)
- **gb** (Gigabytes)
- **tb** (Terabytes)
- **pb** (Petabytes)

Capacity values displayed in units other than bytes might be rounded.

When filtering on capacity, use a unit of bytes, `-unit b`, for exact filtering.

Table 13 provides a list of valid filter attributes, as well as descriptions, qualifiers and wildcards for each object type.

You can use the asterisk (`*`) character as a wildcard character when names are used. The asterisk character can be used either at the beginning or the end of a text string, but not both. Only one asterisk character can be used in a `-filtervalue` parameter.

Table 13. Valid filter attributes

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
cluster	<code>cluster_name</code> or <code>name</code>	=	Yes	The cluster name.
	<code>cluster_unique_id</code> or <code>id</code>	=, <, <=, >, >=	No	The cluster ID.

Table 13. Valid filter attributes (continued)

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
node	<i>node_name</i> or <i>name</i>	=	Yes	The node name.
	<i>id</i>	=, <, <=, >, >=	No	The node ID.
	<i>status</i>	=	No	The status of the node. The following values are valid for node <i>status</i> : <ul style="list-style-type: none"> • adding • deleting • online • offline • pending
	<i>IO_group_name</i>	=	Yes	The I/O group name.
	<i>IO_group_id</i>	=, <, <=, >, >=	No	The I/O group ID.
	<i>hardware</i>	=	No	The following values are valid for <i>hardware</i> type: 8F2, 8F4, 8G4, CF8, and 8A4.
io_grp	<i>HWS_name</i> or <i>name</i>	=	Yes	The I/O group name.
	<i>HWS_unique_id</i> or <i>id</i>	=, <, <=, >, >=	No	The I/O group ID.
	<i>node_count</i>	=, <, <=, >, >=	No	The number of nodes in the I/O group.
	<i>host_count</i>	=, <, <=, >, >=	No	The number of hosts associated with the io_grp.
controller	<i>controller_id</i> or <i>id</i>	=, <, <=, >, >=	No	The controller ID.

Table 13. Valid filter attributes (continued)

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
mdisk	<i>name</i>	=	Yes	The name of the MDisk.
	<i>id</i>	=, <, <=, >, >=	No	The ID of the MDisk.
	<i>controller_name</i>	=	Yes	The name of the controller the MDisk belongs to.
	<i>status</i>	=	No	The status of the MDisk. The following values are valid for MDisk <i>status</i> : <ul style="list-style-type: none"> • online • degraded_ports • degraded_paths • offline • excluded
	<i>mode</i>	=	No	The mode of the MDisk. The following values are valid for MDisk <i>mode</i> : <ul style="list-style-type: none"> • unmanaged • managed • image
	<i>mdisk_grp_name</i>	=	Yes	The MDisk group name.
	<i>mdisk_grp_id</i>	=, <, <=, >, >=	No	The MDisk group ID.
	<i>capacity</i>	=, <, <=, >, >=	No	The capacity. Requires the -unit parameter.
mdiskgrp	<i>name</i>	=	Yes	The MDisk group name.
	<i>storage_pool_id</i> or <i>id</i>	=, <, <=, >, >=	No	The MDisk group ID.
	<i>mdisk_count</i>	=, <, <=, >, >=	No	The number of MDisks in the group.
	<i>vdisk_count</i>	=, <, <=, >, >=	No	The number of VDIs in the group.
	<i>status</i>	=	No	The status of the MDisk group. The valid input options are online , degraded_ports , degraded_paths , excluded , and offline .
	<i>extent_size</i>	=, <, <=, >, >=	No	The extent size. (MB)

Table 13. Valid filter attributes (continued)

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
vdisk	<i>vdisk_name</i> or <i>name</i>	=	Yes	The name of the VDisk.
	<i>vdisk_id</i> or <i>id</i>	=, <, <=, >, >=	No	The ID of the VDisk.
	<i>IO_group_name</i>	=	Yes	The name of the I/O group.
	<i>IO_group_id</i>	=, <, <=, >, >=	No	The ID of the I/O group.
	<i>status</i>	=	No	The status of the VDisk. The valid input options for VDisk status are online , degraded , and offline .
	<i>mdisk_grp_name</i>	=	Yes	The MDisk group name.
	<i>mdisk_grp_id</i>	=, <, <=, >, >=	No	The MDisk group ID.
	<i>capacity</i>	=, <, <=, >, >=	No	The capacity. Requires the -unit argument.
	<i>type</i>	=	No	The VDisk type. The valid value options are seq , striped , and image .
	<i>FC_name</i>	=	Yes	The FlashCopy mapping name.
	<i>FC_id</i>	=, <, <=, >, >=	No	The FlashCopy mapping ID.
	<i>fc_map_count</i>	=, <, <=, >, >=	No	The number of VDisk mappings (either source or target).
	<i>copy_count</i>	=, <, <=, >, >=	No	The number of VDisk mirrored copies.
<i>RC_name</i>	=	Yes	The Metro Mirror relationship name.	
<i>RC_id</i>	=, <, <=, >, >=	No	The Metro Mirror relationship ID.	
vdisk_copy	<i>primary</i>	=	No	Indicates that this copy is the primary copy. The valid values are yes and no .
	<i>status</i>	=	No	The status of the MDisk group. Valid values are online , degraded , or offline .
	<i>sync</i>	=	No	Indicates whether the VDisk copy is synchronized. Valid values are true or false .
	<i>mdisk_grp_name</i>	=	Yes	The name of the MDisk group.
	<i>mdisk_grp_id</i>	=, <, <=, >, >=	No	The ID of the MDisk group.
	<i>type</i>	=	No	The type of the VDisk copy. The valid values are seq , striped , or image .
se_vdiskcopy	<i>mdisk_grp_id</i>	=, <, <=, >, >=	No	The ID of the MDisk group.
	<i>mdisk_grp_name</i>	=	Yes	The name of the MDisk group.
	<i>overallocation</i>	=	No	The percentage of overallocation, which is displayed as a number.
	<i>autoexpand</i>	=	No	Autoexpand flags. The valid values are on and off .

Table 13. Valid filter attributes (continued)

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
	<i>grainsize</i>	=, <, <=, >, >=	No	Space-efficient grain size. The valid values are 32 , 64 , 128 , or 256 .
host	<i>host_name</i> or <i>name</i>	=	Yes	The host name.
	<i>host_id</i> or <i>id</i>	=, <, <=, >, >=	No	The host ID.
	<i>port_count</i>	=, <, <=, >, >=	No	The number of ports.
	<i>iogrp_count</i>	=, <, <=, >, >=	No	The number of I/O groups that are associated with the host.
fcmap	<i>FC_mapping_name</i> or <i>name</i>	=	Yes	The FlashCopy mapping name.
	<i>FC_id</i> or <i>id</i>	=, <, <=, >, >=	No	The FlashCopy mapping ID.
	<i>source_vdisk_name</i>	=	Yes	The source VDisk name.
	<i>source_vdisk_id</i>	=, <, <=, >, >=	No	The source VDisk ID.
	<i>target_vdisk_name</i>	=	Yes	The target VDisk name.
	<i>target_vdisk_id</i>	=, <, <=, >, >=	No	The target VDisk ID.
	<i>group_name</i>	=	Yes	The consistency group name.
	<i>group_id</i>	=, <, <=, >, >=	No	The consistency group ID.
	<i>status</i>	=	No	The mapping status. The following values are valid for fcmap <i>status</i> : <ul style="list-style-type: none"> • idle_or_copied • preparing • prepared • copying • stopped • suspended • stopping • empty
	<i>copy_rate</i>	=, <, <=, >, >=	No	The background copy rate.
fcconsistgrp	<i>name</i>	=	Yes	The consistency group name.
	<i>FC_group_id</i> or <i>id</i>	=, <, <=, >, >=	No	The consistency group ID.
	<i>status</i>	=	No	The consistency group status. The following values are valid for fcconsistgrp <i>status</i> : <ul style="list-style-type: none"> • idle_or_copied • preparing • prepared • copying • stopped • suspended • stopping • empty

Table 13. Valid filter attributes (continued)

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
rrelationship	<i>RC_rel_id</i> or <i>id</i>	=, <, <=, >, >=	No	The Metro Mirror relationship ID.
	<i>RC_rel_name</i> or <i>name</i>	=	Yes	The Metro Mirror relationship name.
	<i>master_cluster_id</i>	=, <, <=, >, >=	No	The master cluster ID.
	<i>master_cluster_name</i>	=	Yes	The master cluster name.
	<i>master_vdisk_id</i>	=, <, <=, >, >=	No	The master VDisk ID.
	<i>master_vdisk_name</i>	=	Yes	The master VDisk name.
	<i>aux_cluster_id</i>	=, <, <=, >, >=	No	The aux cluster ID.
	<i>aux_cluster_name</i>	=	Yes	The aux cluster name.
	<i>aux_vdisk_id</i>	=, <, <=, >, >=	No	The aux VDisk ID.
	<i>aux_vdisk_name</i>	=	Yes	The aux VDisk name.
	<i>primary</i>	=	No	The relationship primary. The following values are valid for <i>primary</i> : <ul style="list-style-type: none"> • master • aux
	<i>consistency_group_id</i>	=, <, <=, >, >=	No	The Metro Mirror consistency group ID.
	<i>consistency_group_name</i>	=	Yes	The Metro Mirror consistency group name.
	<i>state</i>	=	Yes	The relationship state. The following values are valid for <i>state</i> : <ul style="list-style-type: none"> • inconsistent_stopped • inconsistent_copying • consistent_stopped • consistent_synchronized • idling • idling_disconnected • inconsistent_disconnected • consistent_disconnected
<i>progress</i>	=, <, <=, >, >=	No	The progress of the initial background copy (synchronization) for the relationship.	

Table 13. Valid filter attributes (continued)

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
rconsistgrp	<i>group_id</i> or <i>id</i>	=, <, <=, >, >=	No	The consistency group ID.
	<i>name</i>	=	Yes	The consistency group name.
	<i>master_cluster_id</i>	=, <, <=, >, >=	No	The master cluster ID.
	<i>master_cluster_name</i>	=	Yes	The master cluster name.
	<i>aux_cluster_id</i>	=, <, <=, >, >=	No	The aux cluster ID.
	<i>aux_cluster_name</i>	=	Yes	The aux cluster name.
	<i>primary</i>	=	No	The consistency group primary. The following values are valid for <i>primary</i> : <ul style="list-style-type: none"> • master • aux
	<i>state</i>	=	No	The consistency group state. The following values are valid for <i>state</i> : <ul style="list-style-type: none"> • inconsistent_stopped • inconsistent_copying • consistent_stopped • consistent_synchronized • idling • idling_disconnected • inconsistent_disconnected • consistent_disconnected • empty
<i>relationship_count</i>	=, <, <=, >, >=	No	The relationship count.	
user	<i>password</i>	=	No	Specifies if a password is associated with the user. The valid values are yes or no .
	<i>ssh_key</i>	=	No	Specifies if a Secure Shell (SSH) public key is associated with the user. The valid values are yes or no .
	<i>remote</i>	=	No	Specifies if the user authenticates to the cluster using a remote authentication service. The valid values are yes or no .
	<i>usergrp_id</i>	=, <, <=, >, >=	No	The ID of the user group.
	<i>usergrp_name</i>	=	Yes	The name of the user group.
usergrp	<i>role</i>	=	No	The role associated with all users that belong to this user group. The valid values are Monitor , CopyOperator , Service , Administrator , or SecurityAdmin .
	<i>remote</i>	=	No	Specifies if the user group is used to set the role of remote users. The valid values are yes or no .

Table 13. Valid filter attributes (continued)

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
clusterip	<i>port_id</i>	=, <, <=, >, >=	No	The port ID. The valid values are 1 or 2 .
	<i>cluster_name</i>	=	Yes	The cluster name.
	<i>cluster_id</i>	=, <, <=, >, >=	No	The cluster ID.

Chapter 20. Overview of the list dumps commands

The list dumps commands return a list of dumps in the appropriate directory.

SAN Volume Controller dumps are contained in the following directory structure:

- /dumps
- /dumps/audit
- /dumps/cimom
- /dumps/elogs
- /dumps/feature
- /dumps/iostats
- /dumps/iotrace
- /dumps/mdisk

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

An audit log keeps track of action commands that are issued through an SSH session or from the SAN Volume Controller Console. To list the audit log files in the /dumps/audit directory on the current configuration node or on the designated node, issue the `svcinfo lsauditlogdumps` command. To list a specified number of the most recently audited commands, issue the `svctask catauditlog` command. To dump the contents of the audit log to a file on the current configuration node, issue the `svctask dumpauditlog` command. This command also clears the contents of the audit log.

Dumps contained in the /dumps/cimom directory are created by the CIMOM (Common Information Model Object Manager) that runs on the SAN Volume Controller cluster. These files are produced during normal operations of the CIMOM. To list all the files in the /dumps/cimom directory, issue the `svcinfo lscimomdumps` command.

Dumps that are contained in the /dumps/elogs directory are dumps of the contents of the error and event log at the time that the dump was taken. An error or event log dump is created by using the `svctask dumperrlog` command. This dumps the contents of the error or event log to the /dumps/elogs directory. If no file name prefix is supplied, the default `errlog_` is used. The full default file name is `errlog_NNNNNN_YYMMDD_HHMMSS`, where `NNNNNN` is the node front panel name. If the command is used with the `-prefix` parameter, the prefix value is used instead of `errlog`. To list all dumps in the /dumps/elogs directory, issue the `svcinfo lserrlogdumps` command.

Dumps contained in the /dumps/feature directory are dumps of the featurization log. A featurization log dump is created by using the `svctask dumpinternallog` command. This dumps the contents of the featurization log to the /dumps/feature directory to a file called `feature.txt`. Only one of these files exists, so every time the `svctask dumpinternallog` command is run, this file is overwritten. To list all dumps in the /dumps/feature directory, issue the `svcinfo lsfeaturedumps` command.

Dumps that are contained in the /dumps/iostats directory are dumps of the per-node I/O statistics for disks on the cluster. An I/O statistics dump is created

| by using the svctask startstats command. As part of this command, you can specify
| a time interval for the statistics to be written to the file; the default is 15 minutes.
| Every time the time interval is encountered, the I/O statistics that have been
| collected are written to a file in the /dumps/iostats directory. The file names that
| are used for storing I/O statistics dumps are
| Nm_stats_NNNNNN_YYMMDD_HHMMSS,
| Nv_stats_NNNNNN_YYMMDD_HHMMSS, and
| Nn_stats_NNNNNN_YYMMDD_HHMMSS, where NNNNNN is the node front
| panel name for the MDisk, VDisk, or node. To list all dumps in the
| /dumps/iostats directory, issue the svcinfo lsiostatsdumps command.

Dumps that are contained in the /dumps/iotrace directory are dumps of I/O trace data. The type of data that is traced depends on the options specified by the svctask settrace command. The collection of the I/O trace data is started by using the svctask starttrace command. The I/O trace data collection is stopped when the svctask stoptrace command is used. It is when the trace is stopped that the data is written to the file. The file name is *prefix_NNNNNN_YYMMDD_HHMMSS*, where *prefix* is the value entered for the **-filename** parameter in the svctask settrace command, and NNNNNN is the node front panel name. To list all dumps in the /dumps/iotrace directory, issue the svcinfo lsiotrace dumps command.

| Dumps that are contained in the /dumps/mdisk directory are copies of solid-state
| drive (SSD) MDisk internal logs. These dumps are created using the svctask
| triggermdiskdump command. The file name is
| mdiskdump_NNNNNN_MMMM_YYMMDD_HHMMSS, where NNNNNN is the
| front panel name of the node that contains the MDisk, and MMMM is the decimal
| ID of the MDisk. To list all dumps in the /dumps/mdisk directory, issue the
| svcinfo lsmdiskdumps command.

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, are named *NNNNNN.trc*.

To list all dumps in the /dumps directory, issue the svcinfo ls2145dumps command.

The final option available in the list dumps command series is the svcinfo lssoftware dumps command. This command lists the contents of the /home/admin/upgrade directory; files are copied to this directory during software upgrades.

All of the list dumps commands can accept a node identifier as input. If this identifier is not specified, the list of files on the current configuration node are displayed. If the node identifier is specified, the list of files on that node are displayed.

Because files can only be copied from the current configuration node (using secure copy), you can issue the svctask cpdumps command to copy the files from a nonconfiguration node to the current configuration node.

Chapter 21. Information commands

The following commands enable you display specific types of SAN Volume Controller information.

Note: IDs are assigned at run-time by the system and cannot be relied upon to be the same after configuration restoration. Therefore, use object names instead of IDs whenever possible.

caterrlog

The **caterrlog** command displays the contents of the cluster error log and event log.

Syntax

```
svcinfo -- caterrlog [-nohdr] [-delim delimiter]
                  [-config] [-unfixed] [-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.

-config

(Optional) Specifies to list the configuration events.

-unfixed

(Optional) Specifies to list the unfixed errors.

-first number_of_entries_to_return

(Optional) Displays the first *x* number of entries in the log, where *x* is the number that is entered by the user. The value of *x* can be 1 - 256.

Description

This command displays a list of the specified error log entries. When no parameters are used, all error log entries are listed, to a maximum of 256 entries.

The list can be filtered to only include configuration events or unfixed errors by specifying the **-config** or **-unfixed** parameters.

The **-first** parameter results in the display of the first *x* number of records, where *x* is the number that is entered for the **-first** parameter.

An invocation example

```
svcinfo caterrlog -delim :
```

The resulting output

```
id:type:fixed:SNMP_trap_raised:error_type:node_name:sequence_number:
root_sequence_number:first_timestamp:last_timestamp:number_of_errors:error_code
:copy_id
0:cluster:no:no:6:node1:100:100:030407052547:030407052547:1:00981001
0:fc_card:no:no:1:node1:101:101:030407052547:030407052547:1:00073001
1:node:no:no:1:node1:102:102:030407052547:030407052547:1:00074001
0:cluster:no:no:6:node1:103:100:030407052547:030407052547:1:00981001
1:fc_card:no:no:1:node1:104:104:030407052632:030407052632:1:00073003
0:node:no:no:6:node1:105:105:030407082202:030407082717:2:00980500
2:remote:no:no:6:n/a:106:106:030407090117:030407090117:1:00985002
1:node:no:no:5:node1:0:0:030407052546:030407052546:1:00990383
0:cluster:no:no:5:node1:0:0:030407080630:030407080630:1:00990117
0:mdisk_grp:no:no:5:node1:0:0:030407081610:030407081610:1:00990148
128:mdisk_grp:no:no:5:node1:0:0:030407081610:030407081610:1:00990173
1:mdisk_grp:no:no:5:node1:0:0:030407081619:030407081619:1:00990148
0:vdisk:no:no:5:node1:0:0:030407081836:030407081836:1:00990169:0
1:vdisk:no:no:5:node1:0:0:030407081843:030407081843:1:00990169:0
0:vdisk:no:no:5:node1:0:0:030407081854:030407081854:1:00990169:0
0:vdisk:no:no:5:node1:0:0:030407082015:030407082015:1:00990169:0
0:vdisk:no:no:5:node1:0:0:030407082145:030407082145:1:00990169:0
0:vdisk:no:no:5:node1:0:0:030407082148:030407082148:1:00990169:0
0:vdisk:no:no:5:node1:0:0:030407082158:030407082158:1:00990169:0
1:vdisk:no:no:5:node1:0:0:030407082213:030407082213:1:00990169:0
0:host:no:no:5:node1:0:0:030407082441:030407082441:1:00990106
1:host:no:no:5:node1:0:0:030407082457:030407082457:1:00990106
2:host:no:no:5:node1:0:0:030407082523:030407082523:1:00990106
0:flash:no:no:5:node1:0:0:030407082704:030407082704:1:00990184
1:node:no:no:5:node1:0:0:030407082716:030407082716:1:00990501
1:node:no:no:5:node1:0:0:030407082722:030407082722:1:00990501
1:fc_const_grp:no:no:5:node1:0:0:030407083141:030407083141:1:00990204
2:fc_const_grp:no:no:5:node1:0:0:030407083143:030407083143:1:00990204
3:fc_const_grp:no:no:5:node1:0:0:030407083145:030407083145:1:00990204
0:flash:no:no:5:node1:0:0:030407083318:030407083318:1:00990185
0:flash:no:no:5:node1:0:0:030407083355:030407083355:1:00990185
0:flash:no:no:5:node1:0:0:030407085753:030407085753:1:00990185
1:remote:no:no:5:node1:0:0:030407085932:030407085932:1:00990225
2:vdisk:no:no:5:node1:0:0:030407085959:030407085959:1:00990169:0
3:vdisk:no:no:5:node1:0:0:030407090004:030407090004:1:00990169:0
4:vdisk:no:no:5:node1:0:0:030407090013:030407090013:1:00990169:0
2:remote:no:no:5:node1:0:0:030407090106:030407090106:1:00990225
255:rc_const_grp:no:no:5:node1:0:0:030407090323:030407090323:1:00990240
254:rc_const_grp:no:no:5:node1:0:0:030407090327:030407090327:1:00990240
253:rc_const_grp:no:no:5:node1:0:0:030407090333:030407090333:1:00990240
2:remote:no:no:5:node1:0:0:030407090442:030407090442:1:00990226
1:vdisk:no:no:5:node1:0:0:030407090820:030407090820:1:00990182:0
3:vdisk:no:no:5:node1:0:0:030407090825:030407090825:1:00990182:0
```

caterrlogbyseqnum

The **caterrlogbyseqnum** command displays all the errors with the sequence number or root cause number that has been specified by the user.

Syntax

```
svcinfo -- caterrlogbyseqnum -- [-num -- sequence_number] [-root -- root_cause_number] [-nohdr] [-delim -- delimiter]
```

Parameters

-num *sequence_number*

(Required if **-root** is not specified) Specifies the sequence number to view.

-root *root_cause_number*

(Required if **-num** is not specified) Specifies the root sequence number. All errors that are marked with this root cause 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 single error log entry, determined by the sequence number that is specified as the **-num** parameter.

If the **-root** parameter is used, the log is searched for all entries that are marked with a root cause sequence number, as specified. A list of all entries marked with this root cause is displayed.

An invocation example

```
svcinfo caterrlogbyseqnum -num 100 -delim :
```

The resulting output

```
id:type:fixed:SNMP_trap_raised:error_type:node_name:sequence_number:
root_sequence_number:first_timestamp:last_timestamp:number_of_errors:
error_code:copy_id
0:cluster:no:no:6:node1:100:100:030407052547:030407052547:1:00981001:
```

ls2145dumps

Use the **ls2145dumps** command to display a list of error conditions detected by a node, and the associated output files from the **/dumps** directory.

Syntax

```
svcinfo ls2145dumps [-nohdr] [-delim delimiter] [node_id | node_name]
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim delimiter

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

node_id | node_name

(Optional) Displays a list of the available dumps of the given type for the specified node ID or name. If you do not specify a node, the available dumps on the configuration node are listed.

Description

This command displays a list of error conditions detected by a node. If you do not specify a node, the dumps that are available on the configuration node are listed. The command displays the associated output files that are available from the **/dumps** directory.

An invocation example

```
svcinfo ls2145dumps -delim :
```


The resulting output

```
id:2145_filename  
0:000108.trc.old  
1:dump.000108.030328.144007  
2:000108.trc
```

lscimomdumps

Use the **lscimomdumps** command to display a list of files in the **/dumps/cimom** directory.

Syntax



Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, then the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

node_id | *node_name*

(Optional) Specifies the node ID or name to list the available dumps of the given type. If you do not specify a node, the dumps that are available on the configuration node are listed.

Description

This command displays a list of Common Information Model object manager (CIMOM) dumps. These dumps are created when you use a CIM client with the CIMOM of the SAN Volume Controller cluster. If you do not specify a node, the dumps that are available on the configuration node are listed. The command displays files from the **/dumps/cimom** directory.

An invocation example

```
svcinfo lscimomdumps
```

The resulting output

```
id          cimom_filename
0          mkrepositorylog.004565
1          PegasusTrace.004565
```

2	PegasusStandard.004565
3	PegasusAudit.004565
4	PegasusError.004565
5	PegasusDebug.004565

lscopystatus

Use the **lscopystatus** command to determine whether any file copies are currently in progress.

Syntax

```

>> svcinfo -- lscopystatus -- [-nohdr]
<<
[-delim -- delimiter]

```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

Description

This command displays an indicator that shows if a file copy is currently in progress. Only one file can be copied in the cluster at a time.

An invocation example

```
svcinfo lscopystatus
```

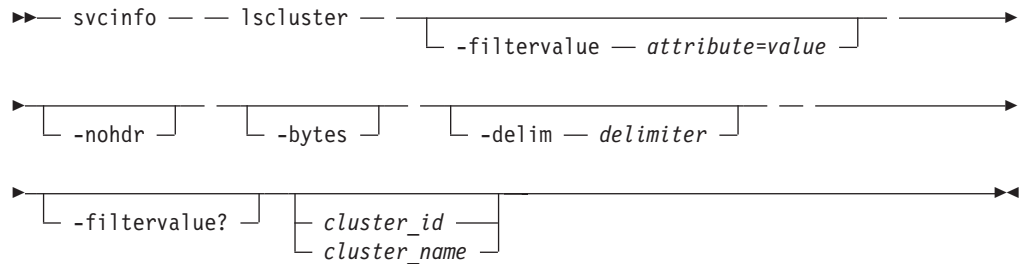
The resulting output

```
status
active
```

lscluster

The **lscluster** command returns a concise list or a detailed view of a cluster.

Syntax



Parameters

-filtervalue *attribute=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are displayed. If a capacity is specified, the units must also be included.

Note: Some filters allow the asterisk character (*) when you enter the command. The following rules apply to the use of wildcard characters with the SAN Volume Controller CLI:

- The wildcard character is an asterisk (*).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, you must enclose the filter entry within double quotation marks (""), as follows:

```
svcinfo lscluster -filtervalue "name=md*"
```

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-bytes

(Optional) Specifies that you want the report to display all capacities as bytes.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

-filtervalue?

(Optional) displays a list of filters that can be applied against this view. The following filter attributes are valid for the **lscluster** command:

- **cluster_name**
- **cluster_unique_id**
- **id**
- **name**

cluster_id | *cluster_name*

(Optional) Specifies the name or ID of a cluster. When you use this parameter, the detailed view of the specific cluster is displayed and any value that you specified by the **-filtervalue** parameter is ignored. If you do not specify the *cluster_id* | *cluster_name* parameter, the concise view of all clusters that match the filtering requirements that are specified by the **-filtervalue** parameter are displayed.

Description

This command displays a concise list or a detailed view of a cluster.

Table 14 provides the attribute values that can be displayed as output view data.

Table 14. Attribute values

Attribute	Possible Values
location	local, remote
statistics status	on, off
email_state	running, stopped, invalid
partnership	fully_configured, partially_configured_local, partially_configured_local_stopped, not present, fully_configured_stopped, fully_configured_remote_stopped, fully_configured_local_excluded, fully_configured_remote_excluded, fully_configured_exceeded

Information about the remote cluster is reported by the `lscluster` command if the `mkpartnership` command has been issued from the local cluster to the remote cluster; for example, if the partnership has been at least partially established from the local cluster.

You can issue the `svcinfolcluster` command to display a detailed view of the cluster.

```
svcinfolcluster -delim : 10030a007e5
```

where `10030a007e5` is the cluster name.

Detailed view shows the fields described for remote clusters only; if the cluster **Location** is **local**, then **Partnership** and **Bandwidth** do not apply (and are not defined or provided). For a remote cluster, these fields indicate the following information:

- **Location:** remote
- **Partnership:**

fully_configured

The `mkpartnership` command has been issued in both directions and the remote cluster is online and available.

partially_configured_local

The `mkpartnership` command has only been issued from the local cluster to the remote cluster. The remote cluster is online and available for partnership.

partially_configured_local_stopped

The mkpartnership command has only been issued from the local cluster to the remote cluster. The chpartnership command with the **stop** parameter has been issued from the local cluster, and the remote cluster is online and available. You need to issue the chpartnership command with the **start** parameter on the local cluster, and mkpartnership on the remote cluster.

Not Present

The mkpartnership command has been issued from the local cluster to the remote cluster, and the remote cluster is not available. Either the remote cluster is offline, or it is not connected to the local cluster.

fully_configured_stopped

The mkpartnership command has been issued in both directions and the remote cluster is online and available. The chpartnership command with the **stop** parameter has been issued from the local cluster.

fully_configured__remote_stopped

The mkpartnership command has been issued in both directions and the remote cluster is online and available. The chpartnership command with the **stop** parameter has been issued from the remote cluster.

fully_configured__local_excluded

The mkpartnership command has been issued in both directions. The local cluster has excluded the connection to the remote cluster due to too many problems, or either cluster in the partnership is unable to sustain the I/O workload for the Metro Mirror or Global Mirror relationships.

fully_configured__remote_excluded

The mkpartnership command has been issued in both directions. The remote cluster has excluded the connection to the local cluster due to too many problems, or either cluster in the partnership is unable to sustain the I/O workload for the Metro Mirror or Global Mirror relationships.

fully_configured__exceeded

There are too many clusters in the cluster network, and the partnership from the local cluster to the remote has been disabled. Refer to the 1710/1720 errors in the cluster error log at the local and remote clusters.

- **Bandwidth:** The bandwidth available on the intercluster link for background copy, in megabytes per second (MBps).

A concise invocation example

```
svcinfolcluster -delim :
```

The concise resulting output

```
id:name:location:partnership:bandwidth:id_alias  
000002006420A162:cluster0:local:::000002006420A162
```

A detailed invocation example

```
svcinfolcluster -delim : cluster1
```

The detailed resulting output

```
id:1521071282978998  
name:cluster1  
location:local  
partnership:fully_configured_stopped  
bandwidth:  
total_mdisk_capacity:1673.3GB
```

```

space_in_mdisk_grps:0
space_allocated to vdisks:0.00MB
total_free_space:1673.3GB
statistics_status:on
statistics_frequency:15
required_memory:8192
cluster_locale:en_US
time_zone:522 UTC
code_level:5.1.0.0
FC_Port_Speed:2Gb
console_IP:9.20.165.12:123
id_alias:1521071282978998
gm_link_tolerance:300
gm_inter_cluster_delay_simulation:0
gm_intra_cluster_delay_simulation:0
email_reply:fred@mycompany.com
email_contact:Fred Higgins
email_contact_primary:01202 123456
email_contact_alternate:44-202-876543-4455
email_contact_location:London Thames Bank
email_state:running
inventory_mail_interval:0
total_vdiskcopy_capacity:40.00GB
total_used_capacity:22.50GB
total_overallocation:67
total_vdisk_capacity:30.00GB
cluster_ntp_ip_address:9.20.199.43
cluster_isns_IP_address:
iscsi_auth_method:none
iscsi_chap_secret:
auth_service_configured:yes
auth_service_enabled:yes
auth_service_url:https://1.2.3.4/login
auth_service_user_name:secadmin
auth_service_pwd_set:yes
auth_service_cert_set:yes
relationship_bandwidth_limit:25

```

lsclustercandidate

The `lsclustercandidate` command lists the clusters that are available for setting up a partnership with the local cluster. This is a prerequisite for creating intercluster Metro or Global Mirror relationships.

Syntax

```

▶▶ svcinfo — — lsclustercandidate — [ -nohdr ] —————▶
▶ [ -delim — delimiter ] —————▶▶

```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

Description

This command displays a list of clusters that are available as candidate partner clusters to form a Metro Mirror or Global Mirror partnership between two clusters.

|
|
|
|
|
|
|
|
|

Output from the command shows the cluster ID, name, and configured status of the remote candidate cluster. The remote candidate cluster forms a partnership with the local cluster when you use the `svctask mkpartnership` command. The remote cluster shows the partnership status as `partially_configured_local_stopped` or `partially_configured_local` when you use the `svcinflscluster` command. The `svcinflsclustercandidate` command displays the configured status of those remote clusters that have formed a partnership with the local cluster.

An invocation example

```
svcinflsclustercandidate
```

The resulting output

id	configured	cluster_name
0000010034E0F430	no	ldcluster26

lsclusterip

The `lsclusterip` command returns a list of the cluster management IP addresses configured for each port.

Syntax

```
svcinflsclusterip [-nohdr] [-delim delimiter] [-filtervalue attribute=value] [-filtervalue? cluster_id] [-filtervalue? cluster_name]
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

-filtervalue *attribute=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are displayed. If a capacity is specified, the units must also be included.

Note: Some filters allow the asterisk character (*) when you enter the command. The following rules apply to the use of wildcard characters with the SAN Volume Controller CLI:

- The wildcard character is an asterisk (*).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, you must enclose the filter entry within double quotation marks (""), as follows:

```
svcinfo lsclusterip -filtervalue "cluster_name=md*"
```

-filtervalue?

(Optional) displays a list of filters that can be applied against this view. The following filter attributes are valid for the `lsclusterip` command:

- **port_id**
- **cluster_name**
- **cluster_id**

cluster_id | *cluster_name*

(Required) Specifies the name or ID of a cluster.

Description

This command displays a list of the cluster management IP addresses configured for each port.

A concise invocation example

```
svcinfo lsclusterip -delim ,
```

The concise resulting output

```
cluster_id,cluster_name,location,port_id,IP_address,service_IP_address,subnet_mask,
gateway,IP_address_6,service_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

```
svcinfo lsclusterip 000002006CC0B71A
```

The detailed resulting output

```
cluster_id 000002006CC0B71A
cluster_name c11
location local
port_id 1
IP_address 192.168.1.2
service_IP_address DHCP
subnet_mask 255.255.255.0
gateway 192.168.1.1
IP_address_6 2001:0db8:85a3:0000:0000:8a2e:0370:7334
service_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
service_IP_address DHCP
subnet_mask 255.255.255.0
gateway 192.168.1.1
IP_address_6 2001:0db8:85a3:0000:0000:8a2e:0370:7334
service_IP_address_6 2001:0db8:85a3:0000:0000:8a2e:0370:7334
gateway_6 2001:0db8:85a3:0000:0000:8a2e:0370:7334
prefix_6 64
```

Iscontroller

The `lscontroller` command returns a concise list, or a detailed view, of controllers that are visible to the cluster.

The list report style can be used to obtain two styles of report:

- A list containing concise information about controllers. (Each entry in the list corresponds to a single controller.)
- The detailed information about a single, user-specified controller.

Syntax

```
svcinfo lscontroller [-filtervalue attrib=value]
                    [-nohdr] [-delim delimiter] [-filtervalue?]
                    [controller_id controller_name]
```

Parameters

-filtervalue *attrib=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are returned. If a capacity is specified, the units must also be included.

Note: Some filters allow the use of a wildcard when you enter the command. The following rules apply to the use of wildcards with the SAN Volume Controller CLI:

- The wildcard character is an asterisk (*).
- The command can contain a maximum of one wildcard, 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:

```
svcinfo lscontroller -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 the valid filter attributes. The following filter attributes for the `svcinfo lscontroller` command are valid:

- `controller_id`
- `id`

controller_id | controller_name

(Optional) Specifies the name or ID of a controller. When you use this parameter, the detailed view of the specific controller is returned and any value that is specified by the **-filtervalue** parameter is ignored. If you do not specify the *controller_id | controller_name* parameter, the concise view displays all controllers matching the filtering requirements that are specified by the **-filtervalue** parameter.

Description

This command returns a concise list, or a detailed view, of controllers visible to the cluster.

The following values are applicable to the data in the output views:

degraded no, yes

To differentiate the name of a storage controller from the name shown on the cluster, list the storage controllers by issuing the `svcinfo lscontroller` command. Record the controller name or ID for the controller that you want to determine. For the controller in question, issue the `svcinfo lscontroller controller name | id` command, where *controller name | id* is the controller name or ID. Record the worldwide node name (WWNN) for the controller. You can use the WWNN to

determine the actual storage controller by launching the native controller user interface, or by using the command line tools it provides to verify the actual controller that has the WWNN.

Notes:

1. The *mdisk_link_count* value is the number of MDisks currently using this controller as their active controller.
2. The *max_mdisk_link_count* value is the highest value that the *mdisk_link_count* has reached since it was last reset to the *mdisk_link_count* value. This value is reset by specific maintenance procedures and when the cluster error log is cleared.
3. The controller *path_count* value is the sum of the MDisk *path_count* values for MDisks currently using this controller port as their active controller port.
4. The controller *max_path_count* value is the highest value that the controller *path_count* has reached since it was last reset to the *path_count* value. This value is reset by specific maintenance procedures and when the cluster error log is cleared.
5. The *allow_quorum* value shows whether the controller is currently enabled to support quorum disks. Quorum support is either enabled or disabled depending on the controller hardware type.

A concise invocation example

```
svcinfolsccontroller -delim :
```

The concise resulting output

```
id:controller_name:ctrl_s/n:vendor_id:product_id_low:product_id_high
7:controller7:3EK0J5Y8:SEAGATE :ST373405:FC
8:controller8:3EK0J6CR:SEAGATE :ST373405:FC
9:controller9:3EK0J4YN:SEAGATE :ST373405:FC
10:controller10:3EK0GKGH:SEAGATE :ST373405:FC
11:controller11:3EK0J85C:SEAGATE :ST373405:FC
12:controller12:3EK0JBR2:SEAGATE :ST373405:FC
13:controller13:3EKYNJF8:SEAGATE :ST373405:FC
14:controller14:3EK0HVTM:SEAGATE :ST373405:FC
```

A detailed invocation example

```
svcinfolsccontroller -delim = 7
```

The detailed resulting output

```
id=7
controller_name=controller7
WWNN=20000004CF2412AC
mdisk_link_count=1
max_mdisk_link_count=1
degraded=no
vendor_id=SEAGATE
product_id_low=ST373405
product_id_high=FC
product_revision=0003
ctrl_s/n=3EK0J5Y8
allow_quorum=no
WWPN=22000004CF2412AC
path_count=1
max_path_count=1
WWPN=21000004CF2412AC
path_count=0
max_path_count=0
```

Iscontrollerdependentvdisks

The **Iscontrollerdependentvdisks** command lists the VDisks that are dependent on the specified controller.

Syntax

```
►► svcinfo — — Iscontrollerdependentvdisks — — [ controller_id_list | controller_name_list ]
```

Parameters

controller_id_list | *controller_name_list*

Specifies one or more controller IDs, controller names, or both. Separate multiple controllers using the colon character (:).

Description

The **Iscontrollerdependentvdisks** command lists the VDisks that are dependent on the status of the specified controllers. If a controller goes offline, the dependent VDisks also go offline. Before taking a controller offline for maintenance, you can use the command to ensure that you do not lose access to any VDisks.

If you have multiple controllers configured as a single subsystem, you must specify all of the controllers in the subsystem, using a single command invocation.

The **Iscontrollerdependentvdisks** command also checks for quorum disks on the specified controller list. If any quorum disks are on the specified controller list, the command returns an error. All quorum disks must be moved before performing any maintenance. After moving quorum disks, reissue the command to list the dependent VDisks.

Note: The command lists the VDisks that are dependent on the controllers at the time the command is run; subsequent changes to your system require rerunning the command.

An invocation example

```
svcinfo Iscontrollerdependentvdisks controller0
```

The resulting output

```
vdisk_id vdisk_name
0 vdisk0
1 vdisk1
2 vdisk2
```

Iscurrentuser

Use the **Iscurrentuser** command to display the name and role of the logged-in user.

Syntax

```
►► svcinfo — — Iscurrentuser — — [ -nohdr ]
```

`-delim — delimiter`

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

Description

This command displays the name and role of the current user.

An invocation example

```
svcinfolcurrentuser
```

The resulting output

```
name superuser
role SecurityAdmin
```

lsdiscoverystatus

Use the **lsdiscoverystatus** command to determine whether a discovery operation is in progress.

Syntax

`svcinfoldiscoverystatus`

`-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 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.

An invocation example

```
svcinfolddiscoverystatus
```

The resulting output

```
status  
inactive
```

lsemailer

The `lsemailer` command returns a concise list or a detailed view of e-mail servers that are configured on the cluster.

Syntax

```
svcinfolddiscoverystatus lsemailer [-nohdr] [-delim delimiter] [email_server_name] [email_server_id]
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by

a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (**:**) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

email_server_name | *email_server_id*

(Optional) Specifies the name or ID of an existing e-mail server that must be listed.

Description

Use this command to display a concise list or a detailed view of e-mail servers that are configured on the cluster.

A concise invocation example

```
svcinfolmailserver -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

```
svcinfolmailserver email0
```

The detailed resulting output

```
id 0
name emailserver0
IP_address 192.135.60.3
port 25
```

lmailuser

The **lmailuser** command generates a report that lists the e-mail event notification settings for all e-mail recipients, an individual e-mail recipient, or a specified type (local or support) of e-mail recipient.

Syntax

```
svcinfolmailuser [-type support|local] [-delim delimiter] [id_or_name]
```

Parameters

-type support | **local**

(Optional) Specifies the types of e-mail recipients you want to view, either customer or support based as determined by the following definitions:

support

Address of the support organization that provides vendor support.

local All other addresses.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, a colon separates all items of data in a concise view; the spacing of columns does not occur. In a detailed view, the data is separated from its header by a colon.

id_or_name

(Optional) Specifies the user ID or user name of the e-mail event recipient for whom you want to see the e-mail notification settings.

Description

When you issue this command, a report is displayed that lists the e-mail event notification settings for all e-mail recipients, an individual e-mail recipient, or a specified type (local or support) of e-mail recipient. The concise and detailed views report the same information.

An invocation example

The following command lists information for all e-mail recipients using the e-mail event notification facility, in a concise view:

```
svcinfolsemailer -delim :
```

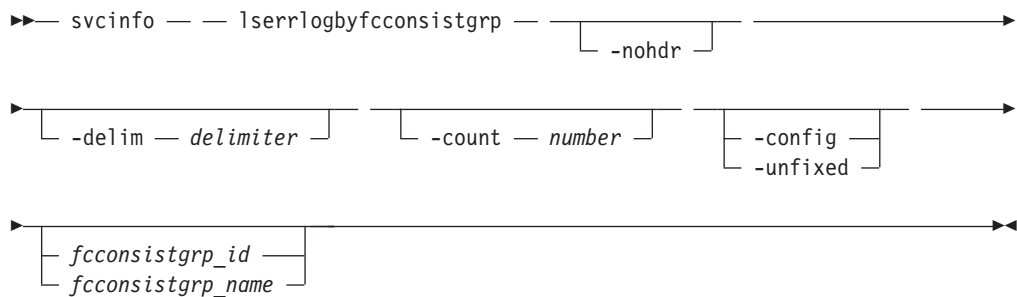
The resulting output

```
id:name:address:user_type:error:warning:info:inventory
1:Support:callhome1@de.ibm.com:support:on:off:off:off
2:Fred:fred_house@my_company.co.uk:local:on:on:on:off
3:Log:our_log@my_company.co.uk:local:on:on:on:on
```

Iserrlogbyfcconsistgrp

The **Iserrlogbyfcconsistgrp** command displays errors and events in the log that are related to FlashCopy consistency groups.

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.

-count *number*

(Optional) Specifies the maximum number of errors or events to list.

-config

(Optional) Lists only configuration events.

-unfixed

(Optional) Lists only unfixed errors.

fcconsistgrp_id | *fcconsistgrp_name*

(Optional) Specifies the object ID that filters the log.

Description

This command displays a list of the errors and events in the log that are related to FlashCopy consistency groups. The list can be filtered further by specifying an object ID or object name. The list displays only the errors and events that have been logged against the specified object. The list can also be filtered to show only the configuration events or the unfixed errors for the given object type or object ID. Similarly, the last *x* number of entries for a given object type or object ID can be listed.

Note: Although there is an object type of *unknown* is displayed in the error log, there is no command available to filter this object type.

An invocation example

```
svcinfo lserrlogbyfcconsistgrp -delim :
```

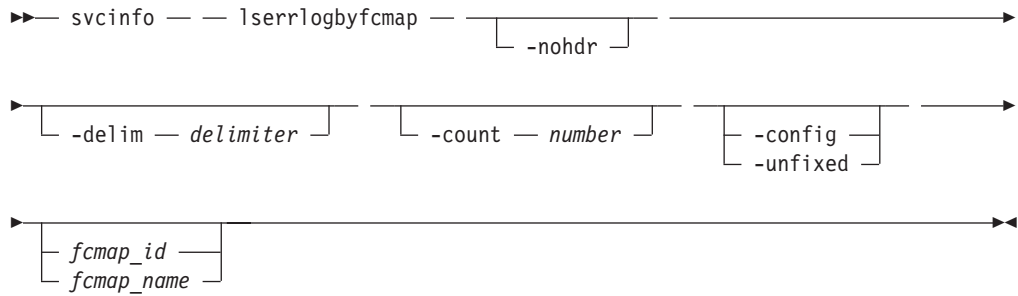
The resulting output

```
id:type:fixed:SNMP_trap_raised:error_type:node_name:sequence_number:
root_sequence_number:first_timestamp:last_timestamp:number_of_errors:error_code
3:fc_const_grp:no:no:5:node1:0:0:030407083145:030407083145:1:00990204
2:fc_const_grp:no:no:5:node1:0:0:030407083143:030407083143:1:00990204
1:fc_const_grp:no:no:5:node1:0:0:030407083141:030407083141:1:00990204
```

lserrlogbyfcmap

The **lserrlogbyfcmap** command displays a list of the errors and events in the log that are related to FlashCopy mappings.

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 even if the **-nohdr** parameter is specified.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

-count *number*

(Optional) Specifies the maximum number of errors or events to list.

-config

(Optional) Lists only configuration events.

-unfixed

(Optional) Lists only unfixed errors.

fcmap_id | *fcmap_name*

(Optional) Specifies the object ID that filters the log.

Description

This command displays a list of the errors and events in the log that are related to FlashCopy mappings. The list can be filtered further by specifying a specific object ID or name. This list displays only the errors and events that have been logged against the specified object. The list can also be filtered to show only the configuration events or the unfixed errors for the given object type or object ID. Similarly, the last *x* entries against a given object type or object ID can be listed.

Note: Although an object type of *unknown* is displayed in the error log, there is no available command to filter this object type.

An invocation example

```
svcinfo lserrlogbyfcmap -delim :
```

The resulting output

```
id:type:fixed:SNMP_trap_raised:error_type:node_name:sequence_number:  
root_sequence_number:first_timestamp:last_timestamp:number_of_errors:error_code  
0:flash:no:no:5:node1:0:0:030407085753:030407085753:1:00990185  
0:flash:no:no:5:node1:0:0:030407083355:030407083355:1:00990185  
0:flash:no:no:5:node1:0:0:030407083318:030407083318:1:00990185  
0:flash:no:no:5:node1:0:0:030407082704:030407082704:1:00990184
```

lserrlogbyhost

The **lserrlogbyhost** command displays a list of the errors and events in the log that are related to hosts.

Syntax

```
➤➤➤ svcinfo — — lserrlogbyhost — —————→  
                                     | -count — number | | -config |  
                                     | -unfixed |  
➤ —————→  
 | -nohdr | | -delim — delimiter | | host_id |  
 | | | | | host_name |
```

Parameters

-count *number*

(Optional) Specifies the maximum number of errors or events to list.

-config

(Optional) Lists only configuration events.

-unfixed

(Optional) Lists only unfixed errors.

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed even if the **-nohdr** parameter is specified.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

host_id | *host_name*

(Optional) Specifies the object ID that filters the log.

Description

This command displays a list of the errors and events in the log that are related to hosts. The list can be filtered further by specifying a specific object ID or name. This list displays only the errors and events that have been logged against the specified object. The list can also be filtered to show only the configuration events or the unfixed errors for the given object type or object ID. Similarly, the last *x* entries against a given object type or object ID can be listed.

Note: Although an object type *unknown* is displayed in the error log, there is no available command to filter this object type.

An invocation example

```
svcinfo lserrlogbyhost -delim :
```

The resulting output

```
id:type:fixed:SNMP_trap_raised:error_type:node_name:sequence_number:
root_sequence_number:first_timestamp:last_timestamp:number_of_errors:error_code
2:host:no:no:5:node1:0:0:030407082523:030407082523:1:00990106
1:host:no:no:5:node1:0:0:030407082457:030407082457:1:00990106
0:host:no:no:5:node1:0:0:030407082441:030407082441:1:00990106
```

lserrlogbyiogrp

The **lserrlogbyiogrp** command displays a list of the errors and events in the log that are related to I/O groups.

Syntax

```
➤— svcinfo — — lserrlogbyiogrp — [ -count — number ] —————➤
[ -config ] [ -unfixed ] [ -nohdr ] [ -delim — delimiter ] —————➤
[ iogroup_id ] [ iogroup_name ] —————➤
```

Parameters

-count *number*

(Optional) Specifies the maximum number of errors or events to list.

-config

(Optional) Lists only configuration events.

-unfixed

(Optional) Lists only unfixed errors.

-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.

iogroup_id | *iogroup_name*

(Optional) Specifies the object ID that filters the log.

Description

This command displays a list of the errors and events in the log that are related to I/O groups. The list can be filtered further by specifying a specific object ID or name. This list displays only the errors and events that have been logged against the specified object. The list can also be filtered to show only the configuration events or the unfixed errors for the given object type or object ID. Similarly, the last *x* entries against a given object type or object ID can be listed.

Note: Although an object type *unknown* is displayed in the error log, there is no available command to filter this object type.

An invocation example

```
svcinfolerrlogbyiogrp -delim :
```

The resulting output

```
id:type:fixed:SNMP_trap_raised:error_type:node_name:sequence_number:  
root_sequence_number:first_timestamp:last_timestamp:number_of_errors:error_code  
1:io_grp:no:no:1:node1:109:109:030407094417:030407094417:1:00000001
```

lserrlogbymdisk

The **lserrlogbymdisk** command displays a list of the errors and events in the log that are related to a specific MDisk.

Syntax

```
svcinfolerrlogbymdisk [-count number] [-config] [-unfixed] [-nohdr] [-delim delimiter] [mdisk_id] [mdisk_name]
```

Parameters

-count *number*

(Optional) Specifies the maximum number of errors or events to list.

-config

(Optional) Lists only configuration events.

-unfixed

(Optional) Lists only unfixed errors.

-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

(Optional) Specifies the object ID that filters the log.

Description

This command displays a list of the errors and events in the log that are related to a specific MDisk. The list can be filtered further by specifying a specific object ID or name. This list displays only the errors and events that have been logged against the specified object. The list can also be filtered to show only the configuration events or the unfixed errors for the given object type or object ID. Similarly, the last *x* entries against a given object type or object ID can be listed.

Note: Although an object type *unknown* is displayed in the error log, there is no available command to filter this object type.

An invocation example

```
svcinfolerrlogbydisk -delim :
```

The resulting output


```
id:type:fixed:SNMP_trap_raised:error_type:node_name:
sequence_number:root_sequence_number:first_timestamp:
last_timestamp:number_of_errors:error_code
11:mdisk:no:no:3:node1:108:108:030407092947:030407092947:1:00000016
11:mdisk:no:no:2:node1:107:107:030407092947:030407092947:1:00000016
```

Iserrlogbydiskgrp

The **Iserrlogbydiskgrp** commands display a list of the errors and events in the log that are related to MDisk groups.

Syntax

```
svcinfolerrlogbydiskgrp [-count number]
```



Parameters

-count *number*

(Optional) Specifies the maximum number of errors or events to list.

-config

(Optional) Lists only configuration events.

-unfixed

(Optional) Lists only unfixed errors.

-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_group_id | mdisk_group_name

(Optional) Specifies the object ID that filters the log.

Description

The **lserrlogbydiskgrp** command displays a list of the errors and events in the log that are related to MDisk groups. The list can be filtered further by specifying a specific object ID or name. This list displays only the errors and events that have been logged against the specified object. The list can also be filtered to show only the configuration events or the unfixed errors for the given object type or object ID. Similarly, the last *x* entries against a given object type or object ID can be listed.

Note: Although an object type *unknown* is displayed in the error log, there is no available command to filter this object type.

An invocation example

```
svcinfolerrlogbydiskgrp -delim :
```

The resulting output

```
id:type:fixed:SNMP_trap_raised:error_type:node_name:sequence_number:
root_sequence_number:first_timestamp:last_timestamp:number_of_errors:error_code
1:mdisk_grp:no:no:5:node1:0:0:030407081619:030407081619:1:00990148
128:mdisk_grp:no:no:5:node1:0:0:030407081610:030407081610:1:00990173
0:mdisk_grp:no:no:5:node1:0:0:030407081610:030407081610:1:00990148
```

Iserrlogbynode

The **Iserrlogbynode** command displays a list of the errors and events in the log that are related to nodes.

Syntax

```
➤➤➤ svcinfo -- Iserrlogbynode -- -count number ➤➤➤
-config
-unfixed -nohdr -delim delimiter ➤➤➤
node_id
node_name ➤➤➤
```

Parameters

-count *number*

(Optional) Specifies the maximum number of errors or events to list.

-config

(Optional) Lists only configuration events.

-unfixed

(Optional) Lists only unfixed errors.

-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_id | node_name

(Optional) Specifies the object ID that filters the log.

Description

This command displays a list of the errors and events in the log that are related to nodes. The list can be filtered further by specifying a specific object ID or name. This list displays only the errors and events that have been logged against the specified object. The list can also be filtered to show only the configuration events or the unfixed errors for the given object type or object ID. Similarly, the last *x* entries against a given object type or object ID can be listed.

Note: Although an object type *unknown* is displayed in the error log, there is no available command to filter this object type.

An invocation example

```
svcinfolerrlogbynode -delim :
```

The resulting output

```
id:type:fixed:SNMP_trap_raised:error_type:node_name:sequence_number:
root_sequence_number:first_timestamp:last_timestamp:number_of_errors:error_code
1:node:no:no:5:node1:0:0:030407082722:030407082722:1:00990501
1:node:no:no:5:node1:0:0:030407082716:030407082716:1:00990501
1:node:no:no:5:node1:0:0:030407052546:030407052546:1:00990383
0:node:no:no:6:node1:105:105:030407082202:030407082717:2:00980500
1:node:no:no:1:node1:102:102:030407052547:030407052547:1:00074001
```

lserrlogbyrconsistgrp

You can use the **lserrlogbyrconsistgrp** command to display the error log by Metro or Global Mirror consistency groups.

Syntax

```
svcinfolerrlogbyrconsistgrp [-count number]
                             [-config] [-unfixed] [-nohdr] [-delim delimiter]
                             [rconsistgrp_id] [rconsistgrp_name]
```

Parameters

-count *number*

(Optional) Specifies the maximum number of errors or events to list.

-config

(Optional) Lists only configuration events.

-unfixed

(Optional) Lists only unfixed errors.

-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.

rcconsistgrp_id | rcconsistgrp_name

(Optional) Specifies the object ID that filters the log.

Description

This command displays a list of the errors and events in the log by Metro or Global Mirror consistency groups. The list can be filtered further by specifying a specific object ID or name. This displays only the errors and events that have been logged against the specified object. The list can also be filtered to show only the configuration events or the unfixed errors for the given object type or object ID. Similarly, the last *x* entries against a given object type or object ID can be listed.

Note: Although an object type *unknown* is displayed in the error log, there is no available command for this object type.

An invocation example

```
svcinfolerrlogbyrcconsistgrp -delim :
```


The resulting output

```
id:type:fixed:SNMP_trap_raised:error_type:node_name:sequence_number:
root_sequence_number:first_timestamp:last_timestamp:number_of_errors:error_code
253:rc_const_grp:no:no:5:node1:0:0:030407090333:030407090333:1:00990240
254:rc_const_grp:no:no:5:node1:0:0:030407090327:030407090327:1:00990240
255:rc_const_grp:no:no:5:node1:0:0:030407090323:030407090323:1:00990240
```

Iserrlogbyrcrelationship

The **Iserrlogbyrcrelationship** command displays a list of the errors and events in the log by Metro Mirror or Global Mirror relationships.

Syntax



Parameters

-count *number*

(Optional) Specifies the maximum number of errors or events to list.

-config

(Optional) Lists only configuration events.

-unfixed

(Optional) Lists only unfixed errors.

-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.

rrelationship_id | rrelationship_name

(Optional) Specifies the object ID that filters the log.

Description

This command displays a list of the errors and events in the log by Metro or Global Mirror relationships. The list can be filtered further by specifying a specific object ID or name. This displays only the errors and events that have been logged against the specified object. The list can also be filtered to show only the configuration events or the unfixed errors for the given object type or object ID. Similarly the last *x* entries against a given object type or object ID can be listed.

Note: Although an object type *unknown* is displayed in the error log, there is no available command to filter this object type.

An invocation example

```
svcinfo lserrlogbyrrelationship -delim :
```

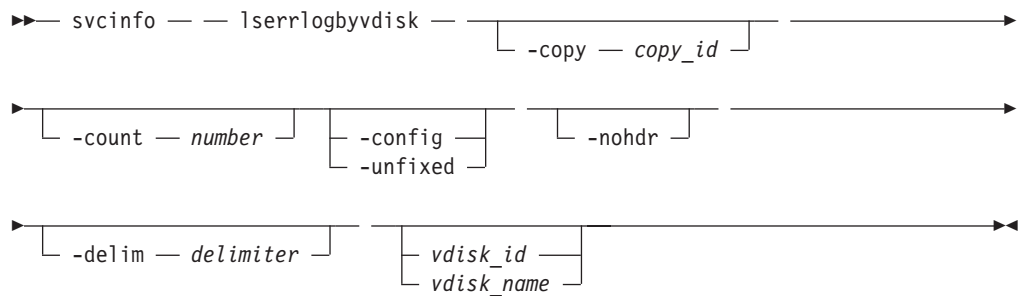
The resulting output

```
id:type:fixed:SNMP_trap_raised:error_type:node_name:sequence_number:
root_sequence_number:first_timestamp:last_timestamp:number_of_errors:error_code
2:remote:no:no:5:node1:0:0:030407090442:030407090442:1:00990226
2:remote:no:no:5:node1:0:0:030407090106:030407090106:1:00990225
1:remote:no:no:5:node1:0:0:030407085932:030407085932:1:00990225
2:remote:no:no:6:n/a:106:106:030407090117:030407090117:1:00985002
```

lserrlogbyvdisk

The **lserrlogbyvdisk** command displays a list of the errors and events in the log by VDIs.

Syntax



Parameters

-copy *copy_id*

(Optional) Displays a list of errors and events for the specified VDisk copy.

-count *number*

(Optional) Specifies the maximum number of errors or events to list.

-config

(Optional) Lists only configuration events.

-unfixed

(Optional) Lists only unfixed errors.

-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*

(Optional) Specifies the object ID that filters the log.

Description

This command displays a list of the errors and events in the log by VDIs. The list can be filtered further by specifying a specific object ID or name. This list displays only the errors and events that have been logged against the specified object. The list can also be filtered to show only the configuration events or the unfixed errors for the given object type or object ID. Similarly, the last *x* entries against a given object type or object ID can be listed.

Note: Although an object type *unknown* is displayed in the error log, there is no available command to filter this object type.

An invocation example

```
svcinfolerrlogbyvdisk -delim :
```

The resulting output

```
id:type:fixed:SNMP_trap_raised:error_type:node_name:type:sequence_number:
root_sequence_number:first_timestamp:last_timestamp:number_of_errors:error
code:copy_id
4:vdisk:no:no:0::vdisk:0:0:021009082703:-300101000000:0:990169:0
1:vdisk:no:no:0::vdisk:0:0:021009081951:-300101000000:0:990182:0
2:vdisk:no:no:0::vdisk:0:0:021009081915:-300101000000:0:990182:0
0:vdisk:no:no:0::vdisk:0:0:021009081835:-300101000000:0:990182:1
```

lserrlogdumps

The **lserrlogdumps** command displays a list of error log dumps in the `/dumps/elogs` directory. These dumps are created as a result of the **svctask dumperlog** command.

Syntax

```
svcinfolerrlogdumps [-nohdr] [-delim delimiter] [node_id | node_name]
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim delimiter

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. 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.

node_id | node_name

(Optional) Specifies the node ID or name to list the available dumps of the given type. If you do not specify a node, the dumps that are available on the configuration node are listed.

Description

This command displays a list of error log dumps. These dumps are created as a result of the **svctask dumperlog** command. An error log dump describes the contents of the error log at the time that the command was issued. If you do not

specify a node, the dumps that are available on the configuration node are listed. The command displays files from the `/dumps/elog`s directory.

An invocation example

```
svcinfo lserrlogdumps
```

The resulting output

id	filename
0	errlog_lynn02_030327_154511
1	aaa.txt_lynn02_030327_154527
2	aaa.txt_lynn02_030327_154559
3	errlog_lynn02_030403_110628

lsfabric

The `lsfabric` command generates a report that displays the connectivity between nodes and other controllers and hosts.

Syntax

```
svcinfo -- lsfabric --
```

-node	node_id_or_name	-port	port_id
-wwpn	wwpn		
-host	host_id_or_name		
-controller	controller_id_or_name		
-cluster	cluster_id_or_name		

Parameters

-node *node_id_or_name*

(Optional) Displays the output for all ports for the specified node. The only parameter that you can specify with the **-node** parameter is the **-port** parameter.

-port *port_id*

(Optional) Displays a concise view of all WWPNS that are logged into the specified port ID and node. The **-port** parameter must be specified with only the **-node** parameter. A valid *port_id* value is a number from 1 - 4 that specifies the port number in the vital product data (VPD) or the hexadecimal WWPNS of the local port.

-wwpn *wwpn*

(Optional) Displays a list of all ports that have a login to the specified WWPNS. You cannot use the **-wwpn** parameter with any other parameter.

-host *host_id_or_name*

(Optional) Specifies a host name or ID. Issuing the `lsfabric` command with the **-host** parameter is equivalent to issuing the `svcinfo lsfabric -wwpn wwpn` command for every configured WWPNS of the specified host. For example, a host with two ports that are zoned to one port of every node in an eight-node cluster produces 16 lines of output. You cannot use the **-host** parameter with any other parameter.

-controller *controller_id_or_name*

(Optional) Specifies a controller ID or name. You cannot use the **-controller**

parameter with any other parameter in this command. Issuing the **lsfabric** command with the **-controller** parameter is equivalent to issuing the **svcinfo lsfabric -wwpn** *wwpn* command for every configured WWPN of the specified controller. For example, a controller with 4 ports connected to a 8 node cluster with 2 counter part SANs produces 64 lines of output.

-cluster *cluster_id_or_name*

(Optional) Specifies a cluster ID or name. You cannot use the **-cluster** parameter with any other parameter. Issuing the **lsfabric** command with the **-cluster** parameter is equivalent to issuing the **svcinfo lsfabric -wwpn** *wwpn* command for every known WWPN in the specified cluster. Output is sorted by remote WWPNs and then cluster WWPNs. This parameter can be used to check the state of connections within the local cluster or between the local and remote cluster. When the local cluster ID or name is specified, each node-to-node connection is listed twice: once from each end. For example, an eight-node cluster with two counterpart SANs produces eight nodes, multiplied by seven other nodes, multiplied by two SANs, multiplied by four point-to-point logins, equals 448 lines of output.

Description

The **lsfabric** command can be issued with any of the parameters to display a limited subset of information. If the command is issued without any parameters, it provides output for every node.

Values for the **Type** and **State** columns are:

state active

The meaning of this value depends on the object that it applies to, as follows:

- **host or controller:** SCSI commands were issued within the last 5 minutes.
- **node:** node ports can see other ports.

state inactive

No transactions have completed within the last 5 minutes.

Note: It can take up to 10 seconds after a command for a controller port to change from inactive to active. It can take up to 5 minutes after a command for a host port to change from inactive to active.

type One of the following values is displayed:

- **host**
- **node**
- **controller**
- **unknown**

You can issue this command to view all the information about the connections that are available to your cluster.

An invocation example

```
svcinfo lsfabric -delim :
```

The resulting output Each row of output contain the following colon-separated columns:

remote_wwpn:remote_nportid: id:node_name:local_wwpn:
local_port:local_nportid:state:name:cluster_name:type

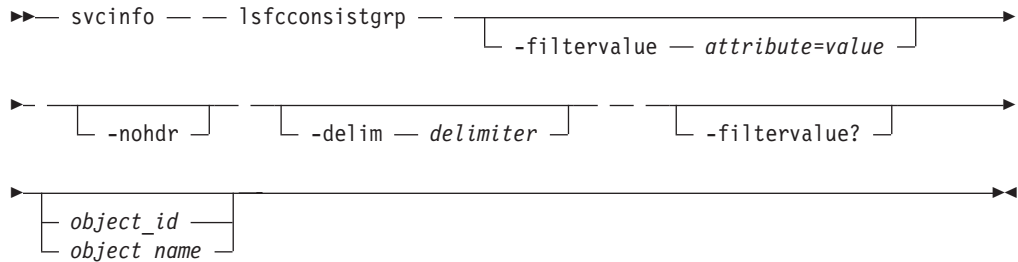
lsfcconsistgrp

The **lsfcconsistgrp** command returns a concise list or a detailed view of FlashCopy consistency groups that are visible to the cluster. This information is useful for tracking FlashCopy consistency groups.

The list report style can be used to obtain two styles of report:

- A list containing concise information about all of the FlashCopy consistency groups on a cluster. (Each entry in the list corresponds to a single FlashCopy consistency group.)
- The detailed information about a single FlashCopy consistency group.

Syntax



Parameters

-filtervalue *attribute=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are returned.

Note: Some filters allow the use of a wildcard when you enter the command. The following rules apply to the use of wildcards with the SAN Volume Controller CLI:

- The wildcard character is an asterisk character (*).
- The command can contain a maximum of one wildcard, which must be the first or last character in the string.
- When you use a wildcard, surround the filter entry with double quotation marks (""), as follows:

```
svcinfo 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 **svcinfc lsfccconsistgrp** command are:

- **name**
- **id**
- **status**
- **FC_group_id**

Description

This command returns a concise list or a detailed view of FlashCopy consistency groups that are visible to the cluster.

The following list provides values of the *status* attribute that are displayed as data in the output views:

status **empty, idle_or_copied, preparing, prepared, copying, stopped, suspended, stopping**

A concise invocation example

```
svcinfc lsfccconsistgrp -delim :
```

The concise resulting output

```
id:name:status
1:ffccg0:empty
2:ffccg1:idle_or_copied
3:ffccg2:idle_or_copied
```

A detailed invocation example

```
svcinfc lsfccconsistgrp -delim : 1
```

The detailed resulting output

```
id:1
name:ffccg0
status:empty
```

A detailed invocation example

```
svcinfc lsfccconsistgrp -delim : fccstgrp0
```

The detailed resulting output

```
id:1
name:FCcgrp0
status:idle_or_copied
autodelete:off
```

```

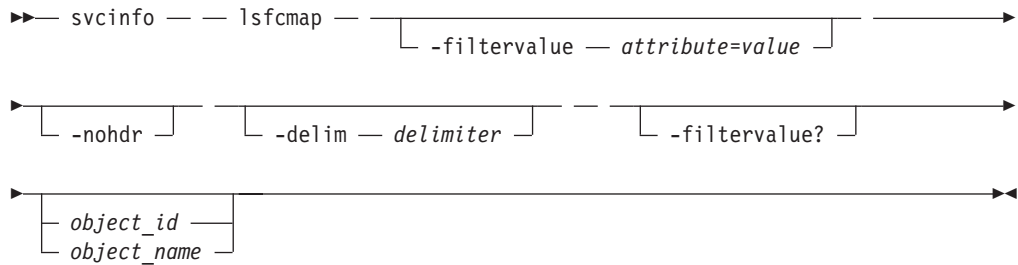
FC_mapping_id:0
FC_mapping_name:fcmap0
FC_mapping_id:1
FC_mapping_name:fcmap1

```

lsfcmmap

The `lsfcmmap` command generates a list containing concise information about all of the FlashCopy mappings that are visible to the cluster, or detailed information for a single FlashCopy mapping.

Syntax



Parameters

-filtervalue *attribute=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are displayed.

Note: Some filters allow the use of a wildcard when you enter the command. The following rules apply to the use of wildcards with the SAN Volume Controller CLI:

- The wildcard character is the asterisk (*), which must be used as the first or last character in the string.
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""), as follows:

```
svcinfo 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 cluster.

The following list shows attribute values that can be displayed as output view data:

status idle_or_copied, preparing, prepared, copying, stopped, suspended or stopping

start_time

Displays the time that the copy was last started. It is in the format `YYMMDDHHMMSS`. If a copy has not been started, a blank line is displayed.

A concise invocation example

```
svcinflsfcmap -delim :
```

The concise resulting output

```
id:name:source_vdisk_id:source_vdisk_name:target_vdisk_id:target_vdisk_name:
group_id:group_name:status:progress:copy_rate:clean_progress:incremental:
partner_FC_id:partner_FC_name:restoring
0:fcmap0:0:vdisk0:1:vdisk1:0:fccstgrp0:idle_or_copied:0:50:0:on:2:fcmap2:no
1:fcmap1:2:vdisk2:3:vdisk3:0:fccstgrp0:idle_or_copied:0:0:100:off:::no
2:fcmap2:1:vdisk1:0:vdisk0:0:fccstgrp1:idle_or_copied:0:0:100:off:0:fcmap0:no
```

A detailed invocation example

```
svcinfc lsfcmap -delim : 0
```

The detailed resulting output

```
| id:0  
| name:fcmap0  
| source_vdisk_id:63  
| source_vdisk_name:vdisk63  
| target_vdisk_id:57  
| target_vdisk_name:vdisk57  
| group_id:  
| group_name:  
| status:idle_or_copied  
| progress:0  
| copy_rate:0  
| start_time:  
| dependent_mappings:0  
| autodelete:off  
| clean_progress:100  
| clean_rate:50  
| incremental:off  
| difference:100  
| grain_size:256  
| IO_group_id:1  
| IO_group_name:io_grp1  
| partner_FC_id:  
| partner_FC_name:  
| restoring:no
```

lsfcmapcandidate

The **lsfcmapcandidate** command lists all of the VDisks that are associated with fewer than 256 FlashCopy mappings.

Syntax

```
▶▶ svcinfo — lsfcmapcandidate — [ -nohdr ] —————▶  
▶ [ -delim — delimiter ] —————▶▶
```

Parameters

-nohdr

(Optional) By default, the heading is displayed for the column of data in a concise style view, and for the item of data in a detailed style view. The **-nohdr** parameter suppresses the display of the heading.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, a colon character (:) separates all items of data in a concise

view; the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

Description

This command returns a list of VDisks that are associated with fewer than 256 FlashCopy mappings.

An invocation example

```
svcinflsfmapcandidate
```

The resulting output

```
id  
2  
3  
4
```

lsfmapprogress

The **lsfmapprogress** command returns the progress of the background copy of a FlashCopy mapping. This is displayed as a percentage completed value.

Syntax

```
svcinflsfmapprogress [-nohdr] [-delim delimiter] fcmapi fcmaname
```

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 (:).

fcmapi | fcmaname

(Required) Specifies that you want the report to display the progress of the background copy for the designated FlashCopy mapping.

Description

This command reports a percentage for the progress of the background copy being done on the specified FlashCopy mapping.

An invocation example

```
svcinfolsfcmapprogress 0
```

The resulting output

```
id          progress
0           0
```

lsfcmdependentmaps

The **lsfcmdependentmaps** command displays all the FlashCopy mappings that are dependent on the user specified mapping.

Syntax

```
svcinfolsfcmdependentmaps [-nohdr]
                             [-delim delimiter] [fc_id | fc_name]
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim delimiter

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

fc_id | fc_name

(Required) Specifies the name or ID of the FlashCopy mapping to list the dependent maps for.

Description

This command returns a list of dependent FlashCopy mappings. This command can be used to determine the list of FlashCopy mappings that would also stop if you stopped a mapping using the **-force** parameter.

There is a `dependent_mapping_count` field in the FlashCopy map detailed view (displayed when you process the **lsfcm** command) that you can use as an indicator as to whether there are any dependent mappings in progress. If the count is zero, there are no dependent copies.

Note: If a period of time elapses between the time you process the **lsfcmap** command and the **lsfcmapdependentmaps** command, there could be a difference between the actual number of dependent mappings being processed and the number that was reported by the **lsfcmap** command.

A concise invocation example

```
svcinfolsfmapdependentmaps -delim : 2
```

The resulting output

```
fc_id:fc_name  
1:fcmap1  
3:fcmap3
```

lsfeaturedumps

The **lsfeaturedumps** command displays a list of dump files in the **/dumps/feature** directory. These dump files are created as a result of the **svctask dumpinternallog** command.

Syntax

```
svcinfolsffeaturedumps [-nohdr] [-delim delimiter] [node_id | node_name]
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim delimiter

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

node_id | node_name

(Optional) Specifies the node ID or node name to list the available dumps of the given type. If you do not specify a node, the dumps that are available on the configuration node are listed.

Description

This command displays a list of featurization dumps. These dumps are created as a result of issuing the **svctask dumpinternallog** command. A featurization dump

file describes the contents of the featurization log at the time that the command was issued. If you do not specify a node, the dumps that are available on the configuration node are listed. The command displays files from the `/dumps/feature` directory.

Issue the `svcinfolfeaturedumps` command to display 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.

An invocation example

```
svcinfolfeaturedumps
```

The resulting output

```
id          feature_filename
0           feature.txt
```

lsfreeextents

The `lsfreeextents` command lists the number of free extents that are available on a specified MDisk.

Syntax

```
svcinfolfeatureextents [-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

```
svcinfo lsfreeextents 2
```

The resulting output

```
id 2
number_of_extents 4372
```

lshbaportcandidate

The **lshbaportcandidate** command lists all of the unconfigured, logged-in host bus adapter (HBA) ports. This information is used to find open HBA ports.

Syntax

```
▶▶▶ svcinfo — — lshbaportcandidate — — [ -nohdr ] —————▶▶▶
▶ [ -delim — delimiter ] —————▶▶▶
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

Description

This command returns a list of unconfigured, logged in HBA ports.

Note: The **svcinfo lshbaportcandidate** command presents a list of host HBA ports that are logged in to nodes. However, there are situations when the information that is presented might include host HBA ports that are no longer logged in or even part of the SAN fabric. For example, a host HBA port is unplugged from a switch but **svcinfo lshbaportcandidate** still shows the WWPN that is logged in to all nodes. If this occurs, the incorrect entry is

removed when another device is plugged in to the same switch port that previously contained the removed host HBA port.

An invocation example

```
svcinfo lshbaportcandidate
```

The resulting output

```
id  
210100E08B2520D4
```

lshost

The `lshost` command generates a list with concise information about all the hosts visible to the cluster and detailed information about a single host.

Syntax

```
svcinfo -- lshost [ -filtervalue -- attrib=value ]  
[ -nohdr ] [ -delim delimiter ] [ -filtervalue? ]  
[ object_id ]  
[ object_name ]
```

Parameters

-filtervalue *attrib=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are returned. If a capacity is specified, the units must also be included.

Note: Some filters allow the use of a wildcard when you enter the command. The following rules apply to the use of wildcards with the SAN Volume Controller CLI:

- The wildcard character is an asterisk (*).
- The command can contain a maximum of one wildcard.
- When using a wildcard character, you must enclose the filter entry within double quotation marks (" "), as follows:

```
svcinfo lshost -filtervalue "name=md*"
```

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the

-delim parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

object_id | object_name

(Optional) Specifies the name or ID of an object. When you use this parameter, the detailed view of the specific object is returned and any value that is specified by the **-filtervalue** parameter is ignored. If you do not specify the *object_id | object_name* parameter, the concise view of all objects matching the filtering requirements that is specified by the **-filtervalue** parameter are displayed.

-filtervalue?

(Optional) Specifies that you want your report to display any or all of the list of valid filter attributes. The valid filter attributes for the `svcinfo lshost` command are:

- `host_name`
- `host_id`
- `port_count`
- `name`
- `id`
- `iogrp_count`

Description

This command returns a concise list or a detailed view, of hosts visible to the cluster.

The following list provides the different states for a host:

offline

The host is offline if one or more I/O groups with VDisk mappings does not have a login for the specified WWPN.

degraded

The host is degraded if one or more nodes with VDisk mappings do not have a login for the specified WWPN.

inactive

The host is inactive if all the nodes with VDisk mappings have a login for the specified WWPN, however, no nodes have seen any SCSI commands from the WWPN in the last 5 minutes.

active The host is active if all the nodes with VDisk mappings have a login for the specified WWPN, however, at least one node has seen SCSI commands from the WWPN in the last 5 minutes.

If a host does not have any VDisk mappings, then it is either reported as offline or inactive.

Note: The `svcinfo lshost` command presents a list of host HBA ports that are logged in to nodes. However, there are situations when the information presented can include host HBA ports which are no longer logged in or even part of the SAN fabric. For example: A host HBA port is unplugged from a switch but `scvinfo lshost` still shows the WWPN logged in to all

nodes. If this occurs, the incorrect entry is removed when another device is plugged in to the same switch port that previously contained the removed host HBA port.

A concise invocation example

```
svcinfo lshost
```

The concise resulting output

id	name	port_count	iogrp_count
0	mchost20	1	4
1	mchost30	1	4
2	mchost200	1	4
3	mchost40	1	4
4	mchost240	1	4
5	mchost170	1	4
6	mchost120	1	4
7	mchost60	1	4
8	mchost180	1	4
9	mchost13	2	4

A detailed invocation example

```
svcinfo lshost mchost13
```

The detailed resulting output

```
id 9
name mchost13
port_count 2
type generic
mask 1111
iogrp_count 4
iscsiname mc13host
node_logged_in_count 0
state offline
iscsiname mchost13
node_logged_in_count 0
state offline
```

lshostiogrp

The **lshostiogrp** command displays a list of all the I/O groups that are associated with a specified host.

Syntax

```
svcinfo -lshostiogrp [-nohdr] [-delim delimiter]
[host_id]
[host_name]
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

host_id | *host_name*

(Required) The name or ID of the host for which the list of I/O groups is required.

Description

This command displays a list of all the I/O groups that are mapped to the specified host.

An invocation example

```
svcinfolshostiogrp -delim : hostone
```


The resulting output

```
id:name
0:io_grp0
1:io_grp1
```

lshostvdiskmap

The `lshostvdiskmap` command displays a list of VDisks that are mapped to a given host. These are the VDisks that are recognized by the specified host.

Syntax



Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the

-delim parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

host_id | *host_name*

(Optional) Specifies the host in terms of its ID or name. The command displays a list of all the virtual disks that are mapped to the specified host and the SCSI ID by which they are mapped. If neither a host ID or name are entered, the command displays a list of all recognized host-to-VDisk mappings.

Description

This command displays a list of VDisk IDs and names. These are the VDIsks that have been mapped to the specified host; that is, they are visible to the specified host. The SCSI LUN ID is also displayed. This SCSI LUN ID is the ID by which the VDisk is recognized by the host.

Each VDisk that is exported by the cluster is assigned a unique vpath number. This number identifies the VDisk and determines which VDisk corresponds to the volume that the hosts recognize. This procedure can only be performed using the command-line interface.

For the specified volume, find the vpath serial number by issuing the datapath query device command. Find the host that is defined to the cluster that corresponds with the host that you are working with.

1. The WWPNs are an attribute of the HBA. You can find these by looking at the device definitions stored by your operating system. For example, on AIX they are in the ODM, in Windows they are in the Device Manager details for the given HBA.
2. Verify which host is defined to the cluster that these ports belong to. The ports are stored as part of the detailed view, so you must list each host in turn by issuing the following command:

```
svcinfolshost host_name | host_id
```

where *host_name* | *host_id* is the name or ID of the host. Check for matching WWPNs.

Note: Name your hosts accordingly. For example, if the actual host is called **orange**, also name the host that is defined to the cluster **orange**.

When you have the *hostname* defined to the cluster and the *vpath serial number*, issue the following command:

```
svcinfolshostvdiskmap hostname
```

where *hostname* is the name of the host. A list is displayed. Look for the VDisk UID that matches the *vpath serial number* and record the VDisk name or ID.

An invocation example

```
svcinfolshostvdiskmap -delim : 2
```

The resulting output

```

| id:name:SCSI_id:vdisk_id:vdisk_name:vdisk_UID
| 2:host2:0:10:vdisk10:6005076801958001500000000000000A
| 2:host2:1:11:vdisk11:6005076801958001500000000000000B
| 2:host2:2:12:vdisk12:6005076801958001500000000000000C
| 2:host2:3:13:vdisk13:6005076801958001500000000000000D
| 2:host2:4:14:vdisk14:6005076801958001500000000000000E

```

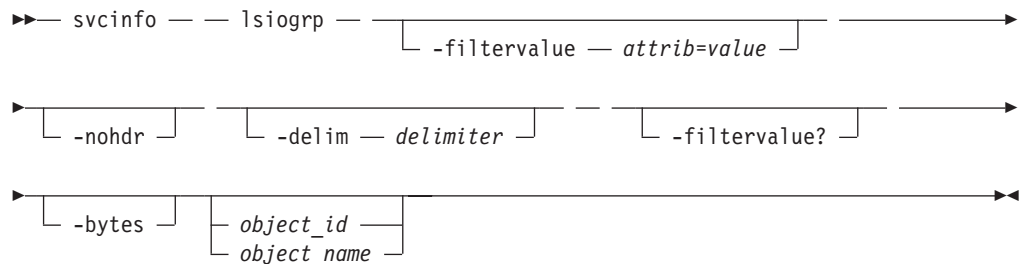
Isiogrp

The **Isiogrp** command returns a concise list or a detailed view of I/O groups visible to the cluster.

The list report style can be used to obtain the following two styles of report:

- A list containing concise information about all the I/O groups that are visible to the cluster. Each entry in the list corresponds to a single I/O group.
- The detailed information about a single I/O group.

Syntax



Parameters

-filtervalue *attrib=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are returned. If a capacity is specified, the units must also be included.

Note: Some filters allow the use of a wildcard when you enter the command. The following rules apply to the use of wildcard characters with the SAN Volume Controller CLI:

- The wildcard character is an asterisk (*), which must be the first or last character in the string.
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""), as follows:

```
svcinfo 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 **svcinflsiogrp** command:

- **HWS_name**
- **HWS_unique_id**
- **node_count**
- **vdisk_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 cluster.

A concise invocation example

```
svcinflsiogrp -delim :
```

The concise resulting output

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

A detailed invocation example

```
svcinflsiogrp -delim : 0
```

The detailed resulting output

```
id:0
name:io_grp0
node_count:1
vdisk_count:51
host_count:0
flash_copy_total_memory:3.0MB
flash_copy_free_memory:1.0MB
remote_copy_total_memory:6.5MB
remote_copy_free_memory:2.8MB
mirroring_total_memory:1.0MB
mirroring_free_memory:0.3MB
```

Isiogrphost

The **Isiogrphost** command displays a list of the hosts that are mapped to a specified I/O group.

Syntax

```
svcinfo -- Isiogrphost [-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 **Isiogrphost** command displays a list of hosts that are mapped to a specified I/O group.

An invocation example

```
svcinfo Isiogrphost -delim : 0
```

The resulting output

```
id:name
0:hostzero
1:hostone
```

Isiogrpcandidate

Use the **Isiogrpcandidate** command to list the I/O groups that can have nodes added to them.

Syntax

```
svcinfo -- lsiogrpcandidate -- [-nohdr]
[-delim delimiter]
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

Description

This command displays a list of I/O groups to which nodes can be added. Only the I/O group IDs are displayed.

An invocation example

```
svcinfo lsiogrpcandidate
```

The resulting output


```
id
0
1
2
3
4
```

lsiostatsdumps

The `lsiostatsdumps` command displays a list of dumps in the `/dumps/iostats` directory. Use the `svctask startstats` command to create these dumps.

Syntax

```
svcinfo -- lsiostatsdumps -- [-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.

node_id | *node_name*

(Optional) Specifies the node ID or name to list the available dumps of the given type. If you do not specify a node, the dumps that are available on the configuration node are listed.

Description

This command displays a list of I/O statistics dumps. These dumps are created when you issue the `svctask startstats` command. If you do not specify a node, the dumps that are available on the configuration node are listed. The command displays files from the `/dumps/iostats` directory.

An invocation example

```
svcinfo lsiostatsdumps
```


The resulting output


```
|
|      id          iostat_filename
|      0          Nv_stats_mala75_031123_072426
|      1          Nm_stats_mala75_031123_072425
|      2          Nn_stats_mala75_031123_072424
```

Isiotracedumps

Use the **Isiotracedumps** command to display a list of files in the `/dumps/iotrace` directory.

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_id | *node_name*

(Optional) Specifies the node ID or name to list the available dumps of the given type. If you do not specify a node, the dumps that are available on the configuration node are listed.

Description

This command displays a list of I/O trace dumps. These dumps are created when you issue the `svctask settrace` command. If you do not specify a node, the dumps that are available on the configuration node are listed. The command displays files from the `/dumps/iotrace` directory.

An invocation example

```
svcinfolsiotracedumps
```

The resulting output


```


id          iotrace_filename
0           c1_mala75_030405_092155
1           c2_mala75_030405_092156
2           c3_mala75_030405_092158
3           c4_mala75_030405_092159
4           c5_mala75_030405_092201
  
```

Isiscsiauth

The `lsiscsiauth` command lists the Challenge Handshake Authentication Protocol (CHAP) secret configured for authenticating an entity to the SAN Volume Controller cluster.

Syntax





Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

-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:

```
svcinfolsiscsiauth -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 cluster. The command also displays the configured iSCSI authentication method. The `iscsi_auth_method` field can have values of `none` or `chap`.

When you create an iSCSI host using the `svctask mkhost` command with the `iscsiname` parameter, the host is initially configured with the authentication method as `none`, and no CHAP secret is set. To set a CHAP secret for authenticating the iSCSI host with the SAN Volume Controller cluster, use the `svctask chhost` command with the `chapsecret` parameter.

A invocation example

```
svcinfolsiscsiauth
```

The resulting output

type	id	name	iscsi_auth_method	iscsi_chap_secret
host	0	mchost20	none	
host	1	mchost30	none	
host	2	mchost200	none	
host	3	mchost40	none	
host	4	mchost240	none	
host	5	mchost170	none	
host	6	mchost120	none	
host	7	mchost60	none	
host	8	mchost180	none	
host	9	mchost13	none	
host	10	newhost	none	

Islicense

The `Islicense` command displays current license settings for cluster features.

Syntax

```
svcinfollicense [-nohdr] [-delim delimiter]
```

Parameters

-nohdr

(Optional) Suppresses the display of these headings. By default, headings are displayed for each column of data (in a concise style view providing general information about objects of a particular type) and for each item of data (in a detailed style view providing much more information about a specific object of a particular type).

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The `-delim` parameter overrides this behavior. Valid input for the `-delim` parameter is a one-byte character. If you enter `-delim :`, a colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

Description

The **lslicense** command displays license settings for cluster features, including FlashCopy, RemoteCopy, and Virtualization settings. The displayed output lists feature enablement and capacities.

Use the **chlicense** command to change the feature license settings. Because the feature license settings are entered when the cluster is first created, you must only update the settings if you have changed your license.

An invocation example

```
svcinfolicense
```

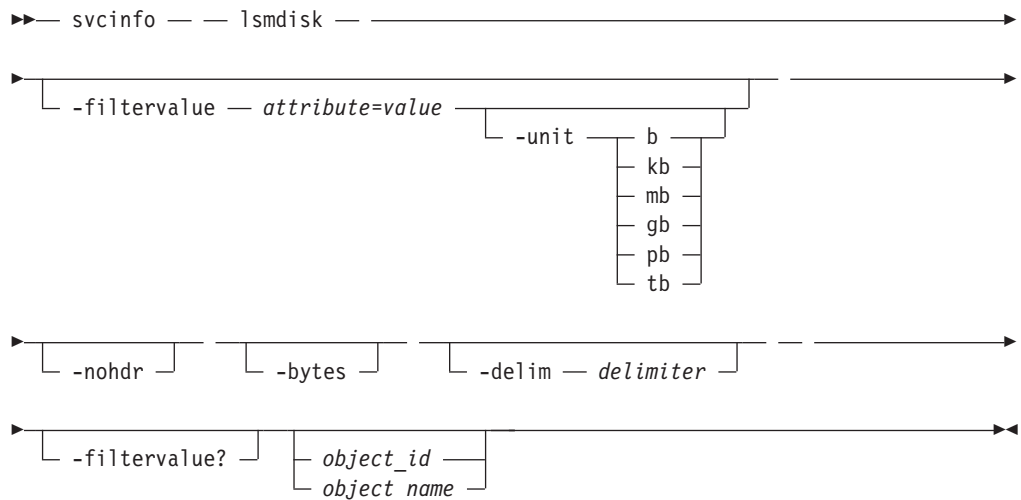
The resulting output

```
used_flash 4.73
used_remote 0
used_virtualization 21.12
license_flash 5
license_remote 0
license_virtualization 32
license_physical_disks 0
license_physical_flash off
license_physical_remote off
```

lsmdisk

The **lsmdisk** command returns a concise list or a detailed view of managed disks (MDisks) visible to the cluster. It can also list detailed information about a single MDisk.

Syntax



Parameters

-filtervalue attribute=value

(Optional) Specifies a list of one or more filter attributes matching the specified values; see **-filtervalue?** for the supported attributes. Only objects with a value that matches the filter attribute value are returned. If **capacity** is specified, the units must also be included.

Note: Some filters allow the use of a wildcard when entering the command. The following rules apply to the use of wildcards with the SAN Volume Controller CLI:

- The wildcard character is an asterisk (*).
- The command can contain a maximum of one wildcard, which must be the first or last character in the string.
- When using a wildcard character, you must enclose the filter entry within double quotation marks (""), as follows:

```
svcinfolsmdisk -filtervalue "name=md*"
```

-filtervalue?

(Optional) Includes all of the valid filter attributes in the report. The following filter attributes are valid for the `svcinfolsmdisk` command:

- id
- name
- status
- mode
- mdisk_grp_id
- mdisk_grp_name
- capacity
- quorum_index
- block_size
- controller_name
- ctrl_type
- ctrl_WWNN
- controller_id
- path_count
- max_path_count
- ctrl_LUN_#
- UID
- preferred_WWPN
- active_WWPN
- node_id
- node_name
- location

Any parameters specified with the **-filtervalue?** parameter are ignored.

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-bytes

(Optional) Specifies that you want the report to display all capacities as bytes. Capacity values displayed in units other than bytes might be rounded. When filtering on capacity, use a unit of bytes, **-unit b**, for exact filtering.

-delim delimiter

(Optional) By default in a concise view, all columns of data are space

separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

object_id | object_name

(Optional) Specifies the name or ID of an object. When you use this parameter, the detailed view of the specific object is returned and any value that is specified by the **-filtervalue** parameter is ignored. If you do not specify the *object_id | object_name* parameter, the concise view displays all objects matching the filtering requirements that are specified by the **-filtervalue** parameter.

Description

This command returns a concise list or a detailed view of MDisks visible to the cluster. Table 15 provides the potential output for MDisks.

Table 15. MDisk output

Attribute	Values
status	online, offline, excluded, degraded_paths, degraded_ports
mode	unmanaged, managed, image
quorum_index	0, 1, 2, or blank if the MDisk is not being used as a quorum disk
block_size	512, 524 bytes in each block of storage
ctrl_type	4, 6, where 6 is a solid-state drive (SSD) attached inside a node and 4 is any other device

Note: The automatic discovery performed by the cluster does not write anything to an unmanaged MDisk. It is only when you add an MDisk to an MDisk group, or use an MDisk to create an image mode virtual disk, that the system uses the storage.

To see which MDisks are available, issue the `svctask detectmdisk` command to manually rescan the fibre-channel network for any new MDisks. Issue the `svcinfo lsmdiskcandidate` command to show the unmanaged MDisks. These MDisks have not been assigned to an MDisk group.

Notes:

1. The MDisk *path_count* value is the number of nodes through which the MDisk is currently accessible.
2. The MDisk *max_path_count* value is the highest value that the *path-count* has reached since the MDisk was initially discovered or since the MDisk was last accessible through all nodes in the cluster.
3. If an SSD is attached inside a node:
 - The *node_id* value is the ID of the node that the MDisk is in.
 - The *node_name* value is the name of the node that the MDisk is in.
 - The *location* value is the slot number printed on the drive bay.

The following define the status fields:

| **Online**
| The MDisk is online and available.
| **Degraded ports**
| There are one or more MDisk port errors.
| **Degraded paths**
| One or more paths to the MDisk have been lost; the MDisk is not online to
| every node in the cluster.
| **Offline**
| All paths to the MDisk are lost.
| **Excluded**
| The MDisk is excluded from use by the cluster; the MDisk port error count
| exceeded the threshold.

A concise invocation example

```
svcinfo lsmdisk -delim :
```

The concise resulting output

```
id:name:status:mode:mdisk_grp_id:mdisk_grp_name:capacity:  
ctrl_LUN_#:controller_name:UID  
0:mdisk0:online:managed:0:mdiskgrp0:68.4GB:0000000000000000:controller0:  
20000004cf2422aa00000000000000000000000000000000000000000000000000000  
1:mdisk1:online:managed:0:mdiskgrp0:68.4GB:0000000000000000:controller1:  
20000004cf1fd19d0000000000000000000000000000000000000000000000000000  
2:mdisk2:online:managed:0:mdiskgrp0:68.4GB:0000000000000000:controller2:  
20000004cf2425310000000000000000000000000000000000000000000000000000  
3:mdisk3:online:managed:0:mdiskgrp0:68.4GB:0000000000000000:controller3:  
20000004cf1fd7a00000000000000000000000000000000000000000000000000000  
4:mdisk4:online:unmanaged:::68.4GB:0000000000000000:controller4:  
20000004cf2413080000000000000000000000000000000000000000000000000000  
5:mdisk5:online:unmanaged:::68.4GB:0000000000000000:controller5:  
20000004cf2412ca0000000000000000000000000000000000000000000000000000
```

A detailed invocation example

```
svcinfo lsmdisk -delim : 2
```

The detailed resulting output

```
id:2  
name:mdisk2  
status:online  
mode:unmanaged  
mdisk_grp_id:  
mdisk_grp_name:  
capacity:68.4GB  
quorum_index:  
block_size:512  
controller_name:controller2  
ctrl_type:4  
ctrl_WWNN:20000004CF242531  
controller_id:2  
path_count:1  
max_path_count:1  
ctrl_LUN_#:0000000000000000  
UID:20000004cf2425310000000000000000000000000000000000000000000000000000  
preferred_WWPN:22000004CF242531  
active_WWPN:22000004CF242531  
node_id  
node_name  
location
```

lsmdiskdumps

The `lsmdiskdumps` command displays a list of dumps in the `/dumps/mdisk` directory.

Syntax

```
svcinfo lsmdiskdumps [-nohdr] [-delim delimiter] [node_id | node_name]
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

node_id | *node_name*

(Optional) Specifies the node ID or name to list the available dumps of the given type. If you do not specify a node, the dumps that are available on the configuration node are listed.

Description

This command displays a list of managed disk (MDisk) error data. If you do not specify a node, the dumps that are available on the configuration node are listed. The command displays files from the `/dumps/mdisk` directory.

An invocation example

```
svcinfo lsmdiskdumps
```


The resulting output

```
id          mdisk_filename
0           mdiskdump_desty5_2_090625_215452
```

lsmdisklba

The `lsmdisklba` command lists the MDisk and logical block address (LBA) for the specified VDisk LBA.

Syntax



Parameters

-vdisklba *vdisklba*

(Required) Specifies the 64-bit hexadecimal logical block address (LBA) on the VDisk.

-copy *id*

(Optional) Specifies the VDisk copy ID to list the MDisk and LBA for. If this parameter is not specified, the command lists MDisks and LBAs for all VDisk copies.

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

vdisk_id | *vdisk_name*

(Required) Specifies the VDisk name or ID.

Description

The `lsmdisklba` command returns the logical block address (LBA) of the MDisk that is associated with the VDisk LBA. For mirrored VDIsks, the command lists the MDisk LBA for both the primary and the copy.

If applicable, the command also lists the range of LBAs on both the VDisk and MDisk that are mapped in the same extent, or for space-efficient disks, in the same grain. If a space-efficient VDisk is offline and the specified LBA is not allocated, the command displays the VDisk LBA range only.

Table 16 on page 259 summarizes the data that can be returned with this command.

Table 16. `lsmdisklba` command output

Field	Fully allocated, single copy VDisk	LBA not allocated on space-efficient VDisk	Mirrored VDisk with one normal copy and one offline space-efficient copy	
			Normal copy	Space-efficient copy
<code>copy_id</code>	yes	yes	yes	yes
<code>mdisk_id</code>	yes	no	yes	no
<code>mdisk_name</code>	yes	no	yes	no
<code>type</code>	allocated	unallocated	allocated	offline
<code>mdisk_lba</code>	yes	no	yes	no
<code>mdisk_start</code>	yes	no	yes	no
<code>mdisk_end</code>	yes	no	yes	no
<code>vdisk_start</code>	yes	yes	yes	yes
<code>vdisk_end</code>	yes	yes	yes	yes

An invocation example

```
svcinfolsmdisklba -vdisk 0 -vdisklba 0x0
```

The resulting output

```
copy_id 0
mdisk_id 1
mdisk_name mdisk1
type allocated
mdisk_lba 0x00090000
mdisk_start 0x00090000
mdisk_end 0x000907FF
vdisk_start 0x00000000
vdisk_end 0x000007FF
```

lsmdiskcandidate

The `lsmdiskcandidate` command lists all of the unmanaged MDisk by MDisk ID.

Syntax

```
svcinfolsmdiskcandidate [-nohdr]
                        [-delim delimiter]
```

Parameters

`-nohdr`

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The `-nohdr` parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

`-delim delimiter`

(Optional) By default in a concise view, all columns of data are

space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

Description

This command displays a list of MDisks that are unmanaged. Only the MDisk IDs are displayed.

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

Note: The automatic detection performed by the cluster does not write anything to a unmanaged MDisk. It is only when you instruct the cluster to add an MDisk to a managed disk group or use a MDisk to create an image mode virtual disk that the storage is actually used.

Check to see which MDisks are available by issuing the **svctask detectmdisk** command to manually scan the fibre-channel network for any MDisks. Issue the **svcinfolismdiskcandidate** command to show the unmanaged MDisks. These MDisks have not been assigned to an MDisk group. Alternatively, you can issue the **svcinfolismdisk** command to view all of the MDisks.

An invocation example

```
svcinfolismdiskcandidate
```


The resulting output

```
id  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14
```

lismdiskextent

The **lismdiskextent** command displays the extent allocation between managed disks and virtual disks. The output lists a VDisk ID, VDisk copy ID, and the number of extents.

Syntax



Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

mdisk_name | *mdisk_id*

(Required) Specifies the specific object ID or name of the given type.

Description

The command displays a list, in which each entry contains a VDisk ID, VDisk copy ID, and the number of extents. These VDisk copies are using extents on the specified MDisk. The number of extents being used on each MDisk is also shown.

Every VDisk copy is constructed from one or more MDisks. At times, you might have to determine the relationship between the two objects. The following procedure allows you to determine the relationships.

To determine the relationship between VDisk copies and MDisks, issue the following command for each VDisk copy:

```
svcinfo lsvdiskmember vdisk_name | vdisk_id
```

where *vdisk_name* | *vdisk_id* is the name or ID of the VDisk copy. This displays a list of IDs that correspond to the MDisks that make up the VDisk copy.

To determine the relationship between VDisk copies and MDisks and the number of extents that are provided by each MDisk, you must use the command-line interface. For each VDisk copy, issue the following command:

```
svcinfo lsvdiskextent vdisk_name | vdisk_id
```

where *vdisk_name* | *vdisk_id* is the name or ID of the VDisk copy. This displays a table of MDisk IDs and the corresponding number of extents that each MDisk is providing as storage for the given VDisk copy.

To determine the relationship between MDisks and VDisk copies, issue the following command for each MDisk:

```
svcinfolsmdiskmember mdisk_name | mdisk_id
```

where *mdisk_name* | *mdisk_id* is the name or ID of the MDisk. This displays a list of IDs that correspond to the VDisk copies that are using this MDisk.

To determine the relationship between MDisks and VDisk copies and the number of extents that are used by each VDisk copy, you must use the command-line interface. For each MDisk, issue the following command:

```
svcinfolsmdiskextent mdisk_name | mdisk_id
```

where *mdisk_name* | *mdisk_id* is the name or ID of the MDisk. This command displays a table of VDisk copy IDs and the corresponding number of extents that are being used by each VDisk copy.

An invocation example

```
svcinfolsmdiskextent -delim : mdisk0
```

The resulting output

```
id:number_of_extents:copy_id  
1:1:1
```

lsmdiskgrp

The `lsmdiskgrp` command returns a concise list or a detailed view of MDisk groups visible to the cluster.

Syntax

```
svcinfolsmdiskgrp [-filtervalue -attrib=value] [-nohdr] [-bytes] [-delim delimiter] [-filtervalue?] [object_id | object_name]
```

Parameters

-filtervalue *attrib=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are returned. If a capacity is specified, the units must also be included.

Note: Some filters allow the use of a wildcard when you enter the command. The following rules apply to the use of wildcards with the SAN Volume Controller CLI:

- The wildcard character is an asterisk (*).
- The command can contain a maximum of one wildcard, which must be the first or last character in the string.
- When using a wildcard, you must enclose the filter entry within double quotation marks (""), as follows:

```
svcinfolsmdiskgrp -filtervalue "name=md*"
```

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-bytes

(Optional) Specifies that you want the report to display all capacities as bytes.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

object_id | object_name

(Optional) Specifies the name or ID of an object. When you use this parameter, the detailed view of the specific object is returned and any value specified by the **-filtervalue** parameter is ignored. If you do not specify the *object_id | object_name* parameter, the concise view of all objects matching the filtering requirements specified by the **-filtervalue** parameter are displayed.

-filtervalue?

Display a list of valid filter attributes. The valid filters for the **svcinfolsmdiskgrp** command are:

- name
- storage_pool_id
- mdisk_count
- vdisk_count
- extent_size
- status
- id

Description

This command returns a concise list or a detailed view of MDisk groups visible to the cluster.

Command output includes values for the following attributes:

status The state of the MDisk with the highest-priority status in the group, excluding image mode MDisks.

VDisk_count

The number of VDisk copies that are in the MDisk group.

capacity

The total amount of MDisk storage that is assigned to the MDisk group.

free_capacity

The amount of MDisk storage that is assigned to the MDisk group that is unused. MDisk storage can be used for cluster quorum data and VDIs.

real_capacity

The amount of MDisk storage that is assigned to the MDisk group that is assigned to VDIs.

virtual_capacity

The total virtual size of all the VDisk copies that are associated with the MDisk group. This is the same as the real_capacity value unless you have configured space-efficient VDisk copies in this MDisk group.

used_capacity

The total used size of all the VDisk copies that are associated with the MDisk group. This is the same as the real_capacity value unless you have configured space-efficient VDisk copies in this MDisk group.

overalllocation

Expressed as a percentage, the ratio of the virtual_capacity value to the capacity. An MDisk group overalllocation of over 100 is only possible if you have configured space-efficient VDisk copies.

warning

This field is a percentage. A warning is generated when the amount of space in the MDisk group that has been assigned exceeds this level.

The following define the status fields, from lowest to highest priority:

Online

The MDisk is online and available.

Degraded ports

There are one or more MDisk port errors.

Degraded paths

One or more paths to the MDisk have been lost; the MDisk is not online to every node in the cluster.

Offline

All paths to the MDisk are lost.

Excluded

The MDisk is excluded from use by the cluster; the MDisk port error count exceeded the threshold.

A concise invocation example

```
svcinflsmdiskgrp -delim :
```

The concise resulting output

```
id:name:status:mdisk_count:vdisk_count:capacity:extent_size:free_capacity:
virtual_capacity:used_capacity:real_capacity:overalllocation:warning
0:mdiskgrp0:degraded_paths:4:0:34.2GB:16:34.2GB:0:0:0:0:0
1:mdiskgrp1:online:4:6:200GB:16:100GB:400GB:75GB:100GB:200:80
```

A detailed invocation example

```
svcinfo lsmdiskgrp -delim : 0
```

The detailed resulting output

```
id:0
name:mdiskgrp0
status:online
mdisk_count:2
vdisk_count:26
capacity:33.3GB
extent_size:16
free_capacity:26.8GB
virtual_capacity:6.00GB
used_capacity:10.56GB
real_capacity:6.00GB
overallocation:18
warning:0
```

lsmdiskmember

The **lsmdiskmember** command displays a list of VDisks that are using extents on the specified MDisk. That is, the virtual disks are using extents on the managed disk that are specified by the MDisk ID.

Syntax

```
▶▶▶ svcinfo — — 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 VDisks that use extents of that MDisk.

Description

This command displays a list of virtual disks that use extents on the managed disk that are specified by the ID. The list displays members of the respective object and is independent of the state of the individual members; that is, if they are in offline state, they are still displayed.

Every VDisk is constructed from one or more MDisks. To determine the relationship between VDisk copies and MDisks, issue the following command:

```
svcinfolsvdiskmember vdisk_id | vdisk_name
```

where *vdisk_id* | *vdisk_name* is the name or ID of the VDisk copy. This displays a list of IDs that correspond to the MDisks that make up the VDisk copy.

To determine the relationship between VDisk copies and MDisks and the number of extents that are provided by each MDisk, you must use the command-line interface. For each VDisk copy, issue the following command:

```
svcinfolsvdiskextent vdisk_id | vdisk_name
```

where *vdisk_id* | *vdisk_name* is the name or ID of the VDisk copy. This command displays a table of MDisk IDs and the corresponding number of extents that each MDisk provides as storage for the VDisk copy.

To determine the relationship between MDisks and VDisk copies, issue the following command:

```
svcinfolsmdiskmember mdisk_id | mdisk_name
```

where *mdisk_id* | *mdisk_name* is the name or ID of the MDisk. This command displays a list of IDs that correspond to the VDisk copies that are using this MDisk.

To determine the relationship between MDisks and VDisk copies and the number of extents that are used by each VDisk copy, you must use the command-line interface. For each MDisk *mdisk_id* | *mdisk_name*, issue the following command:

```
svcinfolsmdiskextent mdisk_id | mdisk_name
```

where *mdisk_id* | *mdisk_name* is the name or ID of the MDisk. This command displays a table of VDisk copy IDs and the corresponding number of extents that are being used by each VDisk copy.

An invocation example

```
svcinfolsmdiskmember -delim : 1
```

The resulting output

```
id:copy_id
0:0
1:0
2:0
3:0
4:0
5:0
6:0
```

lsmigrate

The **lsmigrate** command displays the progress of all current data migration operations.

Syntax

```
svcinfo 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.

An invocation example

```
svcinfo lsmigrate -delim :
```

The resulting output

```
migrate_type:MDisk_Group_Migration
progress:96
migrate_source_vdisk_index:33
migrate_target_mdisk_grp:4
max_thread_count:4
migrate_source_vdisk_copy_id:1
```

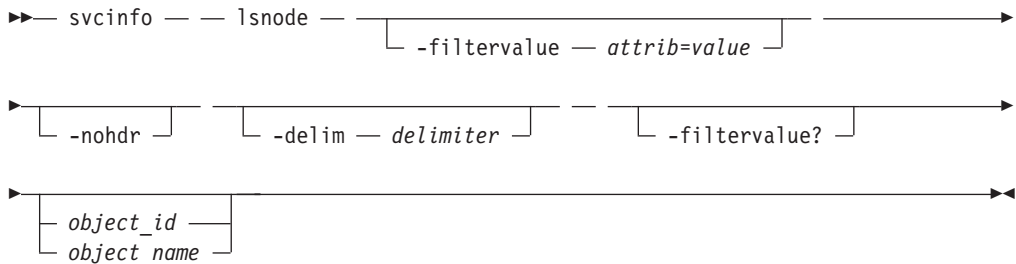
lsnode

The **lsnode** command returns a concise list or a detailed view of nodes visible to the cluster.

The list report style can be used to obtain two styles of report:

- A list containing concise information about all the nodes on a cluster. (Each entry in the list corresponds to a single node.)
- The detailed information about a single node.

Syntax



Parameters

-filtervalue *attrib=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are returned. If a capacity is specified, the units must also be included.

Note: Some filters allow the use of a wildcard when you enter the command. The following rules apply to the use of wildcards with the SAN Volume Controller CLI:

- The wildcard character is an asterisk (*).
- The command can contain a maximum of one wildcard.
- When using a wildcard, you must enclose the filter entry within double quotation marks (""), as follows:

```
svcinfo lsnode -filtervalue "name=md*"
```

-filtervalue?

Displays a list of valid filter attributes for the **-filtervalue** *attribute=value* parameter. The valid filters for the **svcinfo lsnode** command are:

- node_name
- id
- status
- IO_group_name
- IO_group_id
- name
- hardware

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by

a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

object_id | object_name

(Optional) Specifies the name or ID of an object. When you use this parameter, the detailed view of the specific object is returned and any value that is specified by the **-filtervalue** parameter is ignored. If you do not specify the *object_id | object_name* parameter, the concise view of all objects matching the filtering requirements that is specified by the **-filtervalue** parameter are displayed.

Description

This command returns a concise list or a detailed view of nodes visible to the cluster. Table 17 provides the possible values that are applicable to the attributes that are displayed as data in the output views.

Table 17. *svcinfo lsnode* attribute values

Attribute	Value
status	offline flushing pending online adding deleting
config_node	no yes
port_status	active inactive not installed
hardware	8F2 8F4 8G4 CF8 8A4 other

To determine a node's WWPNs, enter the following command:

```
svcinfo lsnode
```

Note: Record the node name or ID because you use this value in the next step. For the node or nodes in question, issue the following command:

```
svcinfo lsnode node_name | node_id
```

where *node_name* | *node_id* is the node name or ID.

Note: Record the four port IDs (WWPNs).

Attention: If the node is in the adding state, the WWPN is displayed as 0000000000000000. Once the node has successfully become a member of the cluster, the state changes to online and the WWPN is displayed correctly.

A concise invocation example

```
svcinfo lsnode
```

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
81 dvt104607 1000849036 5005076801000070 online 0 io_grp0 no 20400002042400C6
8G4 iqn.1986-03.com.ibm:2145.1dcluster-65.dvt104607 dvt104607
83 h1cn114263 1000849028 500507680100056C online 2 io_grp2 no 2040000204240088
8A4 iqn.1986-03.com.ibm:2145.1dcluster-65.h1cn114263 h1cn114263
84 h1cn114302 100083B159 50050768010000FC online 3 io_grp3 no 2040000203481149
```

```

8A4 iqn.1986-03.com.ibm:2145.ldcluster-65.hlcn114302 hlcn114302
85 hlcn114252 1000849020 5005076801000D79 online 0 io_grp0 no 2040000204240080
8A4 iqn.1986-03.com.ibm:2145.ldcluster-65.hlcn114252 hlcn114252
82 dvt101952 100062L097 500507680100055C online 1 io_grp1 yes 2040000182700247
8F4 iqn.1986-03.com.ibm:2145.ldcluster-65.dvt101952 dvt101952
86 dvt101036 100062L102 50050768010000D0 online 1 io_grp1 no 2040000182701002
8F4 iqn.1986-03.com.ibm:2145.ldcluster-65.dvt101036 dvt101036
87 hlcn114163 1000840052 5005076801000023 online 2 io_grp2 no 20400002047C0142
8A4 iqn.1986-03.com.ibm:2145.ldcluster-65.hlcn114163 hlcn114163
88 hlcn114414 100083B158 500507680100004A online 3 io_grp3 no 2040000203481148
8A4 iqn.1986-03.com.ibm:2145.ldcluster-65.hlcn114414 hlcn114414

```

A detailed invocation example

```
svcinfn lsnode -delim : node1
```

The detailed resulting output

```

id:1
name:node1
UPS_serial_number:1000849042
WWNN:5005076801007B24
status:online
IO_group_id:0
IO_group_name:io_grp0
partner_node_id:
partner_node_name:
config_node:yes
UPS_unique_id:10000000000007B24
port_id:5005076801107B24
port_status:active
port_speed:2
port_id:5005076801207B24
port_status:inactive
port_speed:2
port_id:0000000000000000
port_status:not_installed
port_speed:2
port_id:0000000000000000
port_status:not_installed
port_speed:2
hardware:other
iscsi_name:iqn.1986-03.com.ibm:2145.rack8.dvt110683
iscsi_alias:rapchikvgadekar
failover_active:no
failover_name:node1
failover_iscsi_name:iname1
failover_iscsi_alias:ialias1

```

lsnodecandidate

The lsnodecandidate command lists all of the nodes that are available to add to the cluster.

Syntax

```

>>> svcinfn -- lsnodecandidate -- [-nohdr]

```

```

>>> [-delim delimiter]

```


Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

Description

This command displays a list of nodes that are available to add to the cluster. This includes nodes that are not already part of a cluster, but are compatible with the cluster software level. Nodes with hardware types that are incompatible with the installed software are not listed.

An invocation example

```
svcinfolsnodecandidate -delim :
```

The resulting output

```
id: panel_name: UPS_serial_number: UPS_unique_id: hardware
1: 146355: 10L3ASH: 202378101C0D18D8: 8G4
```

lsnodedependentvdisks

The `lsnodedependentvdisks` command displays the virtual disks (VDisks) that are dependent on the status of the specified node.

Syntax

```
▶▶— svcinfolsnodedependentvdisks — [ -nohdr ] —▶▶
▶ [ -delim delimiter ] [ node_id node_name ] —▶▶
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. Using the **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

node_id | *node_name*

(Required) Specifies the node by its ID or name.

Description

The `lsnodedependentvdisks` command lists the VDisks that are dependent on the status of the specified node. If the node goes offline or is removed from the cluster, the dependent VDisks also go offline. Before taking a node offline or removing it from the cluster, you can use the command to ensure that you do not lose access to any VDisks.

The `lsnodedependentvdisks` command also checks the expected state of quorum disks. If all available quorum disks are only currently accessible through the specified node, the command returns an error. The quorum disks should be moved to MDisks that are accessible through all nodes.

Various scenarios can produce node-dependent VDisks. The following examples are common scenarios in which the `lsnodedependentvdisks` command will return node-dependent VDisks:

1. The node contains solid-state drives (SSDs) and also contains the only synchronized copy of a mirrored VDisk.
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 VDisks 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 VDisk mirror synchronizations between SSD MDisks to complete. To resolve (2-4), bring any offline MDisks online and repair any degraded paths.

Note: The command lists the VDisks that are dependent on the nodes at the time the command is run. Subsequent changes to your system require running the command again.

An invocation example

```
svcinfo lsnodedependentvdisks node1
```

The resulting output

<i>vdisk_id</i>	<i>vdisk_name</i>
0	vdisk0
1	vdisk1

lsnodevpd

The `lsnodevpd` command displays the vital product data (VPD) for each node.

Syntax

```
svcinfo -- lsnodevpd [-nohdr] [-delim delimiter]
node_id | node_name
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim delimiter

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. Using the **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

node_id | node_name

(Required) Specifies the node by its ID or name.

Description

This command displays the VPD for the specified node. Each field is reported on a new line. All fields are strings. The VPD is split into sections. Each section has a section heading. The number of fields in that section follows each heading. Each section is separated by an empty line.

For example:

```
section name:3 fields
field1:value
field2:value
field3:value
```

```
new section:x fields
```

```
...
```

Some sections contain information about multiple objects of that type. Each object within the section is separated by an empty line.

For example:

section name:4 fields
object1 field1:value
object1 field2:value

object2 field1:value
object2 field2:value

new section: x fields
...

Note: For 8F4, 8G4, and 8A4 nodes, the VPD displays the device serial number of the fibre-channel card as N/A.

An invocation example

```
svcinfo lsnodevpd 1
```

The resulting output

```
id 1  
  
system board: 21 fields  
part_number 43V7072  
system_serial_number KD1438A  
number_of_processors 4  
number_of_memory_modules 6  
number_of_fans 6  
number_of_FC_cards 1  
number_of_scsi/ide_devices 2  
BIOS_manufacturer IBM Corp.  
BIOS_version -[D6E124AUS-1.01]-  
BIOS_release_date 04/30/2009  
system_manufacturer IBM  
system_product IBM System x -[2145CF8]-  
version 00  
planar_manufacturer IBM  
planar_product 49Y6498  
planar_version (none)  
power_supply_part_number 39Y7201  
CMOS_battery_part_number 33F8354  
frame_assembly_part_number  
ethernet_cable_part_number  
service_processor_firmware 1.01  
  
processor: 6 fields  
processor_location Processor 1  
manufacturer Intel(R) Corporation  
version Intel(R) Xeon(R) CPU E5530 @ 2.40GHz  
speed 2400  
status Enabled  
CPU_part_number 46D1266  
  
memory module: 96 fields  
part_number 44T1493  
device_location DIMM01  
bank_location BANK01  
size (MB) No Module Installed  
manufacturer Not Specified  
serial_number Not Specified  
  
part_number 44T1493  
device_location DIMM02  
bank_location BANK02  
size (MB) 4096  
manufacturer Samsung
```

```

serial_number 99062848

part_number 44T1493
device_location DIMM03
bank_location BANK03
size (MB) 4096
manufacturer Samsung
serial_number C7062848
...

fan: 12 fields
part_number 43V6929
location location1

part_number 43V6929
location location2

part_number 43V6929
location location3
...


Adapter card: 18 fields
card_type FC card
part_number 31P1337
port_numbers 1 2 3 4
location 0
device_serial_number 11S31P1333YM10MY96A206
manufacturer IBM
device QE8
card_revision 2
chip_revision 2.0

card_type SAS card
part_number 44E8690
port_numbers 1 2 3 4
location 0
device_serial_number 11S31P1299YM10MY948004
manufacturer IBMHUR
device Capri-PMC8001
card_revision Y
chip_revision 1.1

Fibre channel SFP: 48 fields
part_number 17P9211
manufacturer JDSU
device PLRXPLVCSH4921
serial_number C915EB06V
supported_speeds 2,4,8
connector_type LC
transmitter_type SN
wavelength 850
max_distance_by_cable_type OM1:20,OM2:50,OM3:150
hw_revision 1
port_number 1
WWPN 500507680140350d
...

device: 15 fields
part_number 31P1339
bus USB
device 0
model IBM USB Endeavour
revision 1.0
serial_number NA
approx_capacity 0
hw_revision 0

```

Parameters

-filtervalue *attrib=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are returned. If a capacity is specified, the units must also be included.

Note: Some filters allow the use of a wildcard when you enter the command. The following rules apply to the use of wildcards with the SAN Volume Controller CLI:

- The wildcard character is an asterisk (*).
- The command can contain a maximum of one wildcard, which must be the first or last character in the string.
- When using a wildcard, enclose the filter entry within double quotation marks (""), as follows:

```
svcinfolspportip -filtervalue "node_name=md*"
```

-filtervalue?

(Optional) Displays the valid filter attributes. The following filter attributes for the `svcinfolspportip` 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, either 1 or 2. If omitted, a concise view is displayed for all ports. When you use this parameter, the detailed view of the specified port is returned and any value that is specified by the **-filtervalue** parameter is ignored. If you do not use the *ethernet_port_id*

parameter, the concise view displays all ports matching the filtering requirements that are specified by the **-filtervalue** parameter.

Description

This command lists all port IP addresses for each node in the cluster. The concise view displays two rows of output for each ethernet port. Each node has two ethernet ports.

Use the `svcinfo lspportip` command with the optional **ethernet_port_id** parameter to display a detailed view of the specified port.

Both output rows for a port show the MAC address of that port if it can be determined. If the node and the ethernet link are online, the rows also show the speed and duplex state of the link. The duplex field can have values of `Half` or `Full`, or it is blank if the node is offline.

The first row for each port shows any iSCSI addresses that have been configured for that port and are not failed over to a different node. The failover field on this row is set to `no`. The second row for each port shows any iSCSI addresses that have been configured for the partner node, or for the local node with failover, and that are active on the port. The failover field on this row is set to `yes`.

The state field is set to `unconfigured` if there are no iSCSI addresses configured on the port. The state field is set to `offline` if there are configured addresses but the link is down, and `online` if the link is up. Any offline rows indicate a potential problem.

A concise invocation example

```
svcinfo lspportip -delim :
```

The concise resulting output

```
id:node_id:node_name:IP_address:mask:gateway:IP_address_6:
prefix_6:gateway_6:MAC:duplex:state:speed:failover
1:1:dvt101794:9.71.47.129:255.255.254.0:9.71.46.1:::00:14:
5e:33:51:92:Half:online:100Mb/s:no
1:1:dvt101794:::00:14:5e:33:51:92:Half:online:100Mb/s:yes
2:1:dvt101794:::00:14:5e:33:51:93::unconfigured::no
2:1:dvt101794:::00:14:5e:33:51:93::unconfigured::yes
1:2:dvt101760:9.71.47.83:255.255.254.0:9.71.46.1:::00:14:5e:
7e:2a:58:Half:online:100Mb/s:no
1:2:dvt101760:::00:14:5e:7e:2a:58:Half:online:100Mb/s:yes
2:2:dvt101760:::00:14:5e:7e:2a:59::unconfigured::no
2:2:dvt101760:::00:14:5e:7e:2a:59::unconfigured::yes
1:3:dvt101761:9.71.47.253:255.255.254.0:9.71.46.1:::00:14:5e:
33:50:fa:Half:online:100Mb/s:no
1:3:dvt101761:::00:14:5e:33:50:fa:Half:online:100Mb/s:yes
2:3:dvt101761:::00:14:5e:33:50:fb::unconfigured::no
2:3:dvt101761:::00:14:5e:33:50:fb::unconfigured::yes
1:4:dvt101786:9.71.47.227:255.255.254.0:9.71.46.1:::00:14:5e:
33:50:da:Half:online:100Mb/s:no
1:4:dvt101786:::00:14:5e:33:50:da:Half:online:100Mb/s:yes
2:4:dvt101786:::00:14:5e:33:50:db::unconfigured::no
2:4:dvt101786:::00:14:5e:33:50:db::unconfigured::yes
1:5:destiny35:9.71.47.69:255.255.254.0:9.71.46.1:::00:21:5e:09:
21:44:Full:online:1Gb/s:no
1:5:destiny35:::00:21:5e:09:21:44:Full:online:1Gb/s:yes
2:5:destiny35:::00:21:5e:09:21:46::unconfigured::no
2:5:destiny35:::00:21:5e:09:21:46::unconfigured::yes
1:6:destiny34:9.71.46.239:255.255.254.0:9.71.46.1:::00:21:5e:09:
```



```

|         21:54:Full:online:100Mb/s:no
|         1:6:destiny34:::00:21:5e:09:21:54:Full:online:100Mb/s:yes
|         2:6:destiny34:::00:21:5e:09:21:56::unconfigured::no
|         2:6:destiny34:::00:21:5e:09:21:56::unconfigured::yes

```

A detailed invocation example

```
svcinfolspportip 1
```

The detailed resulting output

```

| id 1
| node_id 1
| node_name dvt101794
| IP_address 9.8.7.1
| mask 255.255.255.0
| gateway 9.0.0.1
| IP_address_6
| prefix_6
| gateway_6
| MAC 00:14:5e:33:51:92
| duplex Half
| state online
| speed 100Mb/s
| failover no

```

lsquorum

The `lsquorum` command lists the managed disks (MDisks) that the cluster is currently using to store quorum data.

Syntax

```

▶▶ svcinfo — lsquorum — [ -nohdr ] [ -delim delimiter ]
▶ [ quorum_index ]

```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **delim** parameter overrides this behavior. Valid input for the **delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified character.

quorum_index

(Optional) Specifies the quorum disk by its index number. The number can be

either 0, 1, or 2. When you use this parameter, a detailed view of the specified disk is returned. If you do not specify a disk, then a concise view of all quorum disks is displayed.

Description

This command displays a concise list or a detailed view of the MDisks that the cluster is currently using to store quorum data. This information can be used to ensure that the quorum candidates are on separate storage subsystems.

A concise invocation example

```
svcinfolsqorum
```

The concise resulting output

```
quorum_index status id name controller_id controller_name active
0 online 987 mdisk78 5 DS8K no
1 online 2001 mdisk1 2 DS4K-1 no
2 online 309 mdisk9 1 IBM-3 yes
```

A detailed invocation example

```
svcinfolsqorum 2
```

The detailed resulting output

```
quorum_index 2
status online
id 309
name mdisk9
controller_id 1
controller_name IBM-3
active yes
```

lsrconsistgrp

The **lsrconsistgrp** command returns a concise list or a detailed view of Metro or Global Mirror consistency groups visible to the cluster.

The list report style can be used to obtain two styles of report:

- A list containing concise information about all the Metro or Global Mirror consistency groups visible to the cluster. (Each entry in the list corresponds to a single Metro or Global Mirror consistency group.)
- The detailed information about a single Metro or Global Mirror consistency group.

Syntax

```
➤➤➤ svcinfol — — lsrconsistgrp — [ -filtervalue — attrib=value ] ➤➤➤
➤ [ -nohdr ] [ -delim — delimiter ] [ -filtervalue? ] ➤
➤ [ object_id ] [ object_name ] ➤➤➤
```

Parameters

-filtervalue *attrib=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are displayed. If a capacity is specified, the units must also be included.

Note: Some filters allow the use of a wildcard when you enter the command. The following rules apply to the use of wildcards with the SAN Volume Controller CLI:

- The wildcard character is an asterisk (*).
- The command can contain a maximum of one wildcard, which must be the first or last character in the string.
- When using a wildcard, you must enclose the filter entry with double quotation marks (""), as follows:

```
svcinfo lsrcconsistgrp -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 **svcinfo lsrcconsistgrp** command are valid:

- group_id
- name
- master_cluster_id
- master_cluster_name
- aux_cluster_id
- aux_cluster_name
- primary
- state

- relationship_count
- id
- copy_type

Description

This command returns a concise list or a detailed view of Metro or Global Mirror consistency groups that are visible to the cluster.

Table 18 provides possible values for the attributes that are displayed as data in the output views.

Table 18. lsrcconsistgrp command output values

Attribute	Value
primary	n/a, master, aux
state	inconsistent_stopped, inconsistent_copying, consistent_stopped, consistent_synchronized, idling, idling_disconnected, inconsistent_disconnected consistent_disconnected, empty
freeze_time	The time in YY/MM/DD/HH/MM format.
status	online, primary_offline, secondary_offline
sync	in_sync, out_of_sync
copy_type	metro, global, empty_group

Note: The names of the Metro or Global Mirror relationships and consistency groups might be blank if the relationship or consistency groups are intercluster and the cluster partnership is disconnected.

A concise invocation example

```
svcinfo 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

248:jdemo_BA_cons1:0000020060406746:clusterB:0000020061413ABA:clusterA:master:
consistent_stopped:2:global
249:rccstgrp0:0000020061413ABA:clusterA:0000020061413ABA:clusterA::empty:0
:empty_group
250:jdemo_BA_cons2:0000020060406746:clusterB:0000020061413ABA:clusterA:master:
inconsistent_stopped:1:metro
251:BA_cons1:0000020060406746:clusterB:0000020061413ABA:clusterA:master:
consistent_stopped:4:metro
252:AB_cons2:0000020061413ABA:clusterA:0000020060406746:clusterB::empty:0
:empty_group
253:AB_cons1:0000020061413ABA:clusterA:0000020060406746:clusterB:aux:
consistent_stopped:3:global
254:AA_cons2:0000020061413ABA:clusterA:0000020061413ABA:clusterA::empty:0
:empty_group
255:AA_cons1:0000020061413ABA:clusterA:0000020061413ABA:clusterA:master:
consistent_synchronized:2:global
```

A detailed invocation example

```
svcinfo 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:inconsistent_stopped
relationship_count:1
freeze_time:
status:online
sync:
copy_type:metro
RC_rel_id:2
RC_rel_name:aaa

```

lsrrelationship

The `lsrrelationship` command returns a concise list or a detailed view of Metro or Global Mirror relationships visible to the cluster.

The list report style can be used to obtain two styles of report:

- A list containing concise information about all the Metro or Global Mirror relationships visible to the cluster. (Each entry in the list corresponds to a single Metro or Global Mirror relationship.)
- The detailed information about a single Metro or Global Mirror relationship.

Syntax

```

>> svcinfo -- lsrrelationship -- [-filtervalue -- attrib=value]

```

[-nohdr] [-delim -- delimiter] [-filtervalue?]

[object_id] [object_name]

Parameters

-filtervalue *attribute=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are returned. If a capacity is specified, the units must also be included.

Note: Some filters allow the use of a wildcard when you enter the command. The following rules apply to the use of wildcards with the SAN Volume Controller CLI:

- The wildcard character is an asterisk (*).
- The command can contain a maximum of one wildcard, which must be the first or last character in the string.
- When using a wildcard, you must enclose the filter entry with double quotation marks (" "), as follows:

```

svcinfo lsrrelationship -filtervalue "name=md*"

```

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed even if the **-nohdr** parameter is specified.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

object_id | object_name

(Optional) Specifies the name or ID of an object. When you use this parameter, the detailed view of the specific object is returned and any value that is specified by the **-filtervalue** parameter is ignored. If you do not specify the *object_id | object_name* parameter, the concise view of all objects matching the filtering requirements that are specified by the **-filtervalue** parameter are displayed.

-filtervalue?

(Optional) Specifies that you want your report to display any or all of the list of valid filter attributes. The valid filter attributes for the **svcinfo lsrelationship** command are:

- RC_rel_id
- RC_rel_name
- master_cluster_id
- master_cluster_name
- master_vdisk_id
- master_vdisk_name
- aux_cluster_id
- aux_cluster_name
- aux_vdisk_id
- aux_vdisk_name
- primary
- consistency_group_id
- consistency_group_name
- state
- progress
- copy_type

Description

This command returns a concise list or a detailed view of Metro or Global Mirror relationships visible to the cluster.

Table 19 provides possible values for the attributes that are displayed as data in the output views.

Table 19. *lsrrelationship* command attributes and values

Attribute	Value
primary	n/a, master, aux
state	inconsistent_stopped, inconsistent_copying, consistent_stopped, consistent_synchronized, idling, idling_disconnected, inconsistent_disconnected, consistent_disconnected
progress	0-100, n/a
freeze time	The time in YY/MM/DD/HH/MM format.
status	online, primary_offline, secondary_offline
sync	n/a, in_sync, out_of_sync
copy_type	metro, global

Note: The names of the Metro or Global Mirror relationships and consistency groups can be blank if the relationship or consistency groups are inter-cluster and the cluster partnership is disconnected.

A concise and detailed invocation example

```
svcinfo lsrrelationship -delim : -filtervalue name=j*
```

The concise and detailed 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
45:jrel_AB1:0000020061413ABA:clusterA:45:jdisk_B8:0000020060406746:clusterB:38:j
disk_B1:master:::consistent_stopped:50:metro
48:jrel_AB2:0000020061413ABA:clusterA:48:jdisk_A4:0000020060406746:clusterB:41:j
disk_B4:master:::consistent_synchronised:50:metro
49:jrel_BA_1:0000020060406746:clusterB:42:jdisk_B5:0000020061413ABA:clusterA:49:j
disk_A5:master:248:jdemo_BA_cons1:consistent_stopped:50:metro
50:jrel_BA_2:0000020060406746:clusterB:43:jdisk_B6:0000020061413ABA:clusterA:
50:jdisk_A6:master:248:jdemo_BA_cons1:consistent_stopped:50:metro
```

A detailed invocation example

```
svcinfo lsrrelationship -delim : AB_2
```

The detailed resulting output

```
id:9
name:AB_2
master_cluster_id:0000020061413ABA
master_cluster_name:clusterA
master_vdisk_id:9
master_vdisk_name:stripe9
aux_cluster_id:0000020060406746
aux_cluster_name:clusterB
aux_vdisk_id:9
aux_vdisk_name:stripe9_b
primary:master
consistency_group_id:
consistency_group_name:
state:consistent_stopped
bg_copy_priority:50
progress:
```

```
freeze_time:2006/05/05/08/26/46
status:secondary_offline
sync:in_sync
copy_type:metro
```

lsrrelationshipcandidate

The **lsrrelationshipcandidate** command lists VDIs that are eligible to form Metro or Global Mirror relationships. You can list eligible VDIs that are on the local or remote cluster.

Syntax

```
►► svcinfo — — lsrrelationshipcandidate —————►►
|
| ┌ -master ──┬─ master_vdisk_id ──┬─┬─ -aux ──┬─ aux_cluster_id ──┬─
| │           │ └─ master_vdisk_name ┘ │ │           │ └─ aux_cluster_name ┘
| └──────────┘ └──────────────────┘ └──────────┘ └──────────────────┘
|
| ┌ -nohdr ┘ ┌ -delim ── delimiter ──┘
| └────────┘ └────────────────────────┘
|
|◀◀
```

Parameters

-master *master_vdisk_id* | *master_vdisk_name*

(Required) Specifies a particular VDisk to use as the master VDisk. The command finds candidates that match the size of this VDisk. If you are requesting candidate VDIs on the local cluster, this command also matches the *io_group*.

-aux *aux_cluster_id* | *aux_cluster_name*

(Required) Specifies a remote cluster with VDisk candidates for an intercluster relationship. If you do not specify this parameter, the candidates on the local cluster are displayed.

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

Description

This command displays a list of VDIs that can be either the master or the auxiliary disk for a Metro or Global Mirror relationship. VDisk IDs and names are displayed.

An invocation example

```
svcinfo lsrrelationshipcandidate -delim :
```

The resulting output

```
id:vdisk_name  
0:vdisk0  
4:vdisk4
```

lsrrelationshipprogress

You can use the **lsrrelationshipprogress** command to display the progress of the background copy of a Metro Mirror or Global Mirror relationship as a percentage. When the initial background copy process for a relationship has completed, null is displayed for the progress of that relationship.

Syntax

```
►— svcinfo — — lsrrelationshipprogress — —————►  
└── -nohdr ─┘  
  
└── -delim — delimiter ─┘ └── rrelationship_id ─┘  
└── rrelationship_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.

rrelationship_id | *rrelationship_name*

Specifies the object ID or name of the specified type.

Description

This command displays the progress of the background copy of a Metro Mirror or Global Mirror relationship as a percentage.

An invocation example

```
svcinfo lsrrelationshipprogress -delim : 0
```

The resulting output

lsrepairsevdiskcopyprogress

The **lsrepairsevdiskcopyprogress** command lists the repair progress for space-efficient VDisk copies.

Syntax

```
▶▶ svcinfo — lsrepairsevdiskcopyprogress — [ -nohdr ] —————▶▶
▶ [ -delim — delimiter ] [ -copy — id ] [ vdisk_name ] [ vdisk_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.

-copy *id*

(Optional) Lists the repair progress for the specified copy.

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

vdisk_name | *vdisk_id*

(Optional) Specifies the virtual disk name or ID to list repair progress for. You must specify this parameter last on the command line. If you do not enter this parameter, the command lists progress for all space-efficient copies in the cluster.

Description

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

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

An invocation example

```
svcinfo lsrepairsevdiskcopyprogress -delim :
```

The resulting output

```
vdisk_id:vdisk_name:copy id:task:progress:estimated_completion_time  
0:vdisk0:0:repairing:50:070301120000  
0:vdisk0:1:repairing:51:070301120000  
1:vdisk1:0:repairing:32:070301153500
```

An invocation example

```
svcinfo lsrepairsevdiskcopyprogress -delim : vdisk0
```

The resulting output

```
vdisk_id:vdisk_name:copy id:task:progress:estimated_completion_time  
0:vdisk0:0:repairing:50:070301120000  
0:vdisk0:1:repairing:51:070301120000
```

An invocation example

```
svcinfo lsrepairsevdiskcopyprogress -delim : -copy 1 vdisk0
```

The resulting output

```
vdisk_id:vdisk_name:copy id:task:progress:estimated_completion_time  
0:vdisk0:1:repairing:51:070301120000
```

lsrepairvdiskcopyprogress

The **lsrepairvdiskcopyprogress** command displays the progress of mirrored VDisk repairs and validations.

Syntax

```
➤➤—svcinfo— —lsrepairvdiskcopyprogress— —————→  
                                         ┌ -copy — id ─┐  
➤┐───────────────────────────────────────────────────────────────────────────────────➤  
└─vdisk_name─┘  
└─vdisk_id──┘
```

Parameters

-copy *id*

(Optional) Specifies the VDisk copy ID to list repair progress for. If you do not specify this parameter, progress is displayed for all copies.

vdisk_name | *vdisk_id*

(Optional) Specifies the virtual disk name or ID to list repair progress for. You must specify this parameter last on the command line.

Description

The **lsrepairvdiskcopyprogress** command displays the progress of repairs and validations being made to mirrored VDIs. Use this command to track progress after running the **svctask repairvdiskcopy** command. You can specify a VDisk copy using the **-copy** *id* parameter. To display the VDIs that have two or more copies with an active task, specify the command with no parameters; it is not possible to have only one VDisk copy with an active task.

The command displays progress for the following types of VDisk copies:

- All VDisk copies display the same task; validate, medium or resync, depending on the specified parameter.
- All VDisk copies display the same percentage and estimated completion time.
- If specified, non-mirrored VDIs are displayed as a single copy with a blank task; they are not displayed in the full concise view.
- Once a task completes, the task is blank for all copies.
- If the task is blank, the percentage and the completion time are also blank.

An invocation example

```
svcinfolrepairvdiskcopyprogress -delim :
```

The resulting output

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

An invocation example

```
svcinfolrepairvdiskcopyprogress -delim : vdisk0
```

The resulting output

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

An invocation example

```
svcinfolsvdiskcopyrepairprogress -delim : -copy 0 vdisk0
```

The resulting output

```
vdisk_id:vdisk_name:copy id:task:progress:estimated_completion_time
0:vdisk0:0:medium:50:070301120000
```

lsrmvdiskdependentmaps

The `lsrmvdiskdependentmaps` command displays all FlashCopy mappings that must be stopped for the specified VDisk to be deleted.

Syntax

```
svcinfol - -lsrmvdiskdependentmaps -nohdr
-vdelim - delimiter - vdisk_name vdisk_id
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

| **-delim delimiter**

| (Optional) By default in a concise view, all columns of data are
| space-separated. The width of each column is set to the maximum possible
| width of each item of data. In a detailed view, each item of data has its own
| row, and if the headers are displayed, the data is separated from the header by
| a space. The **-delim** parameter overrides this behavior. Valid input for the
| **-delim** parameter is a one-byte character. If you enter **-delim :** on the
| command line, the colon character (:) separates all items of data in a concise
| view; for example, the spacing of columns does not occur. In a detailed view,
| the data is separated from its header by the specified delimiter.

| *vdisk_name* | *vdisk_id*

| (Required) Specifies the name or ID of the VDisk for which the FlashCopy
| mappings are displayed.

Description

| This command returns a list of the FlashCopy mappings that must be stopped
| before the specified VDisk can be deleted. Any mappings that are returned in the
| list for the VDisk are automatically stopped when the VDisk is deleted with the
| **force** option.

An invocation example

```
svcinflsrmvdiskdependentmaps -delim : 0
```

The resulting output

```
id:name  
2:fcmap2  
5:fcmap5
```

lsroute

| The **lsroute** command displays the IP routing table.

Syntax

| **svcinflsroute**

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 **svcinflsroute** command is equivalent to the Linux **route** command.

An invocation example

```
svcinflsroute
```

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

```

lssevdiskcopy

The **lssevdiskcopy** command lists the space-efficient copies of the specified VDisks.

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) Specifies the VDisk copy to list space-efficient copies for. You must specify a *vdisk_name* | *vdisk_id* value with this parameter.

-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 **svcinfo lssevdiskcopy** command are valid:

- mdisk_grp_id
- mdisk_grp_name
- overallocation
- autoexpand
- grainsize

vdisk_name | *vdisk_id*

(Optional) Specifies the virtual disk name or ID to list space-efficient copies for.

You must specify this parameter last on the command line. If you do not enter this parameter, the command lists all space-efficient VDisk copies in the cluster.

Description

The **lssevdiskcopy** command lists all space-efficient copies of the specified VDisk. If you do not specify a VDisk, the command lists all space-efficient copies in the cluster.

The command provides a concise view of the space-efficient properties of the selected VDisk copies. Run the **svcinfolsvdiskcopy** command to see a concise view of the properties that are common to space-efficient and non-space-efficient VDisk copies. See the description of **svcinfolsvdisk** command for a description of the fields that is shown in the view.

An invocation example

```
svcinfolsvdiskcopy -delim :
```

The resulting output

```
vdisk_id:vdisk_name:copy_id:mdisk_grp_id:mdisk_grp_name:capacity:used_capacity:real_capacity:
free_capacity:overallocation:autoexpand:warning:grainsize
0:vv1:0:0:ppp:16.0GB:5.0MB:4.0GB:15.99GB:400:off:20:32
1:se1:0:0:ppp:16.0GB:1.0GB:4.0GB:15.00GB:400:off:20:32
1:se1:1:0:ppp:16.0GB:2.0GB:8.0GB:14.00GB:200:off:45:256
```

An invocation example

```
svcinfolsvdiskcopy -delim : se1
```

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
1:se1:0:0:ppp:16.0GB:1.0GB:4.0GB:15.00GB:400:off:20:32
1:se1:1:0:ppp:16.0GB:2.0GB:8.0GB:14.00GB:200:off:45:256
```

An invocation example

```
svcinfolsvdiskcopy -delim : -copy 0 0
```


The resulting output

```
id:0
name:vdisk0
capacity:128.0MB
copy_id:0
status:online
sync:yes
primary:yes
mdisk_grp_id:0
mdisk_grp_name:mdiskgrp0
type:striped
mdisk_id:
mdisk_name:
fast_write_state:empty
used_capacity:0.41MB
real_capacity:128.00MB
free_capacity:127.59MB
overallocation:100
autoexpand:off
warning:79
grainsize:32
```

lssnmpserver

The `lssnmpserver` command returns a concise list or a detailed view of SNMP servers that are configured on the cluster.

Syntax



Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

snmp_server_name | snmp_server_id

(Optional) Specifies the name or ID of an existing SNMP server that must be listed.

Description

Use this command to display a concise list or a detailed view of SNMP servers that are configured on the cluster.

A concise invocation example

```
svcinfo 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

```
svcinfo 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

```

lssoftwaredumps

Use the **lssoftwaredumps** command to display a list of software packages from the **/home/admin/upgrade** directory.

Syntax

```

➤➤ svcinfo -- lssoftwaredumps -- [-nohdr]

```

```

➤➤ [-delim delimiter] [-node_id node_name]

```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. 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.

node_id | *node_name*

(Optional) Specifies the node ID or name to list the available dumps of the specific type. If you do not specify a node, the dumps that are available on the configuration node are listed.

Description

This command displays a list of software upgrade packages. These packages are copied as a result of software upgrades. If you do not specify a node, the packages that are available on the configuration node are listed. The command displays files from the **/home/admin/upgrade** directory.

An invocation example

```

svcinfo lssoftwaredumps

```

The resulting output

```
id          software_filename
0          s1_mala75_030405_092143
1          s2_mala75_030405_092145
2          s3_mala75_030405_092146
```

lssoftwareupgradestatus

The `lssoftwareupgradestatus` command displays the status of a software upgrade.

Syntax

```
▶▶ svcinfo — — lssoftwareupgradestatus — — [nohdr] ▶▶
```

Parameters

`-nohdr`

(Optional) Suppresses the display of headings.

Description

The `lssoftwareupgradestatus` command displays the status of a software upgrade.

Note: If a status of `stalled_non_redundant` is displayed, proceeding with the remaining set of node upgrades might result in offline VDIs. Contact an IBM service representative to complete the manual upgrade.

An invocation example

```
svcinfo lssoftwareupgradestatus
```

The resulting output

```
status
upgrading
```

An invocation example

```
svcinfo lssoftwareupgradestatus
```

The resulting output

```
status
stalled_non_redundant
```


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.

lssyslogserver

The `lssyslogserver` command returns a concise list or a detailed view of syslog servers that are configured on the cluster.

Syntax



Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

syslog_server_name | syslog_server_id

(Optional) Specifies the name or ID of an existing syslog server. When you use this parameter, a detailed view of the specified syslog server is returned. If you do not specify a syslog server name or ID, then a concise view of all syslog servers is displayed.

Description

Use this command to display a concise list or a detailed view of syslog servers that are configured on the cluster.

A concise invocation example

```
svcinfo lssyslogserver -delim :
```

The concise resulting output

```
id:name:IP_address:facility:error:warning:info
0:syslog0:192.135.60.4:0:on:on:on
1:newserver:192.136.70.7:4:on:off:off
```

A detailed invocation example

```
svcinfo lssyslogserver 0
```

The detailed resulting output

```
id 0
name syslog0
IP_address 192.135.60.4
```

facility 0
error on
warning on
info on

Istimezones

The **Istimezones** command lists the time zones that are available on the cluster. Each timezone is assigned an ID that can be used in the **svctask settimezone** command to set the time zone.

Syntax

```
svcinfolstimezones [-nohdr] [-delim delimiter]
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim delimiter

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by a colon character.

Description

This command displays a list of all the time zones that are available on the cluster. Each time zone is assigned an ID. This ID can be used in the **svctask settimezone** command.

An invocation example

```
svcinfolstimezones
```

The resulting output

```
id timezone
0 Africa/Abidjan
1 Africa/Accra
2 Africa/Addis_Ababa
3 Africa/Algiers
4 Africa/Asmera
5 Africa/Bamako
6 Africa/Bangui
```

lsuser

Use the lsuser command to display a list of the users that have been created on the cluster.

Syntax

```
svcinfo -- lsuser [-nohdr] [-delim delimiter] [-filtervalue attribute=value] [-filtervalue?] [userid_or_name]
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

-filtervalue *attribute=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are displayed.

Note: Some filters allow the use of a wildcard when you enter the command. The following rules apply to the use of wildcards with the SAN Volume Controller CLI:

- The wildcard character is the asterisk (*), which must be used as the first or last character in the string.
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""), as follows:

```
svcinfo lsuser -filtervalue "usergrp_name=md*"
```

-filtervalue?

(Optional) Displays the valid filter attributes for the **-filtervalue attribute=value** parameter:

- password
- ssh_key
- remote

- usergrp_id
- usergrp_name

userid_or_name

(Optional) Specifies the ID or name of the user for which the association is being deleted. If this is specified, the detailed view for the specified user is displayed in the output. If you do not specify an ID or name, the concise view is displayed.

Description

This command displays a list of users that have been created on the cluster.

An invocation example

```
svcinflsuser
```


The resulting output

id	name	password	ssh_key	remote	usergrp_id	usergrp_name
0	superuser	yes	no	no	0	SecurityAdmin
1	simon	no	yes	no	2	CopyOperator
2	jane	yes	no	no	3	Service
3	kip	yes	yes	yes		

lsusergrp

Use the lsusergrp command to display a list of the user groups that have been created on the cluster.

Syntax



Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim delimiter

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

-filtervalue *attribute=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are displayed.

Note: Some filters allow the use of a wildcard when you enter the command. The following rules apply to the use of wildcards with the SAN Volume Controller CLI:

- The wildcard character is the asterisk (*), which must be used as the first or last character in the string.
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""), as follows:

```
svcinfolusergrp -filtervalue "role=md*"
```

-filtervalue?

(Optional) Displays the valid filter attributes for the **-filtervalue** *attribute=value* parameter:

- role
- remote

usergrp_id_or_name

(Optional) Specifies the ID or name of the user group to view. If you do not specify an ID or name, all groups are displayed.

Description

This command displays a list of user groups that have been created on the cluster.

An invocation example

```
svcinfolusergrp
```

The resulting output

id	name	role	remote
0	SecurityAdmin	SecurityAdmin	yes
1	Administrator	Administrator	no
2	CopyOperator	CopyOperator	no
3	Service	Service	yes
4	Monitor	Monitor	no
5	support	Service	no

lsvdisk


The **lsvdisk** command displays a concise list or a detailed view of VDisks that are recognized by the cluster.

The list report style can be used to obtain two different styles of report.

- A list containing concise information about all the virtual disks that are recognized by the cluster. (Each entry in the list corresponds to a single virtual disk.)
- The detailed information about a single virtual disk.

Syntax

```
svcinfolusergrp lsvdisk [-filtervalue attribute=value]
```



Parameters

-filtervalue *attrib=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are displayed. If a capacity is specified, the units must also be included.

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-bytes

(Optional) Displays all capacities as bytes. Capacity values displayed in units other than bytes might be rounded. When filtering on capacity, use a unit of bytes, **-unit b**, for exact filtering.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

object_id | object_name

(Optional) Specifies the name or ID of an object. When you use this parameter, the detailed view of the specific object is returned and any value that is specified by the **-filtervalue** parameter is ignored. If you do not specify the *object_id | object_name* parameter, the concise view of all objects matching the filtering requirements that is specified by the **-filtervalue** parameter are displayed.

-filtervalue?

(Optional) Displays a list of valid filter attributes. The following filters for the **svcinfolsvdisk** command are valid:

- `vdisk_name`
- `vdisk_id`
- `vdisk_UID`
- `fc_map_count`
- `copy_count`
- `IO_group_id`
- `IO_group_name`
- `status`

- mdisk_grp_name
- mdisk_grp_id
- capacity
- type
- FC_id
- FC_name
- RC_id
- RC_name
- name
- id

Note: It is not possible to filter the `svcinfolsvdisk` command with `mdisk_grp_name=many` to identify mirrored VDIs. Instead, filter on `copy_count=2`.

Description

This command displays a concise list or a detailed view of attributes for all VDIs and VDI copies in the cluster.

The VDI 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 VDI.
- All synchronized copies for this VDI are in MDisk groups that are offline.
- The VDI is formatting.

If you have a degraded VDI and all of the associated nodes and MDIs are online, call the IBM Support Center for assistance. A VDI is reported as degraded if any of the following occurs:

- One of the nodes in the I/O group is missing.
- One of the nodes in the I/O group cannot access all the MDIs in the MDisk group that the VDI spans. In this case MDIs are shown as degraded and the DMPs for MDIs should be followed to resolve the problem.
- The fast write cache pins data for one or more VDIs in the I/O group and is unable to perform a failback until the situation is resolved. An error log indicating that the cache has pinned data is displayed. Follow the directed maintenance procedures for this error log to resolve the problem. The most common causes of pinned data are the following:
 - One or more VDIs in an I/O group is offline due to an asymmetric failure and has pinned data in the cache. Asymmetric failures can occur because of SAN fabric faults or misconfiguration, back-end controller faults or misconfiguration or because repeated errors has lead to the cluster excluding access to a MDisk through one or more nodes.
 - One or more VDIs in an I/O group is offline due to a problem with a FlashCopy mapping.

The command returns values for the following VDI attributes:

IO_groups_id/name

Specifies the I/O Group that the VDI belongs to.

status The value can be **online**, **offline** or **degraded**.

- mdisk_grp_id/name**
Specifies the name and ID of the MDisk group that the VDisk belongs to. If the VDisk has more than one copy, these fields display **many**.
- type** Specifies the virtualization type of the VDisk. The value can be **striped**, **sequential**, **image** or **many**. The value **many** indicates that the VDisk has more than one copy, which can have different virtualization types.
- capacity**
Specifies the total capacity of the VDisk.
- formatted**
Indicates whether the VDisk was formatted when it was created. The value can be **Yes** or **No**.
- mdisk_id/name**
Specifies the MDisk that is used for sequential and image mode VDIsks. If the VDisk has more than one copy, these fields display **many**.
- FC_id/name**
Specifies the name and ID of the FlashCopy mapping that the VDisk belongs to. The value **many** indicates that the VDisk belongs to more than one FlashCopy mapping.
- RC_id/name**
Specifies the name and ID of the Global Mirror or Metro Mirror relationship that the VDisk belongs to.
- vdisk_UID**
Specifies the UID of the VDisk.
- throttling**
Specifies the throttle rate of the VDisk.
- preferred_node_id**
Specifies the ID of the preferred node for the VDisk.
- fast_write_state**
Specifies the cache state for the VDisk. The value can be **empty**, **not_empty**, **corrupt**, or **repairing**. A cache state of **corrupt** indicates that the VDisk requires recovery by using one of the **recovervdisk** commands. A cache state of **repairing** indicates that repairs initiated by a **recovervdisk** command are in progress.
- cache** Specifies the cache mode of the VDisk. The value can be **readwrite** or **none**.
- udid** Specifies the unit number for the VDisk. Only OpenVMS hosts require a unit number.
- fc_map_count**
Specifies the number of FlashCopy mappings that the VDisk belongs to.
- sync_rate**
Specifies the rate for synchronization for mirrored copies.
- The command returns values for the following VDisk copy attributes:
- copy_id**
Specifies a system-assigned identifier for the VDisk copy. The value can be **0** or **1**.
- status** The value can be **online** or **offline**. A copy is offline if all nodes cannot access the MDisk group that contains the copy.

sync Indicates whether the VDisk copy is synchronized.

primary

Indicates whether the VDisk copy is the primary copy. A VDisk has exactly one primary copy. The value can be **Yes** or **No**.

mdiskgrp_id/name

Specifies the name and ID of the MDisk group that the VDisk copy belongs to.

type Specifies the virtualization type of the VDisk. The value can be **striped**, **sequential** or **image**.

mdisk_id/name

Specifies the MDisk that is used for sequential and image mode VDIsks.

fast_write_state

Specifies the cache state of the VDisk copy. The value can be **empty**, **not_empty**, **corrupt**, or **repairing**. The value is always empty for non-space-efficient copies. A cache state of **corrupt** indicates that the VDisk is space-efficient and requires repair that is initiated by a **recovervdisk** command or the **repairsevdiskcopy** command.

used_capacity

Specifies the portion of **real_capacity** that is being used to store data. For non-space-efficient copies, this value is the same as the VDisk capacity. If the VDisk copy is space-efficient, the value increases from zero to the **real_capacity** value as more of the VDisk is written to.

real_capacity

Specifies the amount of physical storage that is allocated from an MDisk group to this VDisk copy. If the VDisk copy is not space-efficient, the value is the same as the VDisk capacity. If the VDisk copy is space-efficient, the value can be different.

free_capacity

Specifies the difference between the **real_capacity** and **used_capacity** values.

overalllocation

Expressed as a percentage, specifies the ratio of VDisk capacity to **real_capacity** values. This value is always **100** for non-space-efficient VDIsks.

autoexpand

Specifies whether **autoexpand** is enabled on a space-efficient VDisk. The value can be **on** or **off**.

warning

Expressed as a percentage, for space-efficient VDisk copies only. A warning is generated when the ratio of **used_capacity** to VDisk capacity reaches the specified level.

grainsize

For space-efficient VDisk copies, specifies the grain size chosen for the VDisk copy when it was created.

A concise invocation example

```
svcinfo lsvdisk -delim :
```

The concise 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
0:vdisk0:0:io_grp0:degraded:0:mdiskgrp0:16.0MB:striped:::::60050768017F06BF780000000000000:1
1:vdisk1:0:io_grp0:degraded:0:mdiskgrp0:16.0MB:striped:::::60050768017F06BF780000000000000:1
2:vdisk2:0:io_grp0:degraded:0:mdiskgrp0:16.0MB:striped:::::60050768017F06BF780000000000000:1
3:vdisk3:0:io_grp0:degraded:0:mdiskgrp0:16.0MB:striped:::::60050768017F06BF780000000000000:1

```

A detailed invocation example

```
svcinfc lsvdisk -delim : 251
```

The detailed resulting output

```

id:251
name:i0vd163
IO_group_id:0
IO_group_name:io_grp0
status:online
mdisk_grp_id:3
mdisk_grp_name:vind1
capacity:16.0MB
type:striped
formatted:no
mdisk_id:
mdisk_name:
FC_id:
FC_name:
RC_id:
RC_name:
vdisk_UID:6005076801A0002C8000000000000078B
throttling:0
preferred_node_id:1
fast_write_state:empty
cache:readwrite
udid:
fc_map_count:0
sync_rate:50
copy_count:1

copy_id:0
status:online
sync:yes
primary:yes
mdisk_grp_id:3
mdisk_grp_name:vind1
type:striped
mdisk_id:
mdisk_name:
fast_write_state:empty
used_capacity:16.00MB
real_capacity:16.00MB
free_capacity:0.00MB
overallocation:100
autoexpand:
warning:
grainsize:

```

lsvdiskcopy

The **lsvdiskcopy** command lists VDisk copy information.

Syntax

```

▶▶ svcinfo — — lsvdiskcopy — — [ -copy copy_id ]

```



Parameters

-copy *copy_id*

(Optional) Specifies the VDisk copy to list information for. You must specify a *vdisk_name* | *vdisk_id* value with this parameter.

-filtervalue?

(Optional) Displays a list of valid filter attributes. The following filters for the **svcinfolsvdiskcopy** command are valid:

- primary
- status
- sync
- mdisk_grp_id
- mdisk_grp_name
- type

vdisk_name | *vdisk_id*

(Optional) Specifies the VDisk to list copy information for. You must specify this parameter last on the command line. If you specify a *vdisk_name* | *vdisk_id* value only, all copies for the VDisk are listed.

Description

The **lsvdiskcopy** command lists information for VDisk copies. If you specify the command with no parameters, all VDIsks and copies in the cluster are listed.

An invocation example

```
svcinfolsvdiskcopy -delim :
```

The resulting output

```
vdisk_id:vdisk_name:copy_id:status:sync:primary:mdisk_grp_id:mdisk_grp_name:
capacity:type
0:vd1:0:online:yes:yes:1:mdisk_group_1:20GB:striped
0:vd1:1:offline:no:no:2:mdisk_group_2:20GB:striped
1:vd2:0:online:yes:yes:mdisk_group_2:100GB:image
```

An invocation example

```
svcinfolsvdiskcopy -delim : vd1
```

The resulting output

```
vdisk_id:vdisk_name:copy_id:status:sync:primary:mdisk_grp_id:mdisk_grp_name:
capacity:type
0:vd1:0:online:yes:yes:1:mdisk_group_1:20GB:striped
0:vd1:1:offline:no:no:2:mdisk_group_2:20GB:striped
```

An invocation example

```
svcinfolsvdiskcopy -copy 0 -delim : vv1
```

The resulting output

```
id:0
name:vdisk0
capacity:128.0MB
copy_id:0
```

```
status:online
sync:yes
primary:yes
mdisk_grp_id:0
mdisk_grp_name:mdiskgrp0
type:striped
mdisk_id:
mdisk_name:
fast_write_state:empty
used_capacity:0.41MB
real_capacity:128.00MB
free_capacity:127.59MB
overallocation:100
autoexpand:off
warning:79grainsize:32
```

lsvdiskdependentmaps

The **lsvdiskdependentmaps** command displays all FlashCopy mappings with target virtual disks (VDisks) that are dependent upon data held on the specified VDisk.

Syntax

```
➤— svcinfo — — lsvdiskdependentmaps — [ vdisk_id | vdisk_name ] —➤
```

Parameters

vdisk_id | *vdisk_name*

(Required) Specifies the name or ID of a virtual disk (VDisk).

Description

The **lsvdiskdependentmaps** command displays FlashCopy mappings that have target VDisks that are dependent upon data held on the specified *vdisk_id* | *vdisk_name*. This can be used to determine whether a FlashCopy mapping can be prepared. Issue the command for the target VDisk *vdisk_id* | *vdisk_name* of the FlashCopy mapping to be prepared. If no FlashCopy mappings are returned, the FlashCopy mapping can be prepared. Any FlashCopy mappings that are returned in the list must be stopped or be in the **idle_or_copied** state, before the new FlashCopy mapping can be prepared.

A concise invocation example

```
svcinfo lsvdiskdependentmaps -delim : 0
```


The concise resulting output

```
id:name
2:fcmap2
5:fcmap5
```

lsvdiskextent

The **lsvdiskextent** command lists the MDisk extents that are provided for the specified VDisks.

Syntax



Parameters

-copy *copy_id*

(Optional) Displays a list of MDisks that are members of the specified VDisk copy.

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

vdisk_name | vdisk_id

(Required) Specifies one or more virtual disk IDs or names.

Description

The **lsvdiskextent** command displays a list of MDisk IDs and the number of extents that each MDisk provides to the specified VDIs.

Each VDisk is constructed from one or more MDisks. To determine the relationship between a VDisk and its MDisks, issue the following command:

```
svcinfo lsvdiskmember vdisk_name | vdisk_id
```

where *vdisk_name | vdisk_id* is the name or ID of the VDisk. This command displays a list of MDisk IDs that make up the VDisk.

To determine the number of extents that are provided by each MDisk, issue the following command:

```
svcinfo lsvdiskextent vdisk_name | vdisk_id
```

where *vdisk_name | vdisk_id* is the name or ID of the VDisk. This command displays a table of MDisk IDs and the corresponding number of extents that each MDisk provides as storage for the given VDisk.

To determine the relationship between MDisks and VDisks, issue the following command for each MDisk:

```
svcinfolsmdiskmember mdisk_name | mdisk_id
```

where *mdisk_name* | *mdisk_id* is the name or ID of the MDisk. This command displays a list of IDs that corresponds to the VDisks that are using this MDisk.

To determine the relationship between MDisks and VDisks, and the number of extents that are used by each VDisk, you must use the command-line interface. For each MDisk, issue the following command:

```
svcinfolsmdiskextent mdisk_name | mdisk_id
```

where *mdisk_name* | *mdisk_id* is the name or ID of the MDisk. This command displays a table of VDisk IDs and the corresponding number of extents that are used by each VDisk.

An invocation example

```
svcinfolsvdiskextent -delim : vdisk0
```

The resulting output

```
id:number_extents  
0:0
```

lsvdiskfcmapcopies

The `lsvdiskfcmapcopies` command displays a list of all FlashCopy mappings with a target VDisk containing a valid copy of the specified VDisk.

Syntax

```
svcinfolsvdiskfcmapcopies [-nohdr] vdisk_name vdisk_id [-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.

vdisk_name | *vdisk_id*

(Required) Specifies the name or ID of the VDisk for which the FlashCopy mappings are displayed.

Description

This command returns a list of the FlashCopy mappings that have a target VDisk with a valid copy of the specified VDisk. The target VDIsks of these mappings can be considered as candidate source VDIsks for mappings to restore from.

The mappings returned are in the copying, idle_copied, or stopping state with 100% progress.

An invocation example

```
svcinfolsvdiskfmapcopies -delim : 0
```

The resulting output

```
id:name:status:progress:difference:start_time:target_vdisk_id:
target_vdisk_name:group_id:group_name
2:fcmap2:copying:80:10:060627083137:10:vdisk10::
5:fcmap5:idle_copied:100:20:060627073130:12:vdisk12:1:fccstgrp1
```

lsvdiskfcmappings

The **lsvdiskfcmappings** command displays a list of FlashCopy mappings to which the VDisk belongs. A VDisk can be part of up to 256 FlashCopy mappings.

Syntax

```
svcinfolsvdiskfcmappings [vdisk_name | vdisk_id]
```

Parameters

vdisk_name | *vdisk_id*

(Required) Specifies the name or ID of the VDisk for which a list of all FlashCopy mappings is required.

Description

The **lsvdiskfcmappings** command returns a list of all FlashCopy mappings that the VDisk is a member of. The list is returned in no particular order.

An invocation example

```
svcinfolsvdiskfcmappings -delim : vdisk2
```

The resulting output

```
fc_id:fc_name
1:fcmap1
3:fcmap3
```

lsvdiskhostmap

Use the `lsvdiskhostmap` command to list the VDisk to the host mapping. These hosts have the specified virtual disk mapped to them; the virtual disk is visible to these hosts.

Syntax

```
svcinfolsvdiskhostmap [-nohdr] [-delim delimiter] vdisk_id | vdisk_name
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim delimiter

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

vdisk_id | vdisk_name

(Required) Specifies the ID or name of the virtual disk. The cluster displays a list of all the hosts to which this virtual disk is mapped and the SCSI ID by which the virtual disk is mapped.

Description

This command displays a list of host IDs and names. These hosts have the specified virtual disk mapped to them; that is, the virtual disk is visible to these hosts. The SCSI LUN ID is also displayed. The SCSI LUN ID is the ID by which the virtual disk is recognized by the host.

Determining the host that a VDisk is mapped to: List the hosts that this VDisk is mapped to, by issuing the following command:

```
svcinfolsvdiskhostmap vdisk_id | vdisk_name
```

where *vdisk_id | vdisk_name* is the name or ID of the VDisk. A list is displayed. Look for the host name or ID to determine which host this VDisk is mapped to. If no data is displayed, the VDisk is not mapped to any hosts.

An invocation example

```
svcinfo lsvdiskhostmap bbb
```

The resulting output

id	name	SCSI_id	host_id	host_name	vdisk_UID
200	bbb	0	9	mchost13	600507680197014B000000000000002A0

lsvdisklba

The `lsvdisklba` command lists the VDisk and logical block address (LBA) for the specified MDisk LBA.

Syntax

```
svcinfo -- lsvdisklba -- -mdisklba -- mdisklba --  
-delim -- delimiter - nohdr -mdisk [mdisk_id | mdisk_name]
```

Parameters

-mdisklba *mdisklba*

(Required) Specifies the 64-bit hexadecimal LBA on the MDisk.

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

-mdisk *mdisk_id* | *mdisk_name*

(Required) Specifies the MDisk name or ID.

Description

The `lsvdisklba` command returns the LBA of the VDisk that is associated with the MDisk LBA.

If applicable, the command also lists the range of LBAs on both the VDisk and MDisk that are mapped in the same extent, or for space-efficient disks, in the same grain.

Table 20 on page 314 provides command output that depends on several variables.

Table 20. *lsvdisklba* command output scenarios

Field	Typical scenario	Quorum disk	Space-efficient metadata	Extent not allocated	Formatting extent	Extent allocated to space-efficient disk, LBA not used on space-efficient disk
<code>copy_id</code>	yes	no	yes	no	yes	yes
<code>vdisk_id</code>	yes	no	yes	no	yes	yes
<code>vdisk_name</code>	yes	no	yes	no	yes	yes
<code>type</code>	allocated	metadata	metadata	unallocated	formatting	unallocated
<code>vdisk_lba</code>	yes	no	no	no	no	no
<code>vdisk_start</code>	yes	no	no	no	no	no
<code>vdisk_end</code>	yes	no	no	no	no	no
<code>mdisk_start</code>	yes	yes	yes	yes	yes	yes
<code>mdisk_end</code>	yes	yes	yes	yes	yes	yes

An invocation example

```
svcinfo lsvdisklba -mdisk 1 -mdisklba 0x0
```

The resulting output

```
vdisk_id
vdisk_name
copy_id
type metadata
vdisk_lba 0x00090000
vdisk_start
vdisk_end
mdisk_start 0x00000000
mdisk_end 0x0000FFFF
```

lsvdiskmember

The **lsvdiskmember** command displays a list of MDisks that are members of the specified VDisk.

Syntax

```

▶▶ svcinfo — — lsvdiskmember — [ -copy — copy_id ] [ -nohdr ]
▶ [ -delim — delimiter ] [ vdisk_id ] [ vdisk_name ]

```

Parameters

-copy *copy_id*

(Optional) Displays a list of MDisks that are members of the specified VDisk copy.

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

vdisk_id | *vdisk_name*

(Required) Displays a list of MDisks that are members of the specified VDisk.

Description

This command displays a list of managed disks, which provide extents that make up the virtual disk that is specified by the ID.

Every VDisk is constructed from one or more MDisks. At times, you might have to determine the relationship between the two objects. The following procedure allows you to determine the relationships.

If you use the **svcinfo lsmdiskmember** command, the concise view displays a list of virtual disks. These are the virtual disks that are using extents on the managed disk that is specified by the ID. The list displays the members of the respective object and is independent of the state of the individual members; that is, if they are in offline state, they are still displayed.

To determine the relationship between VDIsks and MDisks, issue the following command:

```
svcinfo lsvdiskmember vdisk_id | vdisk_name
```

where *vdisk_id* | *vdisk_name* is the name or ID of the VDisk. This displays a list of IDs that correspond to the MDisks that make up the VDisk.

To determine the relationship between VDIsks and MDisks, and the number of extents that are provided by each MDisk, you must use the command-line interface. Issue the following command:

```
svcinfo lsvdiskextent vdisk_id | vdisk_name
```

where *vdisk_id* | *vdisk_name* is the name or ID of the VDisk. This displays a table of MDisk IDs and the corresponding number of extents that each MDisk provides as storage for the specified VDisk.

To determine the relationship between MDisks and VDIsks, issue the following command:

```
svcinfo 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 VDIsks that are using this MDisk.

To determine the relationship between MDisks and VDIs, and the number of extents that are used by each VDisk, you must use the command-line interface. For a specified MDisk, issue the following command:

```
svcinfo lsmdiskextent mdisk_id | mdisk_name
```

where *mdisk_id* | *mdisk_name* is the name or ID of the MDisk. This displays a table of VDisk IDs and the corresponding number of extents that are used by each VDisk.

An invocation example

```
svcinfo lsvdiskmember 1
```

The resulting output

```
id
2
```

lsvdiskprogress

The **lsvdiskprogress** command tracks the progress during new virtual disk formatting.

Syntax

```

>>> svcinfo -- lsvdiskprogress -- [ -nohdr ] ----->
|
| [ -delim delimiter ] [ vdisk_id | vdisk_name ]
|
|----->>>

```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by a colon character.

vdisk_id | *vdisk_name*

(Optional) Specifies the VDisk ID or name. If you do not specify this parameter, the progress of all VDIs currently being formatted is displayed.

Description

This command displays the progress of the format of a new virtual disk as a completed percentage. If the VDisk has multiple copies, the command reports the average progress of the format.

An invocation example

```
svcinfo lsvdiskprogress -delim : 0
```

The resulting output

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

lsvdisksyncprogress

The **lsvdisksyncprogress** command displays the progress of VDisk copy synchronization.

Syntax

```
▶▶—svcinfo— —lsvdisksyncprogress— —  
└─ -copy — id ─┘ ┌─ vdisk_name ─┘  
└─ vdisk_id ─┘
```

Parameters

-copy *id*

(Optional) Specifies the VDisk copy ID to list synchronization progress for. You must also specify a *vdisk_name* | *vdisk_id* value. If you do not specify this parameter, progress is displayed for all copies.

vdisk_name | *vdisk_id*

(Optional) Specifies the virtual disk name or ID to list synchronization progress for.

Description

The **lsvdisksyncprogress** command displays the progress of VDisk copy synchronization. To display the VDisk copies that require synchronization, specify the command with no parameters. Estimated completion time is in the YYMMDDHHMMSS format. The command displays progress for the following types of VDIsks:

- A synchronized copy displays a progress of 100 and a blank estimated completion time.
- An offline copy displays a blank estimated completion time, and a gradually decreasing progress if the VDisk is being written to.
- Nonmirrored VDIsks are displayed as a single copy with a progress of 100, and a blank estimated completion time.

The **lsvdisksyncprogress** command also displays the progress of a mirrored VDisk synchronization. After you create a mirrored VDisk using the **svctask mkvdisk** or **svctask addvdiskcopy** command, you can use the command to monitor the progress of the synchronization.

An invocation example

```
svcinfolsvdisksyncprogress
```

The resulting output

```
vdisk_id vdisk_name copy_id progress estimated_completion_time
0 vdisk0 0 100
0 vdisk0 1 50 070301150000
3 vdisk3 0 72 070301132225
3 vdisk3 1 100
4 vdisk4 0 22 070301160000
4 vdisk4 1 100
8 vdisk8 0 100
8 vdisk8 1 33
```

An invocation example

```
svcinfolsvdisksyncprogress vdisk0
```

The resulting output

```
vdisk_id vdisk_name copy_id progress estimated_completion_time
0 vdisk0 0 100
0 vdisk0 1 50 070301150000
```

showtimezone

Use the **showtimezone** command to display the current time zone settings for the cluster.

Syntax

```
svcinfolshowtimezone [-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 character.

Description

This command displays a single time zone and its associated ID. This is the current time zone setting for the cluster. A list of available time-zones can be viewed by running the `svcinfolstimezones` command. The time zone can be changed by running the `svctask settimezone` command.

An invocation example

```
svcinfolstimezone -delim :
```

The resulting output

```
id:timezone  
522:UTC
```

Chapter 22. Error log commands (Discontinued)

| **Attention:** The svcservicetask error log commands are discontinued. Use the
| svctask cluster diagnostic and service-aid commands to diagnose and find cluster
| problems.

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.

Chapter 23. Licensing commands

The following commands enable you to work with SAN Volume Controller licensed functions.

chlicense

The **chlicense** command changes license settings for cluster features.

Syntax

```
svctask -- chlicense -- -flash capacity_TB --
                        -remote capacity_TB --
                        -virtualization capacity_TB --
                        -physical_flash on | off
                        -physical_remote on | off
                        -physical_disks number
```

Parameters

-flash *capacity_TB*

(Optional) Changes cluster licensing for the FlashCopy feature. To change the licensed capacity for the FlashCopy feature, specify a capacity in terabytes (TB).

-remote *capacity_TB*

(Optional) Changes cluster licensing for the Metro Mirror and Global Mirror feature. To change the licensed capacity for the Metro Mirror and Global Mirror feature, specify a capacity in terabytes (TB).

-virtualization *capacity_TB*

(Optional) Changes cluster licensing for the Virtualization feature. To change the licensed capacity for the Virtualization feature, specify a capacity in terabytes (TB).

-physical_flash **on** | **off**

(Optional) For physical disk licensing, enables or disables the FlashCopy feature. The default value is **off**.

-physical_remote **on** | **off**

(Optional) For physical disk licensing, enables or disables the Metro Mirror and Global Mirror feature. The default value is **off**.

-physical_disks *number*

(Optional) Changes the licensed settings of the cluster for physical disk licensing. Enter the number of physical disks that your cluster is licensed to manage. The correct number is identified in your license.

Note:

- If the **-physical_disks** value is set to zero, the **-physical_flash** and **-physical_remote** values are turned off.
- If the **-physical_disks** value is nonzero, the **-flash**, **-remote**, and **-virtualization** values cannot be set.

- If the **-flash**, **-remote**, or **-virtualization** values are nonzero, the **-physical_flash**, **-physical_remote**, and **-physical_disks** values cannot be set.
- If the **-physical_disks** value is nonzero, only the FlashCopy and RemoteCopy usage is monitored and appropriate error messages are logged.

Description

The `chlicense` command changes license settings for the cluster. Any change that is made is logged as an event in the license setting log.

SAN Volume Controller version 4.3.1 provides two license options: physical disk licensing and capacity licensing. To select physical disk licensing, run the `svctask chlicense` command with the **-physical_disks** parameter. To select capacity licensing, run the `svctask chlicense` command with the **-flash**, **-remote**, or **-virtualization** parameter.

The current license settings for the cluster are displayed in the viewing license settings log panel. These settings show which features you are licensed to use. They also show the storage capacity that is licensed for virtualization. Typically, the license settings log contains entries because feature options must be set as part of the Web-based cluster creation process.

Note: Dumping an empty license settings log produces a file that contains headers, 256 lines of formatted zeros, and two lines that include a checksum operation.

By default, the Copy Services functions are disabled, but this does not stop you from creating and using copy services. However, errors are placed in the license settings log that state that you are using an unlicensed feature. The command-line tool return code also notifies you that you are using an unlicensed feature.

The total virtualized capacity can also be modified with this command. This is the number of terabytes (TB) of virtual disk capacity that can be configured by the cluster.

When you reach 90% capacity, any attempt to create or extend Virtual Disks, Relationships, or Mappings results in a message from the command-line tool. This does not stop you from creating and expanding Virtual Disks, Relationships, or Mappings. When usage reaches or exceeds 100% capacity, errors are placed in the license settings log.

Any error that is placed in the license settings log results in a generic error being placed in the cluster error log. This occurs when you issue a command that violates the license agreement. The return code also notifies you that you are violating the license settings.

An invocation example

```
svctask chlicense -flash 5
```

The resulting output

No feedback

dumpinternallog

The **dumpinternallog** command dumps the contents of the license settings error and event log to a file on the current configuration node.

Syntax

▶— svctask — — dumpinternallog —————▶

Description

This command dumps the contents of the internal license settings error and event log to a file on the current configuration node.

This file is always called **feature.txt** and is created, or overwritten, in the **/dumps/feature** directory on the configuration node.

This file can be requested by IBM service personnel.

Before making any entries, the license settings log contains only zeros. A dump of this log from the **svctask dumpinternallog** command results in an empty file.

An invocation example

```
svctask dumpinternallog
```

The resulting output

No feedback

Chapter 24. Secure Shell key commands (Discontinued)

| **Attention:** The Secure Shell (SSH) key commands are discontinued. Use the user
| management commands to configure remote authentication service and manage
| users and user groups on the cluster.

addsshkey (Discontinued)

| **Attention:** The addsshkey command is discontinued. Use the user management
| commands to configure remote authentication service and manage users and user
| groups on the cluster.

rmallsshkeys (Discontinued)

| **Attention:** The rmallsshkeys command is discontinued. Use the user management
| commands to configure remote authentication service and manage users and user
| groups on the cluster.

rmsshkey (Discontinued)

| **Attention:** The rmsshkey command is discontinued. Use the user management
| commands to configure remote authentication service and manage users and user
| groups on the cluster.

Chapter 25. Service mode commands

Service mode commands perform tasks when the node is in service mode. Some of these tasks are to specify node software, to clean dump directories and to dump the contents of an error log to a file.

When a node is in service mode it is no longer running as part of the cluster. Any information related to the existing cluster state and other nodes is from the last time this node was in the cluster.

applysoftware

The applysoftware command upgrades the node to a new level of software.

Syntax

```
►► — svcservicemodetask — — applysoftware — — — -file — filename_arg — —►►  
└─┬─ -ignore ─┘
```

Parameters

-file *filename_arg*

(Required) Specifies the filename of the new software package.

-ignore

(Optional) Specifies that all prerequisite checking be bypassed and that all hardened data be deleted. Use this parameter with caution.

Description

This command starts the upgrade process of the node to a new level of software. The applysoftware command can be used to apply a level of software to the node in both service and nonservice modes. In service mode the applysoftware command is applied to the specific node in service mode. In nonservice mode, the command is applied to the complete cluster.

The software package as specified by the file name must first be copied on to the current configuration node/node in service mode in the /home/admin/upgrade directory. You can use the PuTTY secure copy (scp) application to copy the file.

The command completes as soon as the upgrade process has successfully begun. The command fails and the upgrade package is deleted if:

- The given package fails an integrity check due to corruption.
- Any node in the cluster has a hardware type not supported by the new software.
- The new software level does not support upgrades from the old software.
- The software level of a remote cluster is incompatible with the new software.

The actual upgrade completes asynchronously.

You can view the contents of the /home/admin/upgrade directory by using the svcinfo/svc servicemodeinfo lssoftware dumps command.

An invocation example

```
svcservicemodetask applysoftware -file newssoftware
```

The resulting output

No feedback

cleardumps

The **cleardumps** command cleans the various dump directories on the node that is in service mode.

Syntax

```
▶▶ svcservicemodetask — —cleardumps— —————▶
```

```
▶ -prefix — directory_or_file_filter —————▶▶
```

Parameters

-prefix *directory_or_file_filter*

(Required) Specifies the directory, files, or both to be cleaned. If a directory is specified, with no file filter, all relevant dump or log files in that directory are cleaned. You can use the following directory arguments (filters):

- /dumps (cleans all files in all subdirectories)
- /dumps/configs
- /dumps/elogs
- /dumps/feature
- /dumps/iostats
- /dumps/iotrace
- /dumps/mdisk
- /home/admin/upgrade

In addition to the directory, you can specify a file filter. For example, if you specify /dumps/elogs/*.txt, all files in the **/dumps/elogs** directory that end in .txt are cleaned.

Note: The following rules apply to the use of wildcards with the SAN Volume Controller CLI:

- The wildcard character is an asterisk (*).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""), as follows:

```
svcservicemodetask cleardumps -prefix "/dumps/elogs/*.txt"
```

Description

This command deletes all the files that match *directory_or_file_filter* value for the node that is in service mode.

You can clean all the dumps directories by specifying **/dumps** as the directory value.

You can clean all the files in a single directory by specifying one of the directory values.

You can list the contents of these directories on the given node by using the **svcservicemodeinfo lsxxxxdumps** commands.

You can use this command to clean specific files in a given directory by specifying a directory or file name. You can use the wildcard (*) as part of your file name.

Note: To preserve the configuration and trace files, any files that match the following wildcard patterns are not cleaned:

- *svc.config*
- *.trc
- *.trc.old

An invocation example

```
svcservicemodetask cleardumps -prefix /dumps/configs
```

The resulting output

No feedback

dumperrlog

The dumperrlog command dumps the contents of the error log to a text file.

Syntax

```
▶▶— svcservicemodetask — — dumperrlog — — [ -prefix — filename_prefix ] ▶▶
```

Parameters

-prefix *filename_prefix*

(Optional) Creates a file name 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

With no parameters, the svcservicemodetask dumperrlog command dumps the cluster error log to a file using a system-supplied prefix of **errlog**, which includes the node ID and time stamp. This information is from when the node was last a part of a cluster. When a file name prefix is provided, the same operation is performed, but the details are stored in the dumps directory within a file with a name that starts with the specified prefix.

A maximum of ten error log dump files are kept on the node. When the 11th dump is made, the oldest existing dump file is overwritten.

Error log dump files are written to the /dumps/elog directory. You can view the contents of this directory using the svcinfo lserrlogdumps command.

Files are not deleted from other nodes until you issue the cleardumps command.

An invocation example

```
svcservicemodetask dumperrlog -prefix testerrorlog
```

The resulting output

No feedback

exit

The **exit** command exits service mode and restarts the node.

Syntax

```
▶▶ svcservicemodetask — — exit —————▶▶
```

Description

This command causes the node to be restarted. The node powers on in standard operating mode and attempts to rejoin the cluster.

At some point during the processing of this command, the SSH (secure shell) client software and the Web server connection that accessed the node are ended as a result of the restart processing.

An invocation example

```
svcservicemodetask exit
```

The resulting output

```
[SSH / webservice connections terminate so an error message to the effect of  
'connection lost' may be displayed, or 'CLIENT RECEIVED SERVER DOWN  
NOTIFICATION']
```

Chapter 26. Service mode information commands

Service mode information commands perform information gathering tasks when the node is in service mode.

| When a node is in service mode it is no longer running as part of the cluster. Any
| information related to the existing cluster state and other nodes is from the last
| time this node was in the cluster.

These commands can only be run on a node that is in service mode. Attempting to run them on a working configuration node will result in the message:

| CMMVC6002E This command can only be run on a node
| that is in service mode.

Attempting to run any of the other `svcin` commands on a node that is in service mode will result in the following message:

| CMMVC6003E This command can not be run on a node
| that is in service mode.

ls2145dumps

The `ls2145dumps` command lists node assert dumps and associated output files that are located in the `/dumps` directory.

Syntax

```
▶▶▶ svc servicemodeinfo -- ls2145dumps -- [ -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*

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 a colon character.

Description

| This command lists node assert dumps and associated output files, which are
| created during an assertion of a node. The command displays the files that are
| located in the **/dumps** directory.

An invocation example

```
svcservicemodeinfo ls2145dumps -delim :
```

The resulting output

```
id:2145_filename  
0:000108.trc.old  
1:dump.000108.030328.144007  
2.000108.trc
```

lscimomdumps

The **lscimomdumps** command lists the dump files in the **/dumps/cimom** directory.

Syntax

```
▶▶ svcservicemodeinfo — — lscimomdumps — — [ -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*

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; 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 a colon character.

Description

| This command displays a list of Common Information Model object manager
| (CIMOM) dumps. These dumps are created when you use a CIM client with the
| CIMOM of the SAN Volume Controller cluster. The command displays the files
| that are located in the **/dumps/cimom** directory.

An invocation example

```
svcservicemodeinfo lscimomdumps
```


The resulting output

id	cimom_filename
0	mkrepositorylog.004565
1	PegasusTrace.004565
2	PegasusStandard.004565
3	PegasusAudit.004565
4	PegasusError.004565
5	PegasusDebug.004565

lsclustervpd

The **lsclustervpd** command returns the vital product data (VPD) for the cluster to which the node belonged.

Syntax

```
svcservicemodeinfo -- lsclustervpd [-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*

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 a colon character.

Description

This command displays the VPD of the cluster to which the node belonged.

An invocation example

```
svcservicemodeinfo lsclustervpd
```

The resulting output

```
| id 000001002FF007E5  
| name cluster1  
| location local  
| partnership  
| bandwidth 0  
| cluster_IP_address 0.0.0.0  
| cluster_service_IP_address 1.1.1.1  
| total_mdisk_capacity 0
```

```

space_in_mdisk_grps 0
space_allocated_to_vdisks 0
total_free_space 0
statistics_status off
statistics_frequency 15
required_memory 2048
cluster_locale en_US
SNMP_setting all
SNMP_community
SNMP_server_IP_address 0.0.0.0
subnet_mask 0.0.0.0
default_gateway 0.0.0.0
time_zone 522 UTC
email_setting all
email_id
code_level 00000000
FC_port_speed 1Gb
console_IP
id_alias
gm_link_tolerance:300
gm_inter_cluster_delay_simulation:0
gm_intra_cluster_delay_simulation:0

```

lserrlogdumps

The **lserrlogdumps** command lists the error log dump files in the `/dumps/elog` directory, which are created when you run the **svctask dumperrlog** command.

Syntax

```

▶▶▶ svcservicemodeinfo — — lserrlogdumps — — [ -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 even if the **-nohdr** parameter is specified.

-delim *delimiter*

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 a colon character.

Description

This command lists error log dump files, which are created when you run the **svctask dumperrlog** command. An error log dump describes the contents of the error log at the time that the command was run. The command displays the files that are located in the **/dumps/elogs** directory.

An invocation example

```
svcserviceinfo lserrlogdumps
```

The resulting output

id	filename
0	errlog_lynn02_030327_154511
1	aaa.txt_lynn02_030327_154527
2	aaa.txt_lynn02_030327_154559
3	errlog_lynn02_030403_110628

lsfeaturedumps

The **lsfeaturedumps** command lists the dump files in the **/dumps/feature** directory, which are created when you run the **svctask dumpinternallog** command.

Syntax

```
svcserviceinfo -- lsfeaturedumps [-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*

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.

Description

This command returns a list of featurization dumps, which are created when you run the **svctask dumpinternallog** command. A featurization dump file describes the contents of the featurization log at the time that the command was executed. The command displays files from the **/dumps/feature** directory.

Issue the **svcinfol sfeaturedumps** command to list the dump files in the **/dumps/feature** 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.

An invocation example

```
svcserviceinfo lsfeaturedumps
```

The resulting output

```
id          feature_filename
0          feature.txt
```

lsiostatsdumps

The **lsiostatsdumps** command lists the dump files in the **/dumps/iostats** directory, which are created when you run the **svctask startstats** command.

Syntax

```
svcserviceinfo -- lsiostatsdumps -- [-nohdr] [-delim delimiter] [node_id | node_name]
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one byte character. If you enter **-delim :** on the command line, the colon character (**:**) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by a colon character.

node_id | node_name

(Optional) Specifies the node ID or name to list the available dumps of the given type. If you do not specify a node, the dumps that are available on the configuration node are listed.

Description

This command displays a list of I/O statistics dumps. These dumps are created when you issue the **svctask startstats** command. If you do not specify a node, the dumps that are available on the configuration node are listed. The command displays files from the **/dumps/iostats** directory.

An invocation example

```
svcserviceinfo lsiostatsdumps
```

The resulting output

```
| id          iostat_filename  
| 0          Nv_stats_mala75_031123_072426  
| 1          Nm_stats_mala75_031123_072425  
| 2          Nn_stats_mala75_031123_072424
```

lsiotracedumps

The **lsiotracedumps** command lists the files that are located in the **/dumps/iotrace** directory.

Syntax

```
▶▶ svcserviceinfo — — lsiotracedumps — — [ -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

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.

Description

```
| This command lists I/O trace dumps, which are created when you run the svctask  
| settrace command. The command displays the files that are located in the  
| /dumps/iotrace directory.
```

An invocation example

```
svcserviceinfo lsiotracedumps
```

The resulting output

id	iotrace_filename
0	c1_mala75_030405_092155
1	c2_mala75_030405_092156
2	c3_mala75_030405_092158
3	c4_mala75_030405_092159
4	c5_mala75_030405_092201

lsmdiskdumps

The `lsmdiskdumps` command displays a list of dumps in the `/dumps/mdisk` directory.

Syntax

```

svcservicemodeinfo -- lsmdiskdumps [-nohdr]
                                [-delim delimiter]
                                [node_id | node_name]

```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

node_id | *node_name*

(Optional) Specifies the node ID or name to list the available dumps of the given type. If you do not specify a node, the dumps that are available on the configuration node are listed.

Description

This command displays a list of managed disk (MDisk) error data. If you do not specify a node, the dumps that are available on the configuration node are listed. The command displays files from the `/dumps/mdisk` directory.

An invocation example

```
svcservicemodeinfo lsmdiskdumps
```

The resulting output

```
id          mdisk_filename
0          mdiskdump_desty5_2_090625_215452
```

lsnodevpd

The `lsnodevpd` command returns the vital product data (VPD) for the specified node.

Syntax

```
svcserviceinfo -- lsnodevpd [-nohdr]
                    [-delim delimiter] [node_id | node_name]
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim delimiter

(Optional) By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by a colon character.

node_id | node_name

(Required) Specifies the node to view in terms of its node ID or name.

Description

This command returns the VPD for the specified node. Each field is reported on a new line. All fields are strings. The VPD is split into sections, each with a section heading. The number of fields in the section follows the 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
thanobject2 field2:value

new section: x fields
...

Note: For 8F4, 8G4, and 8A4 nodes, the VPD displays the device serial number of the FC card as N/A.

An invocation example

```
svcinfo lsnodevpd 1
```

The resulting output

```
id 1  
  
system board: 21 fields  
part_number 43V7072  
system_serial_number KD1438A  
number_of_processors 4  
number_of_memory_modules 6  
number_of_fans 6  
number_of_FC_cards 1  
number_of_scsi/ide_devices 2  
BIOS_manufacturer IBM Corp.  
BIOS_version -[D6E124AUS-1.01]-  
BIOS_release_date 04/30/2009  
system_manufacturer IBM  
system_product IBM System x -[2145CF8]-  
version 00  
planar_manufacturer IBM  
planar_product 49Y6498  
planar_version (none)  
power_supply_part_number 39Y7201  
CMOS_battery_part_number 33F8354  
frame_assembly_part_number  
ethernet_cable_part_number  
service_processor_firmware 1.01  
  
processor: 6 fields  
processor_location Processor 1  
manufacturer Intel(R) Corporation  
version Intel(R) Xeon(R) CPU E5530 @ 2.40GHz  
speed 2400  
status Enabled  
CPU_part_number 46D1266  
  
memory module: 96 fields  
part_number 44T1493  
device_location DIMM01  
bank_location BANK01  
size (MB) No Module Installed  
manufacturer Not Specified  
serial_number Not Specified  
  
part_number 44T1493  
device_location DIMM02  
bank_location BANK02  
size (MB) 4096  
manufacturer Samsung
```



```

serial_number 99062848

part_number 44T1493
device_location DIMM03
bank_location BANK03
size (MB) 4096
manufacturer Samsung
serial_number C7062848
...

fan: 12 fields
part_number 43V6929
location location1

part_number 43V6929
location location2

part_number 43V6929
location location3
...


Adapter card: 18 fields
card_type FC card
part_number 31P1337
port_numbers 1 2 3 4
location 0
device_serial_number 11S31P1333YM10MY96A206
manufacturer IBM
device QE8
card_revision 2
chip_revision 2.0

card_type SAS card
part_number 44E8690
port_numbers 1 2 3 4
location 0
device_serial_number 11S31P1299YM10MY948004
manufacturer IBMHUR
device Capri-PMC8001
card_revision Y
chip_revision 1.1

Fibre channel SFP: 48 fields
part_number 17P9211
manufacturer JDSU
device PLRXPLVCSH4921
serial_number C915EB06V
supported_speeds 2,4,8
connector_type LC
transmitter_type SN
wavelength 850
max_distance_by_cable_type OM1:20,OM2:50,OM3:150
hw_revision 1
port_number 1
WWPN 500507680140350d
...

device: 15 fields
part_number 31P1339
bus USB
device 0
model IBM USB Endeavour
revision 1.0
serial_number NA
approx_capacity 0
hw_revision 0

```

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*

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 a colon character.

Description

|
|
|

This command lists software upgrade packages, which are copied as a result of software upgrades. The command displays files that are located in the **/home/admin/upgrade** directory.

An invocation example

```
svcserviceinfo lsoftwareinfo
```

The resulting output

```
id          software_filename
0           s1_mala75_030405_092143
1           s2_mala75_030405_092145
2           s3_mala75_030405_092146
```

Chapter 27. Controller command

The controller command modifies the name of a storage controller.

chcontroller

The chcontroller command modifies the attributes of a controller.

Syntax

```
svctask -- chcontroller -- [-name new_name]
                             [-allowquorum yes | no] controller_id | controller_name
```

Parameters

-name *new_name*

(Optional) Specifies the new name to be assigned to the controller.

-allowquorum *yes* | *no*

(Optional) Specifies that the controller is allowed or is not allowed to support quorum disks. A value of **yes** enables a suitable controller to support quorum disks. A value of **no** disables a controller from supporting quorum disks, provided that the specified controller is not currently hosting a quorum disk.

controller_id | *controller_name*

(Required) Specifies the controller to modify; use either the controller name or the controller ID.

Description

This command changes the name of the controller that is specified by the *controller_id* | *controller_name* variable to the value that you specify with the **-name** parameter.

Use the optional **-allowquorum** parameter to turn quorum on or off for the specified controller.

Note: If any controller that is associated with an MDisk shows the **allow_quorum** attribute set to **no** with the `svcinfo lscontroller` command, the set quorum action fails for that MDisk. Before using the `svctask chcontroller` command to set the **-allowquorum** parameter to **yes** on any disk controller, check the following Web site to see whether the controller supports quorum.

www.ibm.com/storage/support/2145

You can add a new disk controller system to your SAN at any time. Follow the switch zoning guidelines in the section about switch zoning. Also, ensure that the controller is set up correctly for use with the cluster.

You can create one or more arrays on the controller using RAID-5, RAID-1, or RAID-0+1 (called RAID-10) for maximum redundancy and reliability. If your controller provides array partitioning, you can create a single partition from the entire capacity that is available in the array. Record the LUN number that you assign to each partition. If your disk controller system requires LUN mapping, follow the mapping guidelines to map the partitions or the arrays to the cluster ports.

To add a new disk controller system to a running configuration, ensure that the cluster has detected the new storage MDisks by issuing the **svctask detectmdisk** command. The controller has automatically been assigned a default name. If you are unsure of which controller is presenting the MDisks, issue the **svcinfo 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 unique name by issuing the following command:

```
svctask chcontroller -name newname oldname
```

List the unmanaged MDisks by issuing the following command:

```
svcinfo 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 managed disk group and add only the RAID arrays that belong to the new controller to this MDisk group. Avoid mixing RAID types; for each set of RAID array types (for example, RAID-5 or RAID-1), create a new MDisk group. Assign this MDisk group an appropriate name; if your controller is called FAST650-abc and the MDisk group contains RAID-5 arrays, assign the MDisk a name similar to **F600-abc-R5**. Issue the following command:

```
svctask mkmdiskgrp -ext 16 -name mdisk_grp_name  
-mdisk colon-separated list of RAID-x mdisks returned
```

Note: This creates a new MDisk group with an extent size of 16 MB.

An invocation example

```
svctask chcontroller -name newtwo 2
```

The resulting output

No feedback

Chapter 28. 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.

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

Specify the required parameter, and resubmit the command.

CMMVC5721E *VALUE is not a valid time stamp format. The valid format is MMDDHHmmYYYY.*

Explanation

The specified value is not a valid time-stamp format. The valid format is MMDDHHmmYYYY.

Action

Follow the correct time-stamp format, and resubmit the command.

CMMVC5722E *VALUE is not a valid month.*

Explanation

The specified value is not a valid month.

Action

Specify the correct month (MM), and resubmit the command.

CMMVC5723E *VALUE is not a valid day.*

Explanation

The specified value is not a valid day.

Action

Specify the correct day (DD), and resubmit the command.

CMMVC5724E *VALUE is not a valid hour.*

Explanation

The specified value is not a valid hour.

Action

Specify the correct hour (HH), and resubmit the command.

CMMVC5725E *VALUE is not a valid minute.*

Explanation

The specified value is not a valid minute.

Action

Specify the correct minute (mm), and resubmit the command.

CMMVC5726E *VALUE* are not valid seconds.**Explanation**

The specified value are not valid seconds.

Action

Specify the correct seconds (ss), and resubmit the command.

CMMVC5727E *VALUE* is not a valid filter.**Explanation**

You can filter the output of some views by using the `-filtervalue` parameter. The specified string that you have entered is not a supported value for the `-filtervalue` parameter in this view.

Action

Ensure that you use a supported value for the `-filtervalue` parameter, and resubmit the command.

CMMVC5728E *VALUE* should be in the format `minute:hour:day:month:weekday`.**Explanation**

The specified value should be in the format `minute:hour:day:month:weekday`.

Action

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.

Action

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.

Action

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.

Action

Ensure that you enter a supported combination of parameters and values, and resubmit the command.

CMMVC5732E The shared-memory interface is not available, return code *RETURN_CODE*.

Explanation

You cannot submit CLI commands because the shared memory interface (SMI) is not available.

Action

Ensure that your connection to the cluster is functioning properly. If the problem persists, contact IBM technical support for assistance and report that you have received the specified message and return code.

CMMVC5733E Enter at least one parameter.

Explanation

You must specify at least one parameter for the command that you have submitted.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

Not applicable.

CMMVC5751E The dump file could not be written to.

Explanation

The dump file could not be written to disk.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

Not applicable.

CMMVC5758E Object name already exists.

Explanation

The object name already exists.

Action

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.

Action

Not applicable.

CMMVC5760E Failed to add the node to the cluster member list.**Explanation**

The node could not be added to the cluster.

Action

Not applicable.

CMMVC5761E Failed to delete the node from the cluster member list.**Explanation**

The node could not be deleted from the cluster.

Action

Not applicable.

CMMVC5762E The request did not complete before the timeout period expired.**Explanation**

The operation failed because the timeout period expired.

Action

Resubmit the command.

CMMVC5763E The node failed to go online.**Explanation**

The node failed to go online.

Action

Not applicable.

CMMVC5764E The mode change request is invalid - internal error**Explanation**

The specified mode change is not valid.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

Specify the -force flag to complete the operation, and resubmit the command.

CMMVC5772E The operation requested could not be performed because software upgrade is in progress.

Explanation

The operation failed because a software upgrade is in progress.

Action

Wait for the software upgrade to complete, and resubmit the command.

CMMVC5773E The object selected is in the wrong mode to perform the requested operation.

Explanation

The operation failed because the selected object is in the wrong mode.

Action

Specify the correct mode, and resubmit the command.

CMMVC5774E The userid supplied is not valid.

Explanation

The userid is not valid.

Action

Specify a different userid, and resubmit the command.

CMMVC5775E The directory attribute specified is not valid.

Explanation

The directory attribute is not valid.

Action

Specify a different directory, and resubmit the command.

CMMVC5776E The directory listing could not be retrieved.

Explanation

The directory listing could not be retrieved.

Action

Specify a different directory listing, and resubmit the command.

CMMVC5777E The node could not be added to the I/O Group, because the other node in the I/O 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.

Action

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.

Action

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.

Action

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.

Action

Specify a different cluster ID, and resubmit the command.

CMMVC5782E The object specified is offline.

Explanation

The object is offline.

Action

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 service mode.

Action

None.

CMMVC5784E The cluster name specified is not unique, specify the cluster using the cluster ID.

Explanation

The cluster name is not unique.

Action

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.

Action

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.

Action

Not applicable.

CMMVC5787E The cluster was not created because a cluster already exists.

Explanation

The cluster was not created because a cluster already exists.

Action

Not applicable.

CMMVC5788E The service IP address is not valid.

Explanation

The service IP address is not valid.

Action

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.

Action

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.

Action

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.

Action

Specify the correct entity, and resubmit the command.

CMMVC5792E The action failed because the I/O group is used for recovery.

Explanation

The action failed because the I/O group is used for recovery.

Action

Not applicable.

CMMVC5793E The node was not added to the cluster because the I/O group already contains a pair of nodes.

Explanation

The node was not added to the cluster because the I/O group already contains a pair of nodes.

Action

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.

Action

Specify a node that is contained in the cluster, and resubmit the command.

CMMVC5795E The node was not deleted because a software upgrade is in progress.

Explanation

The node was not deleted because a software upgrade is in progress.

Action

Wait for the software upgrade to complete, and resubmit the command.

CMMVC5796E The action failed because the I/O group that the node belongs to is unstable.

Explanation

A previous configuration command might not yet have completed.

Action

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 VDisks associated with the I/O group, therefore the node could not be deleted.

Action

Not applicable.

CMMVC5798E The action failed because the node is offline.

Explanation

The action failed because the node is offline.

Action

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.

Action

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.

Action

Specify a different entity, and resubmit the command.

CMMVC5801E The upgrade of the cluster software could not proceed because every node in the cluster must be online. Either delete the node that is offline or bring the node online and resubmit the command

Explanation

The upgrade of the cluster software could not proceed because every node in the cluster must be online.

Action

Either delete the node that is offline or bring the node online, and resubmit the command.

CMMVC5802E The upgrade of the cluster software could not proceed because there is an I/O group in the cluster that contains only one node. The software upgrade requires that each node in an I/O group be shut down and restarted. If there is only one node in an I/O group, I/O operations could be lost if I/O operations are not stopped before beginning the software upgrade.

Explanation

The upgrade of the cluster software could not proceed because there is an I/O group in the cluster that contains only one node. The software upgrade requires that each node in an I/O group be shut down and restarted. If there is only one node in an I/O group, I/O operations could be lost if I/O operations are not stopped before beginning the software upgrade.

Action

Either upgrade the cluster using the -force option or specify a different node, and resubmit the command.

CMMVC5803E The entry in the error log was not marked because the error is already fixed or unfixed, or the sequence number could not be found.

Explanation

The entry in the error log was not marked because the sequence number was not found.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

Not applicable.

CMMVC5810E The quorum index number for the managed disk (MDisk) was not set because the MDisk is not available, either because the MDisk is offline or because it is associated with another quorum disk.

Explanation

The MDisk that you specify for this task must be online and cannot already be a quorum disk.

Action

Either change the status of the MDisk to online or specify a different MDisk, and resubmit the command.

CMMVC5811E The quorum index number for the managed disk (MDisk) was not set because the quorum disk does not exist.

Explanation

The quorum index number for the managed disk (MDisk) was not set because the quorum disk does not exist.

Action

Specify a different quorum disk, and resubmit the command.

CMMVC5812E The quorum index number for the managed disk (MDisk) was not set because the MDisk is in the wrong mode.

Explanation

The quorum index number for the managed disk (MDisk) was not set because the MDisk is not in the managed mode.

Action

- Change the mode of the MDisk, and resubmit the command
- Select an MDisk with a mode of managed, and resubmit the command

CMMVC5813E The quorum index number for the managed disk (MDisk) was not set because the MDisk has a sector size that is not valid.

Explanation

The parameter list that was specified is not valid.

Action

Specify a different sector size for the MDisk, 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.

Action

Specify an MDisk that has quorum enabled on all of its associated controllers, and resubmit the command.

CMMVC5815E The managed disk (MDisk) group was not created because an entity that was specified in the command does not exist.

Explanation

The managed disk (MDisk) group was not created because an entity that was specified in the command does not exist.

Action

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.

Action

Specify a different entity, and resubmit the command.

CMMVC5817E The specified managed disk (MDisk) group was invalid.

Explanation

The managed disk (MDisk) group was not renamed because the name was not valid.

Action

Specify a different MDisk group name, and resubmit the command.

CMMVC5818E The managed disk (MDisk) group was not deleted because there is at least one MDisk in the group.

Explanation

The managed disk (MDisk) group was not deleted because there is at least one MDisk in the group.

Action

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 MDisk group because the MDisk is part of another MDisk group.

Action

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 MDisk group because an entity that was specified in the command does not exist.

Action

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 MDisk group because not enough MDisks were included in the list.

Action

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 MDisk group because too many MDisks were included in the list.

Action

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 MDisk group because the MDisk is part of another MDisk group.

Action

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 MDisk group because it does not belong to the MDisk group.

Action

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 MDisk group because a virtual disk (VDisk) is allocated from one or more of the specified MDisks.

Action

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 virtual disk (VDisk) was not created because an entity that was specified in the command does not exist.

Action

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.

Action

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 virtual disk (VDisk) was not created because the I/O group contains no nodes.

Action

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 VDisk was not created because more than one MDisk is specified.

Action

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 virtual disk (VDisk) was not created because no managed disk (MDisk) was specified in the command.

Action

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 virtual disk (VDisk) was not created because the preferred node for I/O operations is not part of the I/O group.

Action

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 virtual disk (VDisk) was not modified because an entity that was specified in the command does not exist.

Action

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 virtual disk (VDisk) was not modified because there are no nodes in the I/O group.

Action

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 virtual disk (VDisk) was not modified because the group is a recovery I/O group.

Action

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 virtual disk (VDisk) was not expanded because an entity that was specified in the command does not exist.

Action

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.

Action

Wait for all commands to complete. Use the `svcinfo lsmigrate` command to view any migrates running in the background.

CMMVC5837E The action failed because the virtual disk (VDisk) is part of a FlashCopy mapping.

Explanation

The action failed because the virtual disk (VDisk) is part of a FlashCopy mapping.

Action

Specify a different VDisk 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 virtual disk (VDisk) is part of a Remote Copy mapping.

Action

Specify a different VDisk 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 virtual disk (VDisk) was not shrunk because an object that was specified in the command does not exist.

Action

Specify a different object, and resubmit the command.

CMMVC5840E The virtual disk (VDisk) was not deleted because it is mapped to a host or because it is part of a FlashCopy or Remote Copy mapping, or is involved in an image mode migrate.

Explanation

The virtual disk (VDisk) was not deleted because it is mapped to a host or because it is part of a FlashCopy or Metro Mirror mapping.

Action

Specify a different VDisk, and resubmit the command.

CMMVC5841E The virtual disk (VDisk) was not deleted because it does not exist.

Explanation

The virtual disk (VDisk) was not deleted because it does not exist.

Action

Specify a different VDisk, 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.

Action

Specify a different entity, and resubmit the command.

CMMVC5843E The virtual disk (VDisk)-to-host mapping was not created because the VDisk does not have a capacity greater than zero bytes.

Explanation

The virtual disk (VDisk)-to-host mapping was not created because the VDisk does not have a capacity greater than zero bytes.

Action

Specify a VDisk in which its capacity is greater than zero bytes, and resubmit the command.

CMMVC5844E The virtual disk (VDisk)-to-host mapping was not created because the SCSI logical unit number (LUN) ID is not valid.

Explanation

The virtual disk (VDisk)-to-host mapping was not created because the SCSI logical unit number (LUN) ID is not valid.

Action

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.

Action

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 virtual disk (VDisk) was not migrated because an object that was specified in the command does not exist.

Action

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 virtual disk (VDisk) was not migrated because its associated managed disk (MDisk) is already in the MDisk group.

Action

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 virtual disk (VDisk) does not exist or it is being deleted.

Action

Specify a different VDisk, 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.

Action

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.

Action

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.

Action

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.

Action

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 VDisk which is using an MDisk group with one of the following problems:

- The target and source MDisk groups have different extent sizes (group migrate).
- The target and source MDisk groups are the same (group migrate).
- The target and source MDisk groups are different (extents migrate).
- Target invalid group (Group Migrate).
- Source invalid group (Group Migrate).

Action

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.

Action

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 virtual disk (VDisk).

Action

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 (MDisk) group.

Explanation

The action failed because the virtual disk (VDisk) does not belong to the specified managed disk (MDisk) group.

Action

Specify a different VDisk, 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 (MDisk) group.

Explanation

The action failed because the managed disk (MDisk) does not exist or it is not a member of the managed disk (MDisk) group.

Action

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 virtual disk (VDisk) is in the wrong mode, the managed disk (MDisk) is in the wrong mode, or both are in the wrong mode.

Action

Check that the VDisk 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 virtual disk (VDisk).

Action

Not applicable.

CMMVC5860E The action failed because there were not enough extents in the managed disk (MDisk) 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 VDisk.

Action

In this case, the MDisk group reports that it has enough free capacity to create the VDisk. You can check the free capacity on each MDisk by issuing the `svcinfolsfreeextents <mdiskname/ID>`. Alternatively, do not specify a stripe set and let the system choose the free extents automatically.

CMMVC5861E The action failed because there were not enough extents on the managed disk (MDisk).

Explanation

The action failed because there were not enough extents on the managed disk (MDisk).

Action

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 virtual disk (VDisk) is being formatted.

Action

Wait for the VDisk 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).

Action

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.

Action

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 virtual disk (VDisk).

Action

Specify a different extent which is in range for the MDisk or VDisk 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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

Specify the correct entity, and resubmit the command.

CMMVC5871E The action failed because one or more of the configured port names is in a mapping.

Explanation

The action failed because one or more of the configured port names is in a mapping.

Action

Specify a port name that is not in a mapping, and resubmit the command.

CMMVC5872E The port (WWPN) was not added to the host object because an object that was specified in the command does not exist.

Explanation

The port (WWPN) was not added to the host object because an object that was specified in the command does not exist.

Action

Specify the correct object, and resubmit the command.

CMMVC5873E No matching WWPN.

Explanation

The action failed because there is no matching worldwide port name.

Action

Not applicable.

CMMVC5874E The action failed because the host does not exist.

Explanation

The action failed because the host does not exist.

Action

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 virtual disk (VDisk) does not exist.

Action

Specify a different VDisk, and resubmit the command.

CMMVC5876E The virtual disk (VDisk)-to-host mapping was not created because the maximum number of mappings has been reached.

Explanation

The virtual disk (VDisk)-to-host mapping was not created because the maximum number of mappings has been reached.

Action

Not applicable.

CMMVC5877E The virtual disk (VDisk)-to-host mapping was not created because the maximum number of SCSI LUNs has been allocated.

Explanation

The virtual disk (VDisk)-to-host mapping was not created because the maximum number of SCSI LUNs has been allocated.

Action

Not applicable.

CMMVC5878E The virtual disk (VDisk)-to-host mapping was not created because this VDisk is already mapped to this host.

Explanation

The virtual disk (VDisk)-to-host mapping was not created because this VDisk is already mapped to this host.

Action

Specify a different VDisk, and resubmit the command.

CMMVC5879E The virtual disk (VDisk)-to-host mapping was not created because a VDisk is already mapped to this host with this SCSI LUN.

Explanation

The virtual disk-to-host mapping was not created because this SCSI LUN is already assigned to another mapping.

Action

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 virtual disk (VDisk)-to-host mapping was not created because the VDisk has a capacity of zero bytes.

Action

Specify a different VDisk, 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.

Action

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 virtual disk (VDisk) already exists.

Action

Specify a different source or target VDisk, 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 virtual disk (VDisk).

Action

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 virtual disk (VDisk) cannot be a member of a Remote Copy mapping.

Action

Specify a different source or target VDisk, 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 virtual disk (VDisk) cannot be a member of a FlashCopy mapping.

Action

Specify a different source or target VDisk, 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 virtual disk (VDisk) is associated with the recovery I/O group.

Action

Specify a different source or target VDisk, 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 virtual disk (VDisk) must not be in router mode.

Action

Specify a different source or target VDisk, 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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

Perform the following steps:

1. Ensure that at least one of the nodes in the mapping's I/O group is online.
2. Fixed unfixed errors in the error log.
3. Follow the Directed Maintenance Procedures.

You might be required to delete and re-add ALL of the FlashCopy maps and Global and Metro Mirror relationships in the I/O group.

Resubmit the command.

CMMVC5919E The FlashCopy mapping or consistency group was not started, either because there are no online nodes in the I/O group or because there are unrecovered FlashCopy mappings or unrecovered Global Mirror or Metro Mirror relationships in the I/O group.

Explanation

This error might be caused by a temporary loss of all of the nodes in the I/O group, which causes all of the FlashCopy mappings and Global and Metro Mirror relationships of the I/O group to be unusable.

Action

Perform the following steps:

1. Ensure that at least one of the nodes in the mapping's I/O group is online.
2. Fixed unfixed errors in the error log.
3. Follow the Directed Maintenance Procedures.

You might be required to delete and re-add ALL of the FlashCopy maps and Global and Metro Mirror relationships in the I/O group.

Resubmit the command.

CMMVC5920E The FlashCopy mapping was not created because the consistency group is not idle.

Explanation

The FlashCopy mapping was not created because the consistency group is not idle.

Action

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.

Action

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 virtual disk (VDisk) is too small.

Action

Specify a different VDisk, and resubmit the command.

CMMVC5923E The FlashCopy mapping cannot be created, either because there are no online nodes in the I/O group or because there are unrecovered FlashCopy mappings or unrecovered Global Mirror or Metro Mirror relationships in the I/O group.

Explanation

This error might be caused by a temporary loss of all of the nodes in the I/O group, which causes all of the FlashCopy mappings and Global and Metro Mirror relationships of the I/O group to be unusable.

Action

Perform the following steps:

1. Ensure that at least one of the nodes in the mapping's I/O group is online.
2. Fixed unfixed errors in the error log.
3. Follow the Directed Maintenance Procedures.

You might be required to delete and re-add ALL of the FlashCopy maps and Global and Metro Mirror relationships in the I/O group.

Resubmit the command.

CMMVC5924E The FlashCopy mapping was not created because the source and target virtual disks (VDisks) are different sizes.

Explanation

The FlashCopy mapping was not created because the source and target virtual disks (VDisks) are different sizes.

Action

Specify a different source and target VDisk 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.

Action

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.

Action

Not applicable.

CMMVC5927E The action failed because the cluster ID is not valid.

Explanation

The action failed because the cluster ID is not valid.

Action

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.

Action

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.

Action

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.

Action

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 virtual disk (VDisk) is locked.

Action

Unlock the master or auxiliary VDisk, 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 virtual disk (VDisk) is a member of a FlashCopy mapping.

Action

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 virtual disk (VDisk) is in the recovery I/O group.

Action

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 virtual disk (VDisk) is in the router mode.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

Not applicable.

CMMVC5940E The cluster that contains the auxiliary virtual disk (VDisk) is unknown.

Explanation

The cluster that contains the auxiliary virtual disk (VDisk) is unknown.

Action

Not applicable.

CMMVC5941E The cluster that contains the master virtual disk (VDisk) has too many consistency groups.

Explanation

The cluster that contains the master virtual disk (VDisk) has too many consistency groups.

Action

Not applicable.

CMMVC5942E The cluster that contains the auxiliary virtual disk (VDisk) has too many consistency groups.

Explanation

The cluster that contains the auxiliary virtual disk (VDisk) has too many consistency groups.

Action

Not applicable.

CMMVC5943E The specified relationship is not valid.

Explanation

The specified relationship is not valid.

Action

Specify the correct relationship, and resubmit the command.

CMMVC5944E The specified consistency group is not valid.

Explanation

The specified consistency group is not valid.

Action

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.

Action

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.

Action

Specify the correct auxiliary cluster, and resubmit the command.

CMMVC5947E The specified master virtual disk (VDisk) is not valid.**Explanation**

The specified master virtual disk (VDisk) is not valid.

Action

Specify the correct master VDisk, and resubmit the command.

CMMVC5948E The specified auxiliary virtual disk (VDisk) is not valid.**Explanation**

The specified auxiliary virtual disk (VDisk) is not valid.

Action

Specify the auxiliary VDisk, and resubmit the command.

CMMVC5949E The specified relationship is unknown.**Explanation**

The specified relationship is unknown.

Action

Specify a different relationship, and resubmit the command.

CMMVC5950E The specified consistency group is unknown.

Explanation

The specified consistency group is unknown.

Action

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.

Action

Not applicable.

CMMVC5952E The relationship and consistency group have different master clusters.

Explanation

The relationship and consistency group have different master clusters.

Action

Not applicable.

CMMVC5953E The relationship and group have different auxiliary clusters.

Explanation

The relationship and group have different auxiliary clusters.

Action

Not applicable.

CMMVC5954E The master and auxiliary virtual disks (VDisks) are different sizes.

Explanation

The master and auxiliary virtual disks (VDisks) are different sizes

Action

Not applicable.

CMMVC5955E The maximum number of relationships has been reached.

Explanation

The maximum number of relationships has been reached.

Action

Not applicable.

CMMVC5956E The maximum number of consistency groups has been reached.

Explanation

The maximum number of consistency groups has been reached.

Action

Not applicable.

CMMVC5957E The master virtual disk (VDisk) is already in a relationship.

Explanation

The master virtual disk (VDisk) is already in a relationship.

Action

Specify a different master VDisk, and resubmit the command.

CMMVC5958E The auxiliary virtual disk (VDisk) is already in a relationship.

Explanation

The auxiliary virtual disk (VDisk) is already in a relationship.

Action

Specify a different auxiliary VDisk, 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.

Action

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.

Action

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.

Action

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.

Action

Specify a different name, and resubmit the command.

CMMVC5963E No direction has been defined.

Explanation

No direction has been defined.

Action

Not applicable.

CMMVC5964E The copy priority is not valid.

Explanation

The copy priority is not valid.

Action

Not applicable.

CMMVC5965E The virtual disks (VDisks) are in different I/O groups on the local cluster.

Explanation

The virtual disks (VDisks) are in different I/O groups on the local cluster.

Action

Not applicable.

CMMVC5966E The master virtual disk (VDisk) is unknown.

Explanation

The master virtual disk (VDisk) is unknown.

Action

Specify a different master VDisk, and resubmit the command.

CMMVC5967E The auxiliary virtual disk (VDisk) is unknown.

Explanation

The auxiliary virtual disk (VDisk) is unknown.

Action

Specify a different auxiliary VDisk, 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.

Action

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.

Action

Perform the following steps:

1. Ensure that at least one of the nodes in the I/O group is online.
2. Fixed unfixed errors in the error log.
3. Follow the Directed Maintenance Procedures.

You might be required to delete and re-add ALL of the FlashCopy maps and Global and Metro Mirror relationships in the I/O group.

Resubmit the command.

CMMVC5970E The Remote Copy relationship was not created because there is not enough memory.

Explanation

The Remote Copy relationship was not created because there is not enough memory.

Action

Not applicable.

CMMVC5971E The operation was not performed because the consistency group contains no relationships.

Explanation

The operation was not performed because the consistency group contains no relationships.

Action

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.

Action

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.

Action

Specify the Force option when starting the consistency group.

CMMVC5974E The operation was not performed because the consistency group is offline.

Explanation

The operation was not performed because the consistency group is offline.

Action

Not applicable.

CMMVC5975E The operation was not performed because the cluster partnership is not connected.

Explanation

The operation was not performed because the cluster partnership is not connected.

Action

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.

Action

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.

Action

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 VDisks since the ConsistentStopped or Idling state has occurred. Because the relationship is no longer synchronized, the state of the relationship is now Stopped.

The -force parameter of the svctask startrelationship command is required when the relationship is not synchronized because consistency would be lost by starting the copy operation. Submitting the svctask startrelationship command on an unsynchronized relationship without using the -force parameter is not supported.

If a relationship is in the InconsistentStopped, InconsistentCopying or ConsistentSynchronized state, the -force parameter is not required, but is supported.

Action

Consider using the -force parameter of the svctask startrelationship command, if appropriate.

CMMVC5980E The operation was not performed because the master and auxiliary clusters are not connected.

Explanation

The operation was not performed because the master and auxiliary clusters are not connected.

Action

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.

Action

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.

Action

Not applicable.

CMMVC5983E dump file was not created. This may be due to the file system being full.

Explanation

dump file was not created. This may be due to the file system being full.

Action

Not applicable.

CMMVC5984E The dump file was not written to disk. This may be due to the file system being full.

Explanation

The dump file was not written to disk. This may be due to the file system being full.

Action

Not applicable.

CMMVC5985E The action failed because the directory that was specified was not one of the following directories: /dumps, /dumps/iostats, /dumps/iotrace, /dumps/feature, /dumps/configs, /dumps/elogs, or /home/admin/upgrade.

Explanation

The action failed because the directory that was specified was not one of the following directories:

- /dumps
- /dumps/iostats
- /dumps/iotrace
- /dumps/feature
- /dumps/configs
- /dumps/elogs
- /home/admin/upgrade

Action

Specify one of the above directories, 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 virtual disk (VDisk) or managed disk (MDisk) failed to return statistics.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

Not applicable.

CMMVC5993E The specified upgrade package does not exist.

Explanation

The specified upgrade package does not exist.

Action

Not applicable.

CMMVC5994E Error in verifying the signature of the upgrade package.

Explanation

The system could not verify the signature of the upgrade package due to the following reasons:

- There is not enough space on the system to copy the file.

- The package is incomplete or contains errors.

Action

If the copy failed with an error indicating that there was insufficient space on the system, free up additional space on your system. Otherwise, ensure that the cluster time and date stamp on the signature is correct. (For example, the time and date cannot be in the future.)

CMMVC5995E Error in unpacking the upgrade package.

Explanation

The most likely cause of this error is lack of system space.

Action

Reboot the node and unpack the upgrade package again.

CMMVC5996E The specific upgrade package cannot be installed over the current version.

Explanation

The upgrade package is not compatible with the current version or the system.

Action

Check the available upgrade packages and find the correct upgrade package for your current version and for your system. If the upgrade package is correct for your system, check the version requirements for the package. You might have to upgrade the current version to an intermediate version before you upgrade to the latest version. (For example, if your current version is 1 and you are trying to upgrade to version 3, you might need to upgrade to version 2 before applying the version 3 upgrade.)

CMMVC5999W Featurization for this facility has not been enabled.

Explanation

Featurization for this facility has not been enabled.

Action

Not applicable.

CMMVC6000W Featurization for this facility has not been enabled.

Explanation

Featurization for this facility has not been enabled.

Action

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.

Action

Create a FlashCopy within the appropriate group.

CMMVC6002E This command can only be run on a node that is in service mode.

Explanation

This command can only be run on a node that is in service mode.

Action

Not applicable.

CMMVC6003E This command can not be run on a node that is in service mode.

Explanation

This command can not be run on a node that is in service mode.

Action

Not applicable.

CMMVC6004E The delimiter value *VALUE* is invalid.

Explanation

The specified value is not a valid delimiter value.

Action

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.

Action

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 MDisk group that is being used as a source and destination for migration operations.

Action

Ensure that the MDisk group 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.

Action

Re-enter the passwords.

CMMVC6008E The key already exists.

Explanation

An attempt was made to load a duplicate SSH key.

Action

Not applicable.

CMMVC6009E Unable to malloc a block of memory in which to copy the returned data.

Explanation

The command line was unable to allocate a block of memory in which to copy the results of the query.

Action

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.

Action

Not applicable.

CMMVC6011E This cluster is part of a remote cluster partnership. Because this upgrade package will make changes to the cluster state, it cannot be applied to the current code level until all remote cluster partnerships are deleted.

Explanation

You have attempted to apply software when a Remote Copy relationship to a remote cluster exists.

Action

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.

Action

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.

Action

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.

Action

Ensure that all parameters have been correctly entered. If this is the case the determine why the object is unavailable, then resubmit the command.

CMMVC6015E A delete request is already in progress for this object.

Explanation

A delete request is already in progress for this object.

Action

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.

Action

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.

Action

Ensure that all input to the CLI is ASCII, then resubmit the command.

CMMVC6018E The software upgrade pre-install process failed.

Explanation

The software upgrade failed as there was an error during the preprocessing. The package is either invalid or corrupted.

Action

Ensure the package is a valid upgrade package. Download the package from the source location again as it might have been corrupted during a network transfer.

CMMVC6019E The software upgrade failed as a node pended as the upgrade was in progress.

Explanation

The software upgrade failed as a node pended as the upgrade was in progress.

Action

Ensure that all nodes are online and available before restarting the upgrade process.

CMMVC6020E The software upgrade failed as the system was unable to distribute the software package to all nodes.

Explanation

The software upgrade failed as the system was unable to distribute the software package to all nodes.

Action

Ensure that all nodes are correctly zoned and that all nodes are online and can see the other nodes in the cluster. You might also want to check the error log.

CMMVC6021E The system is currently busy performing another request. Try again later.

Explanation

The requested action failed as the system is processing another request.

Action

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.

Action

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.

Action

Wait before resubmitting the request.

CMMVC6024E The auxiliary VDisk entered is invalid.**Explanation**

The auxiliary VDisk is entered as a parameter in the command-line interface is not a valid auxiliary VDisk.

Action

Select a valid auxiliary VDisk, and resubmit the command.

CMMVC6025E The RC consistency group Master cluster is not the local cluster.**Explanation**

The auxiliary VDisk is entered as a parameter in the command-line interface is not a valid auxiliary VDisk.

Action

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.

Action

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.

Action

Ensure that the parameters have been entered correctly on the command line.

CMMVC6028E This upgrade package cannot be applied to the current software level because it contains changes to the cluster state and there are remote cluster partnership defined.

Explanation

The action failed because there is a connected remote cluster. The upgrade cannot be applied because it would render the remote cluster at a different code level to the remote cluster.

Action

Ensure that the cluster partnership is unconfigured before resubmitting the command. Ensure that you unconfigure the remote cluster and upgrade the code on it before reconfiguring the cluster partnership.

CMMVC6029E All nodes must have identical code level before a concurrent code upgrade can be performed.

Explanation

The concurrent upgrade failed as two or more nodes were at differing code levels. All nodes must be at the same code level before a software upgrade can be performed.

Action

Use the service mode to bring all nodes to the same level before resubmitting the concurrent upgrade.

CMMVC6030E The operation was not performed because the FlashCopy mapping is part of a consistency group. The action must be performed at the consistency group level.

Explanation

An attempt was made to stop a FlashCopy mapping. This failed as the FlashCopy mapping is part of a consistency group.

Action

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.

Action

Not applicable.

CMMVC6032E The operation was not performed because one or more of the entered parameters is invalid for this operation.

Explanation

An invalid parameter was entered for the command.

Action

If attempting to change the I/O group to which the VDisk belongs, ensure that the VDisk 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.

Action

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.

Action

Not applicable.

CMMVC6035E The action failed as the object already exists.

Explanation

An operation was requested to create an object that already exists.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

Determine the correct number of candidates required for the specific command and re-issue the command.

CMMVC6048E The action failed as the object is busy.**Explanation**

The action failed because the object is busy.

Action

Not applicable.

CMMVC6049E The action failed as the object is not ready.**Explanation**

The action failed because the object is not ready.

Action

Not applicable.

CMMVC6050E The action failed as the command was busy.**Explanation**

The action failed because the command is busy.

Action

Not applicable.

CMMVC6051E An unsupported action was selected.**Explanation**

The action failed because it is not valid with the command.

Action

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.

Action

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

An invalid worldwide port name (WWPN) was specified.

Action

Specify a valid WWPN.

CMMVC6054E The action failed as not all nodes are online.

Explanation

One or more nodes are not online.

Action

Check that each node is online, and resubmit the command.

CMMVC6055E The action failed as an upgrade is in progress.

Explanation

The action failed because a software upgrade is in progress.

Action

Wait for the software upgrade to complete, and resubmit the command.

CMMVC6056E The action failed as the object is too small.

Explanation

The action failed because the object is too small.

Action

Specify a different object, and resubmit the command.

CMMVC6057E The action failed as the object is the target of a FlashCopy mapping.

Explanation

The object is the target of a FlashCopy mapping, thus it cannot be deleted.

Action

Specify an object that is not the target of a FlashCopy mapping, or remove the object from the FlashCopy mapping.

CMMVC6058E The action failed as the object is in the recovery HWS.

Explanation

An attempt was made to perform an operation on a node that is in the recovery IO group.

Action

Get the node into one of the other IO Groups and re-issue 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.

Action

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.

Action

Not applicable.

CMMVC6061E The action failed as the object is being resized.

Explanation

The action failed because the object is being resized.

Action

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 IO groups.

Action

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.

Action

Either add disks to the group and re-issue 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 an invalid name.

Action

Use a name that meets the naming standards and re-issue 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.

Action

Ensure that the object is a member of an appropriate group and re-issue the command.

CMMVC6066E The action failed as the system is running low on memory.

Explanation

The system is running low on memory.

Action

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.

Action

Re-issue the command using a key that does exist.

CMMVC6068E The action failed as there are no free SSH keys.

Explanation

An attempt was made to use an SSH key when there are no free SSH keys.

Action

Upload additional keys and re-issue 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.

Action

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 invalid.

Action

Correct the parameters and re-issue the command.

CMMVC6071E The virtual disk (Vdisk)-to-host mapping was not created because the Vdisk is already mapped to a host.

Explanation

The virtual disk is already mapped to a host.

Action

Not applicable.

CMMVC6073E The maximum number of files has been exceeded.**Explanation**

The maximum number of files has been exceeded.

Action

Not applicable.

CMMVC6074E The command failed as the extent has already been assigned.**Explanation**

The command failed as the extent has already been assigned.

Action

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.

Action

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 VDisk.

Action

Not applicable.

CMMVC6077E WARNING - Unfixed errors should be fixed before applying software upgrade. Depending on the nature of the errors, they might cause the upgrade process to fail. It is highly recommended to fix these errors before proceeding. If a particular error cannot be fixed, contact the support center.

Explanation

Unfixed errors should be fixed before applying software upgrade. Depending on the nature of the errors, they might cause the upgrade process to fail. It is highly recommended to fix these errors before proceeding.

Action

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.

Action

Get the object into a suitable mode and re-issue 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 invalid.

CMMVC6081E Metadata Recovery is busy processing the previous operation.

Explanation

Metadata Recovery is busy processing the previous operation.

CMMVC6082E The attempt to abort metadata recovery failed because the previous operation has completed.

Explanation

The attempt to abort metadata recovery failed because the previous operation has completed.

Action

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.

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.

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.

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.

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.

CMMVC6089E The metadata at the requested lba is flagged as invalid.

Explanation

The metadata at the requested lba is flagged as invalid.

CMMVC6090E The metadata header checksum verification failed.

Explanation

The metadata header checksum verification failed.

CMMVC6091E The metadata region checksum verification failed.

Explanation

The metadata region checksum verification failed.

CMMVC6092E The metadata recovery operation was aborted.

Explanation

The metadata recovery operation was aborted.

CMMVC6093E Metadata recovery internal error - (read only)

Explanation

Metadata recovery internal error - (read only)

CMMVC6095E Metadata recovery encountered the end of the disk.

Explanation

Metadata recovery encountered the end of the disk.

CMMVC6096E Metadata recovery encountered an error from a lower layer - (v1 no resource).

Explanation

Metadata recovery encountered an error from a lower layer - (v1 no resource).

CMMVC6097E Metadata recovery encountered an error from a lower layer - (v1 failure).

Explanation

Metadata recovery encountered an error from a lower layer - (v1 failure).

CMMVC6098E The copy failed as the specified node is the configuration node.

Explanation

The copy failed because the specified node is the configuration node.

Action

Not applicable.

CMMVC6100E *OPTION* not consistent with *ACTION*

Explanation

The specified option is not supported for the specified action.

Action

Remove the option, and resubmit the command.

CMMVC6101E *OPTION* not consistent with *OPTION*

Explanation

The two specified options cannot be used together.

Action

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.

Action

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.

Action

Correct the problem, and resubmit the command.

CMMVC6104E Action *ACTION* not run**Explanation**

An unexpected error has occurred.

Action

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.

Action

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*.

Action

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.

Action

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.

Action

Not applicable.

CMMVC6109E Disk controller system with a WWNN of *WWNN_VALUE* not available.

Explanation

A disk controller system with the specified WWNN has been found. Ensure that the specified disk controller system is available to the cluster.

Action

Ensure that the required disk controller system is available to the cluster, and resubmit the command.

CMMVC6110E Bad code level: *VALUE*.

Explanation

An unexpected error has occurred.

Action

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.

Action

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.

Action

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.

Action

Determine the cause of the problem, and resubmit the command.

CMMVC6114E No help for action *ACTION*.

Explanation

There is no help for the specified action topic.

Action

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.

Action

Not applicable.

CMMVC6116I Feature match for *FEATURE*.

Explanation

The features in the backup configuration file and the target cluster are an exact match.

Action

Not applicable.

CMMVC6117E *FIX_OR_FEATURE* is not available.

Explanation

An unexpected error has occurred.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

Contact IBM technical support for assistance.

CMMVC6123E No *PROPERTY* for *TYPE NAME*.

Explanation

An unexpected error has occurred.

Action

Contact IBM technical support for assistance.

CMMVC6124E No *TYPE* with *PROPERTY VALUE*

Explanation

An unexpected error has occurred.

Action

Contact IBM technical support for assistance.

CMMVC6125E No unique *ID* for *TYPE NAME*

Explanation

An unexpected error has occurred.

Action

Contact IBM technical support for assistance.

CMMVC6126E No *TYPE* with unique *ID VALUE*

Explanation

An unexpected error has occurred.

Action

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.

Action

Specify a different SSH key, and resubmit the command.

CMMVC6128W *DIRECTORY***Explanation**

The files in the specified directory cannot be listed.

Action

Determine why the files cannot be listed, correct the problem, and resubmit the command.

CMMVC6129E **VDisk-to-host map objects have VDisk_UID values that are not consistent****Explanation**

All of the VDisk-to-host map objects do not have the same number for the VDisk LUN instance. Therefore, there is a possibility the backup configuration file is corrupt. The LUN instance number should be the same for all VDisk-to-host map objects that are associated with a specific VDisk. The LUN instance number is incorporated into the VDisk_UID property.

Action

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.

Action

Not applicable.

CMMVC6131E **No location cluster information****Explanation**

An unexpected error has occurred.

Action

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.

Action

Change the state to the required value, and resubmit the command.

CMMVC6133E Required *TYPE* property *PROPERTY* not found

Explanation

An unexpected error has occurred.

Action

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.

Action

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.

Action

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.

Action

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.

Action

After the restore is complete, locate the file containing the key, and perform one of the following actions: (1) Rename the file so that it has the correct name, and resubmit the command. (2) Restore the key manually using the `svctask addsshkey` command.

CMMVC6138E *OPTION* is required

Explanation

An option is missing. The option might be listed as optional, but circumstances make the option mandatory.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

Not applicable.

CMMVC6144W The object with default name *NAME* has been restored as *SUBSTITUTE_NAME*.

Explanation

An object with a default name has been restored with a different name. Ensure that you account for this name change when using the restored cluster in the future. To avoid this problem in the future, choose an appropriate name for each object in the cluster.

Action

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.

Action

Not applicable.

CMMVC6146E Problem parsing *OBJECT_TYPE* data: *LINE***Explanation**

An unexpected error has occurred.

Action

Contact the support center.

CMMVC6147E *TYPE NAME* has a name beginning with *PREFIX*.**Explanation**

An object has been encountered that has a name beginning with the specified reserved prefix. The only valid reason for an object with this kind of name is that a restoration command did not complete successfully.

Action

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.

Action

Correct the problem, and resubmit the command.

CMMVC6149E An action is required.**Explanation**

An action is required to run the command.

Action

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.

Action

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.

Action

Specify a valid option, and resubmit the command.

CMMVC6152E VDisk *VDISK_NAME* instance number *INSTANCE_NUMBER* is not valid.**Explanation**

The VDisk cannot be restored because the instance number, which must be a hexadecimal number, is not valid.

Action

Contact IBM technical support for assistance.

CMMVC6153E *OBJECT* not consistent with *ACTION***Explanation**

The specified object is not supported for the specified action.

Action

Remove the object, and resubmit the command.

CMMVC6154E Required *OBJECT_TYPE* property *PROPERTY_NAME* has a null value.**Explanation**

An unexpected error has occurred.

Action

Contact IBM technical support for assistance.

CMMVC6155I The command *COMMAND* processing has completed successfully.**Explanation**

Only information and warning messages are issued.

Action

Not applicable.

CMMVC6156W *COMMAND* processing completed with errors.**Explanation**

Processing was not successful.

Action

Not applicable.

CMMVC6164E The SVCCONFIG CRON job, which runs overnight on a daily overnight, has failed.**Explanation**

The SVCCONFIG CRON job, which runs overnight on a daily overnight, has failed.

Action

Resolve any hardware and configuration problems that you are experiencing on the 2145 cluster. If the problem persists, contact IBM technical support for assistance.

CMMVC6165E The target is not the original configuration node with a WWNN of *WWNN_VALUE*.**Explanation**

A backup configuration can only be restored to the original configuration node.

Action

Recreate 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.

Action

Resubmit the command from *svconfig restore -prepare*.

CMMVC6181E The target cluster contains an object that has a counterpart in the configuration to be restored, and has the correct ID.

Explanation

The indicated property has an unexpected value.

Action

Check that the correct (matching) backup configuration file (svc.config.backup.xml) is being provided and if it is, use the -force option to ignore the discrepancy. Otherwise, provide the correct file and try again.

CMMVC6182W An object that does not contribute to the fabric of the configuration cannot be restored because its configuration does not permit it to be created.

Explanation

An object that does not contribute to the fabric of the configuration cannot be restored because its configuration does not permit it to be created. For example, a host can only be created if it has at least one port.

Action

Not applicable.

CMMVC6186W The IO group *IO_GROUP_NAME* has been restored with ID *ID_VALUE* instead of *ID_VALUE*.

Explanation

This can occur when the configuration node is different from the node that was used to create the original cluster. This affects the SCSI Inquiry value for the IO group.

Action

Not applicable.

CMMVC6200E The action failed because of incompatible software.

Explanation

The software version on one or more nodes is incompatible with the new version.

Action

Refer to the compatibility requirements for the software version you are adding. Update the cluster to meet the compatibility requirements, and then perform the upgrade.

CMMVC6201E The node could not be added because of incompatible software. The status code is *STATUS_CODE*.

Explanation

The node could not be added because of incompatible software.

Action

Upgrade the software on the node that has been rejected to the same level of software as the cluster to which it will be added, and resubmit the command.

CMMVC6202E The cluster was not modified because the IP address is not valid.

Explanation

An attempt was made to change the IP address of a cluster to an invalid address.

Action

Correct the address and re-issue the command.

CMMVC6203E The action failed because the directory that was specified was not one of the following directories: /dumps, /dumps/iostats, /dumps/iotrace, /dumps/feature, /dumps/config, /dumps/elogs, /dumps/ec or /dumps/pl.

Explanation

An attempt was made to clear a file from, or copy a file to, an invalid directory.

Action

Ensure that the command accesses a valid directory.

CMMVC6204E The action failed as the resulting disk size would be less than, or equal to, zero.

Explanation

An attempt was made to shrink a disk, however the resulting size would have been less than or equal to zero.

Action

Not applicable

CMMVC6205E Metadata recovery can not use the provided MDisk id - invalid or destroyed.

Explanation

Metadata recovery can not use the provided MDisk id - invalid or destroyed.

CMMVC6206E The software upgrade failed as a file containing the software for the specified MCP version was not found.

Explanation

There are two files required to successfully complete a software upgrade. One file contains the files that make up the base operating system, and the other file contains the 2145 software. This message appears if the OS version is incompatible with the 2145 software.

Action

Upload two compatible files, and resubmit the command.

CMMVC6207E The action failed because the virtual disk (VDisk) is part of a Remote Copy mapping.

Explanation

An action was performed against a VDisk that is part of a Remote Copy mapping.

Action

Remove the VDisk from the Remote Copy mapping before resubmitting the command.

CMMVC6208E The action failed because the virtual disk (VDisk) is part of a FlashCopy mapping.

Explanation

An action was performed against a VDisk that is part of a FlashCopy mapping.

Action

Remove the VDisk from the FlashCopy mapping before re-issuing the command.

CMMVC6209E The FlashCopy mapping or consistency group could not be started in a reasonable time. The mapping or group is instead being prepared.

Explanation

The FlashCopy mapping or consistency group could not be started in a reasonable time. The mapping or group is instead being prepared.

Action

Resubmit the command.

CMMVC6210E The command has failed because a virtual medium error exists on the image mode VDisk or copy.

Explanation

When you submit this command, you cannot specify an image mode VDisk that has a virtual medium error on the VDisk or on any copy of the VDisk because the medium errors cannot be maintained on the ejected MDisk image copy.

Action

If an exact image copy is required, ensure that there is no virtual medium error on the image mode VDisk that you specify or on any of its copies, and resubmit the command.

If an exact copy is not required, you can use the `-force` option of the command, but all of the virtual medium errors will be lost.

CMMVC6211E The command failed as a migrate to image was in progress.

Explanation

An attempt was made to execute a command against a VDisk that was involved in a migrate to image operation.

Action

Wait for the migration to complete and re-issue the command.

CMMVC6212E The command failed because data in the cache has not been committed to disk.

Explanation

The command failed because data in the cache has not been committed to disk.

CMMVC6213E You are trying to recover region data that was created by a code level different from the one you are currently running on the node.

Explanation

You are trying to recover region data that was created by a code level different from the one you are currently running on the node.

CMMVC6214E Failed to recreate the cluster you are trying to rebuild.

Explanation

Failed to recreate the cluster you are trying to rebuild.

CMMVC6215E The FlashCopy mapping was not created or modified because the consistency group already contains the maximum number of mappings.

Explanation

An attempt was made to create a FlashCopy mapping in, or move a FlashCopy mapping to, a consistency group that has the maximum number of FlashCopy mappings that it can contain.

Action

Create or move the FlashCopy mapping in another consistency group or remove an existing FlashCopy mapping from the desired group and then re-issue the command.

CMMVC6216E The Remote Copy relationship was not created because the master or auxiliary virtual disk (VDisk) is a member of a Remote Copy mapping.

Explanation

The Remote Copy relationship was not created because the master or auxiliary virtual disk (VDisk) is a member of a Remote Copy mapping.

Action

Select a different VDisk to make up the mapping.

CMMVC6217E The maximum number of hosts for the cluster is already configured.

Explanation

You must remove at least one host definition before you can resubmit the command.

Action

Determine whether the action is required.

If the action is required, review the current configuration to determine whether any current host definitions are not required. Remove at least one host definition that is not required, and resubmit the command.

CMMVC6218E The maximum number of host/IO group pairs for the cluster is already configured.

Explanation

You must remove at least one host/IO group pair definition before you can resubmit the command.

Action

Determine whether the action is required.

If the action is required, review the current configuration to determine whether any current host/IO group pair definitions are not required. Remove at least one host/IO group pair definition that is not required, and resubmit the command.

CMMVC6219E The maximum number of WWPNS for the cluster is already configured.

Explanation

You must remove at least one WWPNS definition before you can resubmit the command.

Action

Determine whether the action is required.

If the action is required, review the current configuration to determine whether any current WWPNS definitions are not required. Remove at least one WWPNS definition that is not required, and resubmit the command.

CMMVC6220E The maximum number of hosts for one or more IO groups is already configured.

Explanation

You must remove at least one host/IO group pair definition from the I/O group that you have specified before you can resubmit the command.

Action

Determine whether the action is required.

If the action is required, review the current configuration to determine whether any current host/IO group pair definitions for the I/O group that you have specified are not required. Remove at least one host/IO group pair definition that is not required from the I/O group that you have specified, and resubmit the command.

CMMVC6221E The maximum number of WPNs for one or more IO groups is already configured.

Explanation

You must remove at least one WWPN definition from the I/O group that you have specified before you can resubmit the command.

Action

Determine whether the action is required.

If the action is required, review the current configuration to determine whether any current WWPN definitions for the I/O group that you have specified are not required. Remove at least one WWPN definition that is not required from the I/O group that you have specified, and resubmit the command.

CMMVC6222E The maximum number of WPNs for the host is already configured.

Explanation

You must remove at least one WWPN definition for the host that you have specified before you can resubmit the command.

Action

Determine whether the action is required.

If the action is required, review the current configuration to determine whether any current WWPN definitions for the host that you have specified are not required. Remove at least one WWPN definition that is not required for the host that you have specified, and resubmit the command.

CMMVC6223E The host does not belong to one or more of the IO groups specified or inferred.

Explanation

The host does not belong to one or more of the IO groups specified or inferred.

Action

Specify a host/IO group combination that is currently defined, and resubmit the command.

CMMVC6224E The host already belongs to one or more of the IO groups specified.

Explanation

The host already belongs to one or more of the IO groups specified.

Action

None.

CMMVC6225E An IO group cannot be removed from a host because of one or more associated VDisks.

Explanation

An IO group cannot be removed from a host because of one or more associated VDisks.

CMMVC6226E The action was not completed because the cluster has reached the maximum number of extents in MDisk Groups.

Explanation

The cluster has reached the maximum number of extents in the MDisk Group; therefore, the action did not complete. You are attempting to use additional extents, for example by creating or expanding a VDisk. The action cannot be initiated because it would cause the maximum number of extents for a cluster to be exceeded.

Action

Free up extents by deleting other VDisks, and resubmit the command.

CMMVC6227I The package installed successfully.

Explanation

The package installed successfully.

Action

None.

CMMVC6228E The cluster was recovered and the CLI functionality is limited until the cause of the failure is determined and any corrective action taken. Contact IBM technical support for assistance.

Explanation

The cluster was recovered and the CLI functionality is limited.

Action

Contact IBM technical support for assistance.

CMMVC6229E The action failed as the SSH key has been revoked.

Explanation

The action failed as the SSH key has been revoked.

CMMVC6230E The action failed as the SSH key index (SSH_LABEL_ID) is invalid.

Explanation

The action failed as the SSH key index (SSH_LABEL_ID) is invalid.

CMMVC6231E The action failed as the audit table is full.

Explanation

The action failed as the audit table is full.

Action

Save the audit log to disk, and resubmit the command.

CMMVC6232E This operation cannot be performed because the cluster is currently aborting the previous software upgrade command.

Explanation

This operation cannot be performed because the cluster is currently aborting the previous software upgrade command.

Action

Wait until the previous software upgrade command has aborted successfully, and resubmit the command.

CMMVC6233E This operation cannot be performed because, either a software upgrade has not been started, or a software upgrade is in progress but is not in a state where it can be aborted.

Explanation

This operation cannot be performed because the software upgrade is making progress.

CMMVC6234E The upgrade cannot be aborted because at least one node has already committed to a new code level.

Explanation

The upgrade cannot be aborted because at least one node has already committed to a new code level.

CMMVC6235E An invalid response has been entered. The command has not been executed. Input is case sensitive. Enter either yes or no.

Explanation

An invalid response has been entered. The command has not been executed. Enter either yes or no.

Action

Enter either yes or no.

CMMVC6236E The command has not completed. A limited availability parameter has been entered without the required environment setting being set.

Explanation

The command has not completed. A limited availability parameter has been entered without the required environment setting being set.

CMMVC6237E The command failed as the remote cluster does not support global mirror.

Explanation

The command failed as the remote cluster does not support global mirror.

CMMVC6238E The copy type differs from other copies already in the consistency group.

Explanation

The copy type differs from other copies already in the consistency group.

Action

Ensure that the copy type of the mapping that you are attempting to add is the same copy type as the mappings in the consistency group to which you are attempting to add the mapping, and resubmit the command.

CMMVC6239E The FlashCopy mapping was not prepared because the mapping or consistency group is in the stopping state. The mapping or consistency group must first complete the stop operation and then be prepared

Explanation

You cannot prepare a FlashCopy mapping or consistency group when the FlashCopy mapping or consistency group is in the stopping state. If you want to prepare a FlashCopy mapping or consistency group, the FlashCopy mapping or consistency group must be in the Stopped or idle_or_copied state.

Action

Wait until the FlashCopy mapping or consistency group reaches the Stopped or idle_or_copied state and then resubmit the command.

CMMVC6240E The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the stopping state.

Explanation

You cannot modify the consistency group of a FlashCopy mapping when the FlashCopy mapping is in the stopping state. If you want to modify the consistency group of a FlashCopy mapping, the FlashCopy mapping must be in the Stopped or idle_or_copied state.

Action

Wait until the FlashCopy mapping reaches the Stopped or idle_or_copied state and then resubmit the command.

CMMVC6241E The FlashCopy mapping was not deleted because the mapping or consistency group is in the stopping state. The mapping or consistency group must be stopped first.

Explanation

You cannot delete a FlashCopy mapping or consistency group when the FlashCopy mapping or consistency group is in the stopping state. If you want to delete a FlashCopy mapping or consistency group, the FlashCopy mapping or consistency group must be in the Stopped or idle_or_copied state.

Action

Wait until the FlashCopy mapping or consistency group reaches the Stopped or idle_or_copied state and then resubmit the command.

CMMVC6242E The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the stopping state. The mapping or consistency group must first complete the stop operation and then be prepared.

Explanation

You cannot start a FlashCopy mapping or consistency group when the FlashCopy mapping or consistency group is in the stopping state. If you want to start a FlashCopy mapping or consistency group, the FlashCopy mapping or consistency group must be in the Prepared state.

Action

Wait until the FlashCopy mapping or consistency group reaches the Stopped or idle_or_copied state and then prepare the FlashCopy mapping or consistency group before starting it.

CMMVC6243E The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is already in the stopping state.

Explanation

A Stop FlashCopy mapping or consistency group task has already been submitted and is still in progress. When the task has completed successfully, the FlashCopy mapping or consistency group state will change to Stopped.

Action

None.

CMMVC6244E The FlashCopy mapping was not created because the source virtual disk (VDisk) cannot be the target for a FlashCopy mapping.

Explanation

A VDisk cannot simultaneously be both the source of a FlashCopy mapping and the target of a FlashCopy mapping. The source VDisk that you have specified is currently defined as the target of a FlashCopy mapping.

Action

You have two options. One option is specify a different source VDisk and resubmit the command. The other option is delete the existing FlashCopy mapping that defines the source VDisk that you have specified as the target VDisk, and resubmit the command.

CMMVC6245E The FlashCopy mapping was not created because the source virtual disk (VDisk) is already in the maximum number of FlashCopy mappings.

Explanation

The number of FlashCopy mappings in which a VDisk can be defined as the source VDisk is limited. The source VDisk that you have specified cannot be defined to another FlashCopy mapping because it is already defined as the source VDisk to the maximum number of FlashCopy mappings.

Action

You have two options. One option is specify a different source VDisk and resubmit the command. The other option is delete one of the existing FlashCopy mappings that contains the source VDisk and resubmit the command.

CMMVC6246E The FlashCopy mapping was not created because the target virtual disk (VDisk) is already a source VDisk in a FlashCopy mapping.

Explanation

A VDisk cannot simultaneously be both the source of a FlashCopy mapping and the target of a FlashCopy mapping. The target VDisk that you have specified is currently defined as the source of a FlashCopy mapping.

Action

You have two options. One option is specify a different target VDisk and resubmit the command. The other option is delete all of the existing FlashCopy mappings that contain the target VDisk that you have specified and resubmit the command.

CMMVC6247E The FlashCopy mapping was not created because the target virtual disk (VDisk) is already a target VDisk in a FlashCopy mapping.

Explanation

A VDisk cannot simultaneously be the target of more than one FlashCopy mapping. The target VDisk that you have specified is currently defined as the target of another FlashCopy mapping.

Action

You have two options. One option is specify a different target VDisk and resubmit the command. The other option is delete the existing FlashCopy mapping that contains the target VDisk that you have specified and resubmit the command.

CMMVC6248E The command failed because the authorization table is full.

Explanation

The command failed because the authorization table is full.

CMMVC6249E The command failed because the authorization record was not found or is already set to the default role.

Explanation

The command failed because the authorization record was not found or is already set to the default role.

CMMVC6250E The command failed because the authorization record is not set to the default role. Use rmauth to set the default role.

Explanation

The command failed because the authorization record is not set to the default role.

Action

Use `rmauth` to set the default role.

CMMVC6251E The command failed because the specified role was not found.

Explanation

The command failed because the specified role was not found.

CMMVC6252E The command failed authorization because the session ssh key is invalid or was deleted.

Explanation

The command failed authorization because the session ssh key is invalid or was deleted.

CMMVC6253E The task has failed because the user's role is not authorized to submit the command.

Explanation

One example of a user role restriction is that a user that has a role of Monitor cannot create a VDisk.

Action

Either log in as a user that has a role that is authorized to submit the task or change the role of the user account that you are using to a role that is authorized to submit the task, and resubmit the task.

CMMVC6254E The command failed because the specified ssh key was not found. NOTE This command must specify an admin key.

Explanation

The command failed because the specified ssh key was not found. NOTE This command must specify an admin key.

CMMVC6255E The command cannot set the authorization record to the default role. Use `rmauth` to set the default role.

Explanation

The command can not set the authorization record to the default role.

Action

Use `rmauth` to set the default role.

CMMVC6263E The command failed because the ssh key already exists or there is a duplicate ssh key.

Explanation

You have attempted to add an ssh key that already exists, and may have a different authorization level associated with it.

Action

Add a different ssh key if the existing ssh key of the same type does not have the authority level that you require.

CMMVC6269E Sendmail error EX_USAGE. A command or configuration line has been used incorrectly.

Explanation

The send e-mail task has failed because a command or a configuration line has been used incorrectly.

Action

Ensure that the e-mail settings are correct, and resubmit the task.

CMMVC6270E Sendmail error EX_DATAERR. Address is wrong, or the message is too large for the mailbox.

Explanation

The send e-mail task has failed because the message sent is too large or a recipient address is incorrect.

Action

Ensure that all addresses are correct and that the message is not too large, and resubmit the task.

CMMVC6271E Sendmail error EX_NOINPUT. An input file (not a system file) did not exist or was not readable.

Explanation

The send e-mail task has failed because a file is missing or cannot be read.

Action

Ensure that the e-mail system is configured correctly. Ensure that access permissions have been specified correctly for all e-mail configuration files, and resubmit the task.

CMMVC6272E Sendmail error EX_NOUSER. The sendmail command could not recognize a specified user.

Explanation

The send e-mail task has failed because the user and domain combination that you specified does not exist.

Action

Specify a defined user and domain combination, and resubmit the task.

CMMVC6273E Sendmail error EX_NOHOST. The sendmail command could not recognize the specified host name.

Explanation

The send e-mail task has failed because the host is not known to the e-mail system.

Action

Ensure that you have configured the SMTP environment correctly. Ensure that you specify a defined host, and resubmit the task.

CMMVC6274E Sendmail error EX_UNAVAILABLE. A required system resource is not available.

Explanation

The send e-mail task has failed because a required system resource is not available.

Action

Ensure that you have configured the SMTP environment correctly, and resubmit the task.

CMMVC6275E Sendmail error EX_SOFTWARE. An internal software error occurred (including bad arguments).

Explanation

The send e-mail task has failed because an incorrect parameter or parameter value has been detected.

Action

Ensure that you have configured the SMTP environment correctly. Specify only supported parameters and parameter values, and resubmit the task.

CMMVC6276E Sendmail error EX_OSERR. A system resource error prevented the sending of an e-mail.

Explanation

The send e-mail task has failed because a system resource error has occurred.

Action

Ensure that you have configured the SMTP environment correctly, and resubmit the task.

CMMVC6277E Sendmail error EX_OSFILE. Failed to open a critical system file.

Explanation

The send e-mail task has failed because a required system file cannot be opened.

Action

Ensure that the e-mail system is configured correctly. Ensure that access permissions have been specified correctly for all e-mail configuration files, and resubmit the task.

CMMVC6278E Sendmail error EX_CANTCREAT. An output file could not be written to by sendmail.

Explanation

The send e-mail task has failed because the system cannot write to a required output file.

Action

Ensure that the e-mail system is configured correctly. Ensure that access permissions have been specified correctly for all e-mail configuration files, and resubmit the task.

CMMVC6279E Sendmail error EX_IOERR. A system I/O error occurred during a sendmail operation. This could be due to a disk failure.

Explanation

The send e-mail task has failed because a write or read I/O operation has failed. This error might be caused by a disk device failure.

Action

Correct the root cause of the I/O failure, and resubmit the task.

CMMVC6280E Sendmail error EX_TEMPFAIL. The sendmail command could not create a connection to a remote system.

Explanation

The send e-mail task has failed because the sendmail application cannot establish a connection to the remote system.

Action

Ensure that the network connection to the remote system is functioning correctly, and resubmit the task.

CMMVC6281E Sendmail error EX_PROTOCOL. The remote system returned something that was incorrect during a protocol exchange.

Explanation

The send e-mail task has failed because an error in the protocol exchange has occurred.

Action

Ensure that the e-mail system is configured correctly. Ensure that you have configured the SMTP environment correctly, and resubmit the task.

CMMVC6282E Sendmail error EX_NOPERM. The user does not have permission to perform the requested operation.

Explanation

The send e-mail task has failed because the User ID does not have authorization to submit the task.

Action

Ensure that authorizations for your User ID in the e-mail and SMTP configurations are correct, and resubmit the task.

CMMVC6283E Sendmail error EX_CONFIG. There is a fatal problem with the sendmail configuration.

Explanation

The send e-mail task has failed because the sendmail configuration is not correct.

Action

Ensure that the e-mail system is configured correctly. Ensure that you have configured the SMTP environment correctly, and resubmit the task.

CMMVC6284E An unknown error occurred. Please ensure your SMTP server is running.

Explanation

The send e-mail task has failed because an unexpected error has occurred.

Action

Ensure that the SMTP server is running, and resubmit the task.

CMMVC6285E The e-mail command timed out. Please check your e-mail server settings as listed on the SVC.

Explanation

The send e-mail task has failed because a command timeout has occurred.

Action

Ensure that your system settings match those recommended in the sendmail application documentation, and resubmit the task.

CMMVC6286E The e-mail service has not been enabled.

Explanation

The send e-mail task has failed because the e-mail application is not enabled.

Action

Enable the e-mail application, and resubmit the task.

CMMVC6287E The user specified does not exist.

Explanation

You must specify a User ID that exists.

Action

Ensure that the User ID that you specify is defined, and resubmit the task.

CMMVC6288E The FlashCopy mapping or consistency group could not be started because a source VDisk is the target of another FC Map that is keeping the VDisk inaccessible.

Explanation

You cannot start a FlashCopy mapping or consistency group when a source VDisk in the FlashCopy mapping or consistency group is the target VDisk of another FlashCopy mapping that is holding the VDisk as inaccessible. The task cannot be initiated because a source VDisk in the FlashCopy mapping or consistency group that you are attempting to start is the target of another FlashCopy mapping that is

either prepared, preparing, stopped or stopping with a progress of less than 100%.

Action

Ensure that you have selected to start the correct FlashCopy mapping or consistency group. Ensure that none of the source VDIs in the FlashCopy mapping or consistency group that you specify are target VDIs in another FlashCopy mapping that is prepared, preparing, stopped or stopping with a progress of less than 100%. Resubmit the task.

CMMVC6289E The command failed because the virtual disk (VDisk) is pending synchronization.

Explanation

This error occurs when at least one of the virtual disk copies is offline.

Action

Fix all of the errors that are associated with the virtual disk copies, and resubmit the command.

CMMVC6290E The command failed because the virtual disk (VDisk) has image mode copies, is pending synchronization and -force has not been specified.

Explanation

This error occurs when at least one of the virtual disk copies is offline.

Action

Perform one of the following actions:

- Fix all of the errors that are associated with the virtual disk copies, and resubmit the command.
- Resubmit the command and specify the -force parameter.

NOTE: When you specify the -force parameter with the command that caused this error, the image mode virtual disk copy is no longer guaranteed to have the correct virtual disk data.

CMMVC6291E The command failed because the virtual disk (VDisk) is pending synchronization and -force has not been specified.

Explanation

The command failed because the virtual disk (VDisk) is pending synchronization and -force has not been specified.

Action

Perform one of the following actions:

- Fix all of the errors that are associated with the virtual disk copies, and resubmit the command.
- Resubmit the command and specify the `-force` parameter.

NOTE: When you specify the `-force` parameter with the command that caused this error, the entire virtual disk copy is resynchronized.

CMMVC6292E The command failed because a repair action is in progress for the virtual disk (VDisk).

Explanation

You cannot submit this command while a repair action is in progress for the virtual disk (VDisk).

Action

Use the `svcinfo lsrepairvdiskcopyprogress` command to view the repair progress. Wait for the virtual disk repair process to complete. If you want the repair process to complete more quickly, increase the rate by submitting an `svctask chvdisk` command. Once the repair action has completed, resubmit the command that caused this error.

CMMVC6296E One or more managed disks (MDisks) have failed validation tests. The first failing MDisk ID is *MDISK_ID*.

Explanation

When you add a managed MDisk to an MDisk group, the new MDisk is validated to ensure that adding it to the MDisk group will not adversely impact the MDisk group status. Either the current status of the MDisk has not allowed the validation to be performed, or the validation has failed. Note: You cannot add Read Only or faulty MDisks to an MDisk group.

Action

- If the MDisk identity has changed since it was last discovered, submit the command-line interface command `svctask detectmdisk`, which might correct the problem.
- Check switch zoning and logical unit presentation on the controller to ensure that the MDisk is physically and logically connected to all of the nodes in this cluster.
- Ensure that the controller settings are correct and that the MDisk logical unit is correctly configured.
- Ensure that the MDisk logical unit state is one that passes the validation. A Read Only or faulty MDisk fails the validation.
- View the cluster error log for more information about the failed validation.

CMMVC6297E One or more managed disks (MDisks) have timed out during validation tests. The first failing MDisk ID is *MDISK_ID*.**Explanation**

When you add a managed MDisk to an MDisk group, the new MDisk is validated to ensure that adding it to the MDisk group will not adversely impact the MDisk group status. The current status of the MDisk permits the validation to be initiated, but the allotted time for the validation process elapsed before the validation process had completed. Note: You cannot add Read Only or faulty MDisks to an MDisk group.

Action

- Ensure that the controller settings are correct and that the MDisk logical unit is correctly configured.
- Ensure that the MDisk logical unit state is one that passes the validation. A Read Only or faulty MDisk fails the validation.
- Check the fibre-channel fabric and storage controller for faults that might reduce the reliability of cluster communication with the MDisk.
- View the cluster error log for more information about the failed validation.

CMMVC6298E The command failed because a target VDisk has dependent FlashCopy mappings.**Explanation**

The target VDisk of the FlashCopy mapping, or the target VDisk 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 VDisk.

Action

Use the `lsvdiskdependentmaps` command and specify the target VDisk to determine which FlashCopy mappings are dependent on the target VDisk. Either wait for these mappings to reach the `idle_or_copied` state, or stop these mappings. Resubmit the command that produced this error.

CMMVC6299E The create failed because the source and target VDIs are members of FlashCopy mappings that have different grain sizes.**Explanation**

All FlashCopy mappings that are in a tree of connected mappings must have the same grain size. The new FlashCopy mapping that you attempted to create would have linked two existing trees that have different grain sizes.

Action

You have three options. The first option is to resubmit the command and specify a different source or target VDisk. The second option is to delete all of the existing mappings that contain the source VDisk and resubmit the command. The third option is to delete all of the existing mappings that contain the target VDisk and resubmit the command.

CMMVC6300E The create failed because the source and target VDisks are members of FlashCopy mappings that belong to different I/O groups.

Explanation

All FlashCopy mappings in a tree of connected mappings must be in the same I/O group. The new FlashCopy mapping that you attempted to create would have linked two existing trees that are in different I/O groups.

Action

You have three options. The first option is to resubmit the command and specify a different source or target VDisk. The second option is to delete all of the existing mappings that contain the source VDisk and resubmit the command. The third option is to delete all of the existing mappings that contain the target VDisk 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.

Action

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 VDisk or the target VDisk, 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.

Action

You have two options. The first option is to resubmit the command and specify a different source or target VDisk. The second option is to delete a sufficient number of the existing FlashCopy mappings in which either the source or the target VDisk 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 VDisk cannot be both the source and the target in a FlashCopy mapping. The FlashCopy mapping was not created because you have specified the same VDisk as both the source and the target.

Action

Resubmit the command and specify source and target VDisks that are not identical.

CMMVC6304E The create failed because the source VDisk does not exist.

Explanation

You must specify an existing VDisk as the source of a FlashCopy mapping. The FlashCopy mapping was not created because the source VDisk that you specified does not exist.

Action

Either create the source VDisk that you specified and resubmit the command, or resubmit the command and specify an existing VDisk as the source.

CMMVC6305E The create failed because the target VDisk does not exist.

Explanation

You must specify an existing VDisk as the target of a FlashCopy mapping. The FlashCopy mapping was not created because the target VDisk that you specified does not exist.

Action

Either create the target VDisk that you specified and resubmit the command, or resubmit the command and specify an existing VDisk as the target.

CMMVC6306E The create failed because the source VDisk is the member of a FlashCopy mapping whose grain size is different to that specified.

Explanation

All FlashCopy mappings that are in a tree of connected mappings must have the same grain size. The FlashCopy mapping was not created because the source VDisk that you specified is either the source or the target VDisk 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.

Action

You have two options. The first option is to delete all of the FlashCopy mappings that contain the source VDisk that you specified where the grain size of the FlashCopy mapping is different from the grain size that you specified, and resubmit the command. The second option is to resubmit the command and do not specify the grain size attribute.

CMMVC6307E The create failed because the target VDisk is the member of a FlashCopy mapping whose grain size is different to that specified.

Explanation

All FlashCopy mappings that are in a tree of connected mappings must have the same grain size. The FlashCopy mapping was not created because the target VDisk that you specified is either the source or the target VDisk 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.

Action

You have two options. The first option is to delete all of the FlashCopy mappings that contain the target VDisk that you specified where the grain size of the FlashCopy mapping is different from the grain size that you specified, and resubmit the command. The second option is to resubmit the command and do not specify the grain size attribute.

CMMVC6308E The create failed because the source VDisk is the member of a FlashCopy mapping whose I/O 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 VDisk that you specified is the source or target VDisk in another FlashCopy mapping and the I/O group of the other FlashCopy mapping is different from the I/O group that you specified.

Action

You have two options. The first option is to delete all of the FlashCopy mappings that contain the source VDisk that you specified where the FlashCopy mapping is in a different I/O group from the I/O group that you specified, and resubmit the command. The second option is to resubmit the command and do not specify the I/O group attribute. If you perform the second option, the default value of the I/O group attribute is used.

CMMVC6309E The create failed because the target VDisk is the member of a FlashCopy mapping whose I/O 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 VDisk that you specified is the source or target VDisk in another FlashCopy mapping and the I/O group of the other FlashCopy mapping is different from the I/O group that you specified.

Action

You have two options. The first option is to delete all of the FlashCopy mappings that contain the target VDisk 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.

Action

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 VDisk 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 VDisk of the FlashCopy mapping that you attempted to create is already the target VDisk of a FlashCopy mapping in the consistency group that you specified.

Action

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 VDisk 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 VDisk of the FlashCopy mapping that you attempted to create is already the source VDisk of a FlashCopy mapping in the consistency group that you specified.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.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: 0:0:0:0:0:0:13.1.68.3, which can be compressed to ::13.1.68.3

Action

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.

Action

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.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

Action

Specify a valid IPv6 gateway address, and resubmit the task.

CMMVC6326E The command has failed because the IPv4 service mode address is not valid.

Explanation

The valid IPv4 address format is d.d.d.d, where d is a decimal value from 0-255.

Action

Specify a valid IPv4 service mode address, and resubmit the task.

CMMVC6327E The command has failed because the IPv6 service mode address is not valid.

Explanation

Valid IPv6 address formats are:

- 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

Action

Specify a valid IPv6 service mode 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: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

Action

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: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

Action

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.

Action

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.

Action

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 e-mail 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.

Action

Either configure the cluster to have an IPv6 cluster management address or use an e-mail 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 e-mail 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.

Action

Either configure the cluster to have an IPv4 cluster management address or use an e-mail 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 e-mail port number supplied is invalid.

Explanation

The value that you entered for an e-mail port number is not a valid e-mail port number.

Action

Specify a valid e-mail 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.

Action

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 space-efficient VDisk copy.

Action

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 space-efficient VDisk copy but you have entered an incorrect value for the `-warning` parameter. The value can either be a percentage of the VDisk capacity or an absolute value that is a multiple of 512 bytes.

Action

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 space-efficient VDisk copy but you have entered an incorrect value for the `-warning` parameter. The warning value cannot be greater than the VDisk capacity.

Action

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 space-efficient VDisk but you did not set the `-size` parameter.

Action

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 space-efficient VDisk copy but you have entered an incorrect value for the `-rsize` parameter. All sizes must be integer multiples of 512 bytes.

Action

Resubmit the command using a supported `-rsize` parameter value.

CMMVC6341E The action failed because the virtual disk (VDisk) copy is not space-efficient.

Explanation

You are attempting to run a command that is valid only for space-efficient VDIs on a VDisk that is not space-efficient.

Action

Specify a space-efficient VDisk, 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 space-efficient VDisk copy, but the command cannot be initiated because it would make the real size less than the size that is currently used.

Action

Determine the used size of the VDisk 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 space-efficient VDisk copy, but the command cannot be initiated because it would make the real size less than zero.

Action

Determine the real size of the VDisk copy, and resubmit the command using a supported `-rsize` parameter value.

CMMVC6344E The repair operation cannot start because the space-efficient virtual disk (VDisk) copy is already being repaired.

Explanation

You are attempting to repair a space-efficient VDisk copy, but the copy is already being repaired.

Action

Specify the correct VDisk and copy parameters, and resubmit the command.

CMMVC6345E The repair operation cannot start because the space-efficient virtual disk (VDisk) copy was created using -import but the cluster could not recognize its format.

Explanation

You are attempting to repair a space-efficient VDisk copy that is reporting corrupt metadata. The cluster cannot repair the VDisk copy because it was not recognized as a valid space-efficient disk when it was imported into this cluster. The most probable cause is that the wrong MDisk was used when the VDisk copy was imported.

Action

Delete the VDisk 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 space-efficient VDisk copy that is reporting corrupt metadata. The cluster cannot repair the VDisk copy because although it was recognized as a valid space-efficient disk when it was imported into this cluster, the real size allocated to the VDisk copy is too small. The most probable cause is that the incorrect value was supplied with -rsize parameter when the VDisk copy was imported.

Action

Delete the VDisk copy. Resubmit the import operation either using a larger value for -rsize, or supplying the -rsize parameter without a value to let the system choose a real size.

CMMVC6347E The specific upgrade package cannot be installed on this hardware level.

Explanation

The version of software that you are attempting to install does not support the hardware level of the configuration node.

Action

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.

Action

Ensure that you specify a supported combination of parameters and parameter values, and resubmit the task.

CMMVC6349E The command failed because the virtual disk cache has been lost. Use the force flag to acknowledge this.

Explanation

You are attempting to move a VDisk that has lost cache data between I/O groups and did not specify the -force flag. You must specify the -force flag when you move a VDisk that has lost cache data between I/O groups.

Action

Specify the -force flag, and resubmit the command.

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 virtual disk mirroring in the I/O Group.

Action

Perform one of the following actions:

- Submit an svctask chiogrp command to increase the bitmap space.
- Remove virtual disk 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 virtual disks (VDisks) are supported for this command.

Action

Perform one of the following actions:

- Submit the appropriate command for a VDisk that is not mirrored.

- Submit an `svctask addvdiskcopy` command to add a copy to the virtual disk, 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 virtual disk (VDisk).

Action

Submit an `svctask rmvdiskcopy` or `svctask splitvdiskcopy` command to decrease the number of virtual disk 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.

Action

Submit an `svcinfo lsvdiskcopy` command to show all of the available copies for this virtual disk. 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.

Action

Use the `svcinfo lsvdisksyncprogress` command to view the synchronization status. Wait for the copy to synchronize. If you want the synchronization process to complete more quickly, increase the rate by submitting an `svctask chvdisk` command. When the copy is synchronized, resubmit the command that caused this error.

CMMVC6355E The command failed because an image mode copy is not synchronized and `-force` was not specified.

Explanation

When you specify an image mode copy for this command, the copy must be synchronized unless you also specify the `-force` parameter.

Action

Perform one of the following actions:

- Use the `svcinfo lsvdisksyncprogress` command to view the synchronization status. Wait for the copy to synchronize. If you want the synchronization process to complete more quickly, increase the rate by submitting an `svctask chvdisk` command. When the copy is synchronized, resubmit the command that caused this error.
- Resubmit the command and specify the `-force` parameter.

NOTE: When you specify the `-force` parameter with the command that caused this error, the image mode copy is no longer guaranteed to have the correct virtual disk data.

CMMVC6356E The command failed because a copy is not synchronized and `-force` was not specified.

Explanation

When you specify a copy for this command, the copy must be synchronized unless you also specify the `-force` parameter.

Action

Perform one of the following actions:

- Use the `svcinfo lsvdisksyncprogress` command to view the synchronization status. Wait for the copy to synchronize. If you want the synchronization process to complete more quickly, increase the rate by submitting an `svctask chvdisk` command. When the copy is synchronized, resubmit the command that caused this error.
- Resubmit the command and specify the `-force` parameter.

NOTE: When you specify the `-force` parameter with the command that caused this error, the entire virtual disk 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.

Action

Perform one of the following actions:

- Use the `svcinfo lsvdisksyncprogress` command to view the synchronization status. Wait for the copy to synchronize. If you want the synchronization process to complete more quickly, increase the rate by submitting an `svctask chvdisk` command. When the copy is synchronized, resubmit the command that caused this error.
- Resubmit the command and specify the `-force` parameter.

NOTE: When you specify the `-force` parameter with the command that caused this error, the created virtual disk is no longer guaranteed to have identical data to the original virtual disk 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.

Action

Use the `svcinfo 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 an `svctask chvdisk` command. When the copy has synchronized, resubmit the command that caused this error.

CMMVC6359E The command failed because there are insufficient online synchronized copies.

Explanation

This error occurs when at least one of the virtual disk copies is offline.

Action

Fix all of the errors that are associated with the virtual disk 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 virtual disk (VDisk).

Action

Use the `svcinfo lsvdisk` command to obtain the virtual disk 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 VDisk or MDisk, but the LBA is too large and does not exist on the disk.

Action

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.

Action

Resolve any MDisk or fabric error log entries, and resubmit the command.

CMMVC6366E One or more nodes in the cluster has hardware that is not supported by the new software package.

Explanation

The version of software that you are attempting to install does not support the hardware in at least one node in the cluster.

Action

Check the release notes for the version of software that you want to install. Upgrade hardware so that all of the hardware in the cluster is supported by the new version of software, and resubmit the task.

CMMVC6367E A remote cluster is running software that is incompatible with the new software package.

Explanation

The version of software that you are attempting to install on the local cluster does not support the version of software that is installed on the remote cluster.

Action

Check the release notes for the version of software that you want to install. Perform one of the following actions:

- Upgrade the software on the remote cluster to a version that is supported by the version of software that you want to install on the local cluster before you upgrade the software on the local cluster.
- Delete the cluster partnership to stop all remote copy relationships between the clusters, and resubmit the task.

CMMVC6368E The new software package might be incompatible with the remote cluster.

Explanation

The software version compatibility between clusters cannot be checked because the remote cluster is not accessible.

Action

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.

CMMVC6369E 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.

Action

Upgrade the FlashCopy storage capacity license to prevent recurrence of this warning message.

CMMVC6370E 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.

Action

Upgrade the Remote Copy storage capacity license to prevent recurrence of this warning message.

CMMVC6372E 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.

Action

Upgrade the virtualized storage capacity license to prevent recurrence of this warning message.

CMMVC6373E 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.

Action

Upgrade the virtualized storage capacity license to prevent recurrence of this warning message.

CMMVC6374E 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.

Action

Upgrade the FlashCopy storage capacity license to prevent recurrence of this warning message.

CMMVC6375E 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.

Action

Upgrade the Remote Copy storage capacity license to prevent recurrence of this warning message.

CMMVC6394E The command failed because an attempt to make the virtual disk cache empty took too long.

Explanation

The failed command must empty the virtual disk cache before attempting the requested action to ensure that data is preserved. The empty VDisk 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 virtual disk cache.

The storage associated with the virtual disk is probably overloaded.

Action

Wait a few minutes to allow the virtual disk cache to empty. Resubmit the command.

Alternatively, you can use the `-force` parameter, if the command supports the `-force` parameter, to bypass the empty virtual disk cache subtask. However, specifying the `-force` parameter will discard cache data for the virtual disk. Only use the `-force` flag with this command if you do not intend to use the existing contents of the virtual disk.

In addition to the above actions, investigate the performance of the network storage devices associated with this virtual disk. 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.

Action

Check for errors in the error log. Follow the directed maintenance procedures to resolve the problem.

CMMVC6400E The command failed because a specified managed disk (MDisk) is already in use.

Explanation

You cannot specify an MDisk for this command if it is already in a managed disk group or is being used as an image mode VDisk.

Action

Specify an MDisk that is not being used as an image mode VDisk and is not in a managed disk group, 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 MDisk group.

Action

Ensure that all of the MDisks that you specify are in the same MDisk group, 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 MDisk group. At least one of the source MDisks that you have specified in the command is not in the required MDisk group.

Action

Ensure that all of the MDisks that you specify are in the MDisk group 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 MDisk group. At least one of the target MDisks that you have specified in the command is not in the required MDisk group.

Action

Ensure that all of the MDisks that you specify are in the MDisk group 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 MDisk groups that you specify for a cross MDisk group migration must be different.

Action

Ensure that the source and target MDisk groups that you specify for a cross MDisk group 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 VDisk and more than one VDisk copy exists.

Action

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 MDisk groups that you have specified in the parameter list does not exist.

Action

Ensure that each of the MDisk groups that you specify exists, and resubmit the command.

CMMVC6407E The command failed because the managed disk group is invalid.

Explanation

At least one MDisk group ID is above the maximum value that is available for the system.

Action

Ensure that each MDisk group 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 MDisk groups that is consistent with the other parameters and parameter values that you specify with the command.

Action

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 MDisk groups that is consistent with the other parameters and parameter values that you specify with the command.

Action

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.

Action

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.

Action

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 an MDisk group extent size that is larger the maximum size.

Action

Specify an MDisk group 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.

Action

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.

Action

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 an MDisk group warning threshold that is equal to or greater than the minimum size.

Action

Specify an MDisk group 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 an MDisk group warning threshold size that is equal to or less than the size of the MDisk group when all of the MDisks have been added, or you must specify an MDisk group warning percentage that is equal to or less than the maximum warning threshold percentage.

Action

Specify valid values for the MDisk group 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 MDisk Group.

Action

Ensure that there is at least one MDisk defined for the MDisk group 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 VDisk that is being resized.

Action

Wait for the resize VDisk 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.

Action

Wait for the delete MDisk operation to complete. Do not include any MDisks that have been deleted in the list of MDisks that you specify, and resubmit the command.

CMMVC6423E The Send Inventory e-mail operation failed because e-mail is not started.

Explanation

The send inventory e-mail functionality has been enabled but the e-mail service has not been started.

Action

Disable the send inventory e-mail functionality or start the e-mail service.

CMMVC6424E The Send Inventory e-mail operation failed because there are no inventory e-mail users.

Explanation

The send inventory functionality has been enabled but no e-mail users with the ability to receive inventory e-mails have been created.

Action

Either turn off the send inventory e-mail functionality or create an e-mail user account that is capable of receiving inventory e-mails. Refer to the documentation for the mke-mailuser command for help on creating e-mail 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.

Action

Not applicable.

CMMVC6426E The command failed because a specified managed disk (MDisk) is already in use.

Explanation

You cannot specify an MDisk that is already configured as an image mode VDisk.

Action

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 VDisk task requires that all of the MDisks that you specify must be in the same MDisk group.

Action

Ensure that all of the MDisks that you specify are in the same MDisk group, 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 MDisk groups that you specify must be in the same MDisk group.

Action

Ensure that all of the source MDisk groups that you specify are in the same MDisk group, 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 MDisk groups that you specify must be in the same MDisk group.

Action

Ensure that all of the target MDisk groups that you specify are in the same MDisk group, and resubmit the task.

CMMVC6430E The command failed because the target and source managed disk groups must be different.

Explanation

The cross MDisk group migration task does not support specifying the same MDisk group to be both the source and target MDisk group.

Action

Specify a source MDisk group and a target MDisk group 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 VDisk and there is more than one copy, you must specify which copy to use as the target copy.

Action

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 MDisk groups that you specify must already exist.

Action

Ensure that all of the MDisk groups that you specify already exist, and resubmit the task.

CMMVC6433E The command failed because the managed disk group is invalid.

Explanation

All of the MDisk group IDs that you specify must have a value that is less than or equal to the maximum supported MDisk group ID value.

Action

Ensure that all MDisk groups have supported ID values. Ensure that all of the MDisk groups 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 MDisk groups than the number that you have specified.

Action

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 MDisk groups than the number that you have specified.

Action

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.

Action

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.

Action

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 MDisk group extent size that you have specified is greater than the supported maximum value.

Action

Specify a supported MDisk group 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.

Action

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.

Action

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 MDisk group warning threshold is less than the minimum supported value.

Action

Specify a supported value for the MDisk group 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 MDisk group warning percentage is greater than the maximum supported value, or the MDisk group warning disk size is greater than the MDisk group capacity.

Action

Specify supported values for MDisk group 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 an MDisk group warning threshold percentage, you must specify an MDisk group that contains at least one MDisk and you must specify a supported value for the MDisk group warning threshold percentage.

Action

Either do not specify an MDisk group warning threshold percentage, or specify a supported value for the MDisk group warning threshold percentage and specify an

MDisk group 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 VDisk that is being resized when you submit this task.

Action

Wait for the resize VDisk task to complete. You can specify the same VDisk and resubmit this task only after the resize VDisk 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.

Action

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 MDisk group and the extent size of the target MDisk group must be identical.

Action

If you want to resubmit this command, ensure that the source and target MDisk groups have the same extent size. If you want to move a VDisk to an MDisk group 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 VDisk that is being migrated.

Action

Either wait until the VDisk migration process has completed and resubmit the task, or specify a VDisk that is not being migrated and resubmit the task.

CMMVC6448E Deleting this node will cause data loss for resources associated with the IO group of this node.

Explanation

This node contains resources which are vital to the IO 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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

Move all quorum disks from the controller to a different storage system using the `svctask setquorum` command, and resubmit this task.

CMMVC6458E The specified controller cannot support quorum disks.**Explanation**

The controller type of the controller that you specified does not support quorum disks.

Action

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 VDisk to itself. The mkrcrelationship command requires that you specify two different VDIsks for the master and auxiliary positions. These can be two VDIsks in the local cluster, or a VDisk in each of two different clusters.

Action

Specify a master VDisk and an auxiliary VDisk 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 managed disk group.

Action

- If you submitted the rmmddisk command and specified a regular MDisk, determine the managed disk group to which the source MDisk is defined, and follow the procedure for bringing the managed disk group online. There will be an entry in the error log for the corresponding managed disk group.
- If you submitted the rmmddisk 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 error log for the corresponding MDisks.
- If you submitted a command to migrate a copy of an image mode VDisk, determine the corresponding source MDisk and follow the procedure for diagnosing problems with MDisks. There will be an entry in the error log for the corresponding MDisks.
- If you submitted any other command to migrate a VDisk copy, determine the managed disk group to which the VDisk is defined, and follow the procedure for bringing the managed disk group online. There will be an entry in the error log for the corresponding managed disk group.

CMMVC6461E The command failed because starting the migration will result in VDIs going offline in the source managed disk group.

Explanation

A migration from an image mode VDisk will use the source managed disk group and the source managed disk group assumes the combined state of the image mode MDisk and the managed disk group. If the online or offline states of the image mode MDisk and the managed disk group are different on different nodes, the source VDisk might go offline or all of the VDIs in the source managed disk group might go offline.

Action

For each node, note the online or offline states of the source VDisk and the source managed disk group. 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 VDIs might go offline.

CMMVC6462E The command failed because starting the migration will result in VDIs going offline because the target managed disk group is offline.

Explanation

The migration process assigns the VDisk an online or offline state based on the states of the source and target managed disk groups. In this case, based on the offline state of the target managed disk group the VDisk 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 error log for the corresponding managed disk group.

Action

For each node, note the online or offline state of the source and target managed disk groups. For each node, if one of these two managed disk groups is online and the other is offline, bring online whichever managed disk group is offline. Taking the online managed disk group offline is not recommended because other VDIs might go offline.

CMMVC6463E The command failed because Starting the migration will result in VDIs going offline because a target MDisk is offline.

Explanation

The VDisk is currently online. The migration process assigns the VDisk an online or offline state based on the states of the source and target MDisks. In this case, based on the offline state of the target MDisk, the VDisk would have been taken offline. The task cannot be initiated because this action is not supported.

Action

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 VDisk size task is in progress.

Action

Wait until the Change VDisk 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 VDisk size task is in progress.

Action

Wait until the Change VDisk 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 VDIsks that you have specified is defined. You cannot define a map that is exactly the same as a map that is already defined.

Action

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 VDisk in the same consistency group.

Action

Specify a target VDisk 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 VDisk 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 VDisk of another FlashCopy mapping that is being restored.

Action

Ensure that the target VDisk in the map that you are attempting to start or prepare is not the source VDisk 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.

Action

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 VDisk of another map that is in one of the following states: copying, stopping, suspended, prepared or preparing.

Action

Ensure that the target VDisk in the map that you are attempting to start or prepare is not the target VDisk of another FlashCopy mapping that is in one of the unsupported states when you submit this task.

CMMVC6471E The Create cluster partnership task cannot be initiated because a cluster in the existing partnership has a downlevel software version that does not support this configuration.

Explanation

One scenario in which this error occurs is when a cluster at version 5.1.0 or later is partnered to a cluster at version 4.3.1 or earlier and you attempt to create another partnership to a cluster at version 5.1.0 to implement multiple cluster mirroring. Software version 4.3.1 does not support multiple cluster mirroring, so adding a partnership to a third cluster is not supported while at least one cluster in the current partnership is at version 4.3.1 or earlier.

Action

Either upgrade the downlevel cluster software version to a version that supports this task or remove the partnership to the cluster that has the downlevel software version, and resubmit the task.

CMMVC6472E The Create cluster partnership task cannot be initiated because the remote cluster with which you are attempting to create a partnership has a downlevel software version that does not support this configuration.

Explanation

The software versions of the clusters in the existing partnership do not support a partnership with a cluster that has the software version of the remote cluster with which you are attempting to create a partnership. If a cluster at version 5.1.0 or later is already in a partnership with another cluster at version 5.1.0 or later, you can only add a partnership to a cluster at version 5.1.0 or later, and cannot add a partnership to a cluster at version 4.3.1 or earlier. If a cluster at version 5.1.0 or later is already in a partnership with another cluster at version 4.3.1 or earlier, you cannot add another partnership while the partnership with the cluster at version 4.3.1 exists. If a cluster is not in a partnership, you can create a partnership between it and a cluster at any version. One scenario in which this error occurs is when you attempt to add a partnership with a remote cluster at version 4.3.1 or earlier to a cluster at software version 5.1.0 or later that is already in partnership with another cluster at software version 5.1.0 or later.

Action

Either upgrade the downlevel cluster software version to a version that supports this task or remove all existing partnerships from the cluster to which you want to partner the cluster that has the downlevel software version, and resubmit the task.

CMMVC6473E The partnership task cannot be initiated because the supported maximum number of accessible remote clusters would be exceeded.

Explanation

With multiple cluster mirroring, you can build a configuration of a chain of clusters. There is a limit to the number of clusters that you can configure in the chain. The task would have resulted in exceeding the supported maximum number of clusters in a chain.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

Either specify a user group or specify the use of the remote authentication service, but not both, and resubmit the task.

CMMVC6491E The task cannot be initiated because you have specified the use of the remote authentication service but you have not specified an ssh key and password.

Explanation

An ssh key and password are both required to use the remote authentication service.

Action

Specify a valid ssh key and a valid password when you submit this task.

CMMVC6492E The task cannot be initiated because you have specified a local user but you have not specified a user group.

Explanation

You must specify a user group when you specify a local user for this task.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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.

Action

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 disk (VDisk) are the same.

Explanation

The action failed because the source and destination VDisk are the same.

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.

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.

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.

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.

Action

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 San Volume Controller (SVC) is configured to use an authentication service to control which users are authorized to access the cluster. An error has occurred while the SVC was attempting to contact the authentication service. The error is probably the result of an incorrect configuration, either of the SVC or of the authentication service. This error occurs if the SSL certificate, user name or password is incorrect.

Action

Ensure that the authentication service is functioning properly. Ensure that the SVC 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 San Volume Controller (SVC) is configured to use an authentication service to control which users are authorized to access the cluster. A timeout has occurred while the SVC 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.

Action

Ensure that the SVC authentication service configuration is correct. Ensure that the Ethernet network between the SVC 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 San Volume Controller (SVC) 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 SVC to refresh its authentication cache. You can force the refresh using the SVC console View Cluster Properties, Remote Authentication panel or by submitting the Command-Line Interface command 'svctask chauthservice -refresh'.

Action

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 SVC to refresh its authentication cache.

If the user name that you are using also has a password configured on the SVC 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 San Volume Controller (SVC) 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.

Action

Either acquire a new authentication token or log in with a user name and password, and resubmit the task.

CMMVC6509E The task has failed because the user name is not configured on the San Volume Controller (SVC) cluster.

Explanation

If the user name is defined on an authentication service and you want to use that service for SVC authentication, you must configure the SVC to use that authentication service.

Action

Ensure that you are using the correct user name.

If the user name is not configured on the SVC and you want to use the SVC to authenticate, create a new user with the user name that you want to use on the SVC.

If the user name is defined on an authentication service and you want to use that service for SVC authentication, configure the SVC to use that authentication service.

Resubmit the task.

CMMVC6510E The task has failed because the password is not correct for the user name that you are using.

Explanation

The password that you are using does not match the password that is configured on the San Volume Controller for the user name that you are using.

Action

Enter the correct password, and resubmit the task.

CMMVC6511E The task has failed because San Volume Controller (SVC) 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 SVC is not configured to use an authentication service or the function is not enabled.

Action

If you want to use an authentication service, configure the SVC to use the service.

If you do not want to use an authentication service, modify the configuration of the user name on the SVC 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.

Action

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.

Action

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.

Action

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.

Action

Either create additional quorum disks or change the configuration so that none of the quorum disks is in the Excluded state. Ensure that none of the quorum disks are in the Excluded state, and resubmit the task.

CMMVC6516E The command has failed because an IPv4 cluster address cannot be removed while remote IPv4 services are configured.

Explanation

The configured cluster IP address protocols determine whether IPv4 or IPv6 or both are enabled on the cluster. If a cluster does not have an IPv4 cluster address the IPv4 protocol stack will not be enabled, and therefore remote services such as e-mail servers or SNMP servers cannot be accessed through an IPv4 address.

Action

If you can only access the service through an IPv4 address and you need to continue to use the service, you will also have to continue to specify an IPv4 cluster address even if you do not intend to manage your cluster through this address.

Otherwise, re-configure the cluster so that all remote services use only IPv6 addresses, and resubmit the task to remove the IPv4 cluster address.

CMMVC6517E The command has failed because an IPv6 cluster address cannot be removed while remote IPv6 services are configured.

Explanation

The configured cluster IP address protocols determine whether IPv4 or IPv6 or both are enabled on the cluster. If a cluster does not have an IPv6 cluster address the IPv6 protocol stack will not be enabled, and therefore remote services such as e-mail servers or SNMP servers cannot be accessed through an IPv6 address.

Action

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 SAN Volume Controller (SVC) cluster.

Explanation

The SVC 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 SVC cluster.

Action

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 SVC cluster.
3. Ensure that at least one user group that is defined for the user on both the authentication service and the SVC 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.

Action

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 SAN Volume Controller (SVC). The current user is defined on an authentication service, and the SVC 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 SVC by using an ssh key, you must define the current user on the SVC 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 SVC.

Action

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 SVC by using an ssh key, define your user account on the SVC 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 SVC.

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.

Action

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.

Action

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, - _ : [] . ~ / %.

Action

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: * : , \ " ' %

Action

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.

Action

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 MDisk group for each VDisk copy that you request.

Action

Specify the same number of unique MDisk groups as the number of VDisk copies that you request, 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.

Action

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.

Action

Ensure that a local MDisk is available, and resubmit the command. To make a local MDisk available for this task, either delete an array on an existing local MDisk or remove a SAN attached MDisk and configure a local MDisk.

CMMVC6530E The command cannot be initiated because the maximum supported number of arrays already exists.

Explanation

The cluster already has the maximum number of arrays that it can support. The command attempted to add a new array.

Action

Remove an array that is no longer needed, and resubmit the command.

CMMVC6531E The command cannot be initiated because array scrubbing is not supported for the RAID geometry of the array that you have specified.

Explanation

Array scrubbing is a background process that detects inconsistency between related parity data in RAID geometries that use multiple parity algorithms. RAID6 is an example of a geometry that supports array scrubbing.

Action

Specify an array that has a RAID geometry that supports array scrubbing, 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.

Action

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 VDisk mirrors or Copy Services relationships in the IO 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.

Action

Select an array member that has an associated LDisk, and resubmit the command.

CMMVC6534E The command cannot be initiated because the LDisk that you have specified does not exist.

Explanation

You have specified an LDisk ID that is not defined. Use the 'svcinfo lsldisk' command to display existing LDisk IDs.

Action

Specify an existing LDisk ID, and resubmit the command.

CMMVC6535E The command cannot be initiated because there are insufficient available LDisks to configure an array using the RAID geometry that you have specified.

Explanation

Each RAID geometry requires a minimum number of available LDisks in order to configure an array using that geometry.

Action

Ensure that there are a sufficient number of available LDisks to accommodate the RAID geometry that you specify, and resubmit the command. You might want to specify a different number of LDisks or a different RAID geometry.

CMMVC6536E The command cannot be initiated because you have specified more LDisks than the specified RAID geometry permits.

Explanation

The number of LDisks that you specify must be within the supported range of number of LDisks that is supported for the RAID geometry that you specify.

Action

Specify a number of available LDisks that is supported for the RAID geometry that you specify, and resubmit the command.

CMMVC6537E The command cannot be initiated because the LDisk that you have specified has a Use property that is not supported for the task.

Explanation

You can submit the 'svcinfo lsldisk' command to display the Use property of existing LDisks. Examples of values of the Use property are:

- None. An LDisk that is not configured to be used.

- Candidate. An LDisk that is available to be configured as part of an array.
- Spare. An LDisk that is configured as part of an array and available to be used if another member of the array fails.
- Member. An LDisk that is configured as member of an array that can contain current data.
- Failed. An LDisk that was formerly a member of an array, but that has been rejected because the LDisk is no longer usable.

Action

Consult the command documentation to determine what LDisk Use values are supported for this command. Ensure that you select an LDisk 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 LDisks that you have specified has a Use property that is not Candidate.

Explanation

Every LDisk that you specify for this command must have a Use property of Candidate. You can submit the 'svcinfo lsldisk' command to display the Use property of existing LDisks.

Action

Ensure that all of the LDisks 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. The RAID state of the array is Degraded.

Explanation

The array must have sufficient redundancy when you submit this command. Because the RAID state of the array is Degraded, the task that you have requested would have taken the array Offline.

Action

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 required a larger number of grains than the supported maximum for the specified grain size.

Action

Either increase the grain size, decrease the requested virtual capacity of the VDisk, or both, and resubmit the task.

CMMVC6541E The task cannot be initiated because the virtual capacity that you have requested for the space-efficient VDisk is larger than the maximum capacity that is supported for the extent size.

Explanation

The extent size of the MDisk group 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 space-efficient VDisk.

Action

Either select a different MDisk group 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 MDisk group 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.

Action

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.

Action

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.

Action

Resubmit the task when the direct-attached managed drive is less busy.

CMMVC6545E The Apply Drive Software task has failed to access the software download image.

Explanation

Either the image file cannot be read, the validation signature is incorrect, the drive type or firmware type is not correct, or the image file has been corrupted.

Action

Reinstall the firmware download image, and resubmit the task. If the problem persists, contact IBM technical support for assistance.

CMMVC6546E A device error was detected during the Apply Drive Software task.

Explanation

The task might have succeeded.

Action

View the error in the error log. Determine the firmware level from the VPD of the node. If the VPD does not show that the downloaded firmware version is installed, resubmit the task.

CMMVC6547W The Download FPGA firmware task has been initiated. The MDisk remains Offline while the task is in progress. Do not remove power from the drive or node while the task in is progress.

Explanation

The task might take approximately fifteen minutes to complete. When the task completes, the drive status changes to Online automatically.

Action

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 Download FPGA firmware task cannot be initiated because the drive that you have selected is a managed disk.

Explanation

You must specify an unmanaged disk when you submit this task. You cannot specify a drive that is a managed disk when you submit this task. Because this task can take fifteen to thirty minutes or more to complete, it is not supported for managed disks.

Action

Migrate all of the data from the managed disk, remove the MDisk from its MDisk group, and resubmit the task.

CMMVC6549E The Authentication task has failed because the authentication service URL that you have specified is not a valid URL.

Explanation

This error might be caused by the authentication service not operating correctly or by an incorrect URL being defined for the authentication service. You can use the 'svctask chauthservice' command to change the URL that is defined in the SVC for the authentication service.

Action

Ensure that the authentication service is operating correctly. Ensure that the authentication service URL that is defined in the SVC 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 SVC has a network address that cannot be resolved. You can use the 'svctask chauthservice' command to change the URL that is defined in the SVC for the authentication service.

Action

Ensure that the authentication service is operating correctly. Ensure that the authentication service URL that is defined in the SVC is correct. Ensure that the network connection between the SVC 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 SVC for authorization by the authentication service is not defined on the authentication service.

Explanation

The authentication service has refused an authentication request from the SVC. You can use the 'svctask chauthservice' command to change the user name or the password that is defined in the SVC for the authentication service.

Action

Ensure that the user name and password combination that is defined in the SVC 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 SVC. You can use the 'svctask chauthservice' command to set the SSL certificate that is defined in the SVC for the authentication service server.

Action

Ensure that the SSL configuration on the authentication service server is correct and that the SSL certificate that is defined in the SVC 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.

Action

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 SVC user name.

Explanation

The SVC user name cannot exceed 256 characters in length, and cannot contain any of the following characters:

- colon :
- percent sign %
- comma ,
- double quote "
- single quote '

Action

Change the definition of the user name in the remote authentication service so that it conforms to the SVC 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.

Action

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 SVC 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.

Action

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 SVC 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.

Action

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 SVC Command-Line Interface (CLI) command help for this command.

Explanation

You are being warned that this command might result in taking VDisks 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.

Action

1. Submit the `'lsnode dependantvdisks'` command to determine which VDisks will go Offline if you resubmit this command using the `-force` flag. If you received this message when you submitted the `'applysoftware'` command, you must submit the `'lsnode dependantvdisks'` command for every node in the cluster; for all other commands you must submit the `'lsnode dependantvdisks'` command for the node that you specified as a parameter in the command that generated this message.
2. This step is required because it is critically important that you understand the implications of using the `-force` flag for the specific command that you have submitted: Refer to the CLI command help to determine what safety precautions are bypassed when you use the `-force` flag. The ignored precautions differ, depending on the command.
3. If you want to bypass the safety precautions when you resubmit the command, you must use the `-force` flag.

CMMVC6559E The Add or Change e-mail 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.

Action

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.

Action

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.

Action

Check the state of the MDisks and complete any outstanding directed maintenance procedures (DMPs). If another set quorum active task might be in progress, wait for several minutes for that task to complete, and resubmit this task. If you have received this error when there is no other set quorum active task in progress, specify a different disk to replace the current active quorum disk and specify the same quorum index number, and resubmit this task.

CMMVC6562E The requested size exceeds the maximum supported value.

Explanation

You have submitted a command that has a size parameter and an associated unit option that has a default value of Megabytes (MB, 2e20 bytes) when the -unit option is not specified. The value that you have specified for the size parameter in combination with the specified or default unit value is greater than the maximum supported size of (2e64 - 1) bytes.

Action

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 e-mail 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 e-mail notification flags set to 'on': -error, -warning, or -info.

Action

Ensure that all of the users that you specify have at least one e-mail 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.

Action

Either specify a new password and resubmit the command, or first modify the password and then resubmit the command to designate remote authentication for this user.

CMMVC6565E The command has failed because the specified node is offline.

Explanation

This command requires that the status of the node that you specify is Online.

Action

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.

Action

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 has failed because no download images were found in the package file.

Explanation

The drive software upgrade package file was unpacked but no download software images were found in the package.

Action

Acquire a valid solid-state drive software upgrade package file, and resubmit the task using the new package file.

CMMVC6568E The Apply Drive Software task has failed because no download images were found in the package file for this drive type.

Explanation

The package file documentation lists the drive types for which there are images.

Action

Acquire a valid solid-state drive software upgrade package file that contains an image for this drive type, and resubmit the task using the new package file.

CMMVC6569E The Apply Drive Software task has failed because no download images were found in the package file of this software type.

Explanation

The package file documentation lists the drive types and software types for which there are images. The value of the `-type` parameter that you enter for software type is case-sensitive.

Action

Ensure that the value that you enter for the `-type` parameter exactly matches the software type that is contained in the solid-state drive software upgrade package file, and resubmit the command.

CMMVC6570E The command was not initiated because the cache mode of the virtual disk (VDisk) is already in the state that you had requested.

Explanation

You have issued a change VDisk cache mode command but requested the current mode, so there would not have been a change. Therefore, the command was ignored.

Action

List the VDisk 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 VDisk 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 VDisk, 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.

Action

Either follow service procedures to bring the I/O group online or specify the -force flag to force the change of the VDisk's cache mode, 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.

Action

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 VDisk 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 VDisk.

Action

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 VDisk 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 VDisk.

Action

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 VDisk 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 VDisk.

Action

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 VDisk 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 VDisk.

Action

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 VDisk 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 VDisk.

Action

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.

Action

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.

Action

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.

Action

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 iSCSI qualified name (IQN) is already assigned or is not valid.

Explanation

The cluster does not support duplicate IQNs. An IQN cannot contain a comma or leading or trailing spaces.

Action

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.

Action

Map the iSCSI host to at least one I/O group, and resubmit the command.

Appendix. Accessibility

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully.

Features

These are the major accessibility features in the SAN Volume Controller Console:

- You can use screen-reader software and a digital speech synthesizer to hear what is displayed on the screen. The following screen reader has been tested: Window-Eyes v6.1.
- You can operate all features using the keyboard instead of the mouse.
- When setting or changing an IP address on the SAN Volume Controller front panel, you can disable the fast increase and decrease address scrolling speed function of the up and down buttons to two seconds. This feature is documented in the topic that discusses initiating cluster creation from the front panel, which is located in the *IBM System Storage SAN Volume Controller Information Center* and the *IBM System Storage SAN Volume Controller Software Installation and Configuration Guide*.

Navigating by keyboard

You can use keys or key combinations to perform operations and initiate many menu actions that can also be done through mouse actions. You can navigate the SAN Volume Controller Console and help system from the keyboard by using the following key combinations:

- To traverse to the next link, button, or topic, press Tab inside a frame (page).
- To expand or collapse a tree node, press → or ←, respectively.
- To move to the next topic node, press V or Tab.
- To move to the previous topic node, press ^ or Shift+Tab.
- To scroll all the way up or down, press Home or End, respectively.
- To go back, press Alt+←.
- To go forward, press Alt+→.
- To go to the next frame, press Ctrl+Tab.
- To move to the previous frame, press Shift+Ctrl+Tab.
- To print the current page or active frame, press Ctrl+P.
- To select, press Enter.

Accessing the publications

You can find the HTML version of the IBM System Storage SAN Volume Controller information at the following Web site:

<http://publib.boulder.ibm.com/infocenter/svcic/v3r1m0/index.jsp>

You can access this information using screen-reader software and a digital speech synthesizer to hear what is displayed on the screen. JAWS version 10 has been tested.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.*

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
3-2-12, Roppongi, Minato-ku, Tokyo 106-8711 Japan*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATIONS "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Corporation
Almaden Research
650 Harry Road
Bldg 80, D3-304, Department 277
San Jose, CA 95120-6099
U.S.A.*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products may be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application

programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at Copyright and trademark information at www.ibm.com/legal/copytrade.shtml.

Adobe and the Adobe logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Intel, Intel logo, Intel Xeon, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Index

Special characters

-filtervalue argument 185

A

accessibility
 keyboard 545
 repeat rate of up and down
 buttons 545
 shortcut keys 545
addhostiogr command 87
addhostport command 88
addmdisk command 127
addnode command 48
addsshkey command 327
addvdiskcopy command 97
administrator role 27
applydisksoftware command 135
applysoftware command 79, 329
audit log commands
 catauditlog 11
 dumppauditlog 13
 lsauditlogdumps 14
 overview 11
authentication
 SSH logins 3

B

backup and restore commands 75
backup command 75
backup commands
 backup 75
 clear 76
 help 77

C

catauditlog command 11
caterrlog command 195
caterrlogbyseqnum command 197
cfgportip command 50
chauthservice command 17
chcluster command 51
chclusterip command 55
chcontroller commands 347
chcurrentuser command 20
chemail command 30
chemailserver command 31
chemailuser command 32
cherrstate command 80
chfconsistgrp command 141
chfcmap command 141
chhost command 89
chiogrp command 57
chlicense command 323
chmdisk command 136
chmdiskgrp command 128
chnode command 59

chpartnership command 155
chrconsistgrp command 156
chrrelationship command 156
chsnmpserver command 34
chsyslogserver command 35
chuser command 20
chusergrp command 22
chvdisk command 100
clear command 76
cleardumps command 60, 330
clearerrlog command 81
CLI (command-line interface)
 configuring PuTTY 5
 preparing SSH client systems 2
cluster commands
 addnode 48
 cfgportip 50
 chcluster 51
 chclusterip 55
 chiogrp 57
 chnode 59
 cleardumps 60
 cpdumps 62
 detectmdisk 63
 ping 64
 rmnode 65
 rmpportip 67
 setclustertime 68
 setpwdreset 69
 settimezone 70
 startstats 70
 stopcluster 72
 stopstats 74
cluster diagnostic and service-aid
 commands
 applysoftware 79
 cherrstate 80
 clearerrlog 81
 dumpperrlog 82
 finderr 82
 overview 79
 setevent 83
 setlocale 83
 svqueryclock 84
 writeserenum 84
 cluster error log
 displaying 195
 command-line interface (CLI)
 configuration 4
 configuring PuTTY 5
 preparing SSH clients 2
 Command-line interface messages
 overview 349
 commands
 addhostiogr 87
 addhostport 88
 addmdisk 127
 addnode 48
 addsshkey 327
 addvdiskcopy 97
 applydisksoftware 135

commands (*continued*)
 applysoftware 79, 329
 backup 75
 catauditlog 11
 caterrlog 195
 caterrlogbyseqnum 197
 cfgportip 50
 chauthservice 17
 chcluster 51
 chclusterip 55
 chcontroller 347
 chcurrentuser 20
 chemail 30
 chemailserver 31
 chemailuser 32
 cherrstate 80
 chfconsistgrp 141
 chfcmap 141
 chhost 89
 chiogrp 57
 chlicense 323
 chmdisk 136
 chmdiskgrp 128
 chnode 59
 chpartnership 155
 chrconsistgrp 156
 chrrelationship 156
 chsnmpserver 34
 chsyslogserver 35
 chuser 20
 chusergrp 22
 chvdisk 100
 clear 76
 cleardumps 60, 330
 clearerrlog 81
 cpdumps 62
 detectmdisk 63
 dumppauditlog 13
 dumpperrlog 82, 331
 dumppinternallog 325
 exit 332
 expandvdisksize 104
 finderr 82
 help 77
 includemdisk 136
 ls2145dumps 198, 333
 lsauditlogdumps 14
 lsaauth 29
 lscimomdumps 199, 334
 lscluster 201
 lsclustercandidate 204
 lsclusterip 205
 lsclustervpd 335
 lscontroller 207
 lscontrollerdependentvdisks 210
 lscopystatus 200
 lscurrentuser 210
 lsdiscoverystatus 211
 lsemailserver 212
 lsemailuser 213
 lserrlogbyfconsistgrp 214

commands (*continued*)

lserrlogbyfcmap 216
 lserrlogbyhost 217
 lserrlogbyiogrp 218
 lserrlogbymdisk 219
 lserrlogbymdiskgp 220
 lserrlogbynode 222
 lserrlogbyrconsistgrp 223
 lserrlogbyrcrelationship 224
 lserrlogbyvdisk 226
 lserrlogdumps 227, 336
 lsfabric 228
 lsfcconsistgrp 230
 lsfcmap 232
 lsfcmapcandidate 234
 lsfcmapdependentmaps 236
 lsfcmapprogress 235
 lsfeaturedumps 237, 337
 lsfreeextents 238
 lshbaportcandidate 239
 lshost 240
 lshostiogrp 242
 lshostvdiskmap 243
 lsiogrp 245
 lsiogrpcandidate 248
 lsiogrphost 247
 lsiostatsdumps 248, 338
 lsio tracedumps 249, 339
 lsiscsiauth 250
 lslicense 252
 lsmdisk 253
 lsmdiskcandidate 259
 lsmdiskdumps 257, 340
 lsmdiskextent 261
 lsmdiskgrp 262
 lsmdisklba 258
 lsmdiskmember 265
 lsmigrate 267
 lsnode 267
 lsnodecandidate 270
 lsnodedependentvdisks 271
 lsnodevpd 273, 341
 lspportip 276
 lsquorum 279
 lsrconsistgrp 280
 lsrcrelationship 283
 lsrcrelationshipcandidate 286
 lsrcrelationshipprogress 287
 lsrepairsevdiskcopyprogress 288
 lsrepairvdiskcopyprogress 289
 lsrmvdiskdependentmaps 290
 lsroute 291
 lssevdiskcopy 292
 lssnmpserver 294
 lssoftwaredumps 295, 344
 lssoftwareupgradestatus 296
 lssshkeys 296
 lssyslogserver 297
 lstimezones 298
 lsuser 299
 lsusergp 300
 lsvdisk 301
 lsvdiskcopy 306
 lsvdiskdependentmaps 308
 lsvdiskextent 309
 lsvdiskfmapcopies 310
 lsvdiskfcmappings 311

commands (*continued*)

lsvdiskhostmap 312
 lsvdisklba 313
 lsvdiskmember 314
 lsvdiskprogress 316
 lsvdisksyncprogress 317
 migrateexts 175
 migratetoimage 176
 migratevdisk 178
 mkauth 27
 mkemailserver 36
 mkemailuser 37
 mkfconsistgrp 143
 mkfcmmap 144
 mkhost 90
 mkmdiskgrp 129
 mkpartnership 158
 mkrconsistgrp 159
 mkrcrelationship 160
 mksnmpserver 38
 mksyslogserver 39
 mkuser 22
 mkusergrp 24
 mkvdisk 106
 mkvdiskhostmap 112
 ping 64
 prestartfconsistgrp 147, 150
 prestartfcmmap 148
 recovervdisk 114
 recovervdiskbycluster 115
 recovervdiskbyiogrp 115
 repairsevdiskcopy 116
 repairvdiskcopy 117
 restore 77
 rmallsshkeys 327
 rmauth 27
 rmemailserver 41
 rmemailuser 41
 rmfconsistgrp 149
 rmfcmmap 149
 rmhost 92
 rmhostiogrp 92
 rmhostport 93
 rmmdisk 130
 rmmdiskgrp 132
 rmnode 65
 rmpartnership 163
 rmpportip 67
 rmrcconsistgrp 163
 rmrcrelationship 164
 rmsnmpserver 42
 rmsshkey 329
 rmsyslogserver 42
 rmuser 25
 rmusergrp 25
 rmvdisk 118
 rmvdiskcopy 120
 rmvdiskhostmap 121
 sendinventoryemail 43
 setclustertime 68
 setdisktrace 179
 setemail 43
 setevent 83, 323
 setlocale 83
 setpwdreset 69
 setquorum 137
 settimezone 70

commands (*continued*)

settrace 180
 showtimezone 318
 shrinkvdisksize 121
 splitvdiskcopy 124
 startemail 43
 startfconsistgrp 150
 startfcmmap 151
 startfconsistgrp 165
 startfcrelationship 167
 startstats 70
 starttrace 182
 stopcluster 72
 stopemail 44
 stopfconsistgrp 152
 stopfcmmap 154
 stoprconsistgrp 169
 stoprcrelationship 171
 stopstats 74
 stoptrace 182
 svqueryclock 84
 switchrconsistgrp 172
 switchrcrelationship 173
 testemail 44
 triggermdiskdump 138
 writeserenum 84
 configuring
 PuTTY 5
 controller commands
 chcontroller 347
 overview 347
 controllers
 changing 347
 command 207, 347
 copyoperator role 27
 cpdumps command 62
 creating users 6
 current time zone 318
D
 data migration progress
 viewing 267
 dependent maps
 viewing 236
 detectmdisk command 63
 diagnostic and service-aid commands
 clearerrlog
 cluster 81
 cluster 79
 applysoftware 79
 cherrstate 80
 setevent 83
 setlocale 83
 svqueryclock 84
 writeserenum 84
 dumperrlog
 cluster 82
 finderr
 cluster 82
 overview 79
 dump files
 listing 198, 333
 lsfeaturedumps 237
 dumpauditlog command 13
 dumperrlog command 82, 331
 dumpinternallog command 325

E

- e-mail and event notification commands
 - chemailserver 31
 - chsnmpserver 34
 - chsyslogserver 35
 - mksnmpserver 38
 - mksyslogserver 39
 - rmemailserver 41
 - rmsnmpserver 42
 - rmsyslogserver 42
- e-mail commands
 - chemail 30
 - chemailuser 32
 - lsemailuser 213
 - mkemailuser 37
 - overview 29
 - rmemailuser 41
 - sendinventoryemail 43
 - setemail 43
 - startemail 43
 - stopemail 44
 - testemail 44
- email and event notification commands
 - mkemailserver 36
- error log commands
 - overview 321
 - setevent 323
- error log dump files
 - viewing 336
- event notification commands
 - overview 29
- exit command 332
- expandvdisksize command 104
- extent allocation
 - viewing 261

F

- featurization settings 252
- filtering
 - FlashCopy
 - consistency groups 230
 - mappings 232, 290, 308, 310
- finderr command 82
- FlashCopy commands
 - chfccoconsistgrp 141
 - chfcmap 141
 - mkfccoconsistgrp 143
 - mkfcmap 144
 - overview 141
 - prestartfccoconsistgrp 147, 150
 - prestartfcmap 148
 - rmfccoconsistgrp 149
 - rmfcmap 149
 - startfccoconsistgrp 150
 - startfcmap 151
 - stopfccoconsistgrp 152
 - stopfcmap 154
- FlashCopy progress 235
- free extents 238

G

- Global Mirror commands
 - chpartnership 155
 - chrconsistgrp 156

Global Mirror commands (continued)

- chrrelationship 156
- mkpartnership 158
- mkrccoconsistgrp 159
- mkrccrelationship 160
- overview 155
- rmpartnership 163
- rmrccoconsistgrp 163
- rmrccrelationship 164
- starttrccoconsistgrp 165
- starttrccrelationship 167
- stoptrccoconsistgrp 169
- stoptrccrelationship 171
- switchchrccoconsistgrp 172
- switchchrccrelationship 173

H

- help command 77
- host bus adaptor, unconfigured
 - viewing 239
- host commands
 - addhostiogrp 87
 - addhostport 88
 - chhost 89
 - mkhost 90
 - overview 87
 - rmhost 92
 - rmhostiogrp 92
 - rmhostport 93
- host I/O group 242
- hosts
 - commands 87
 - viewing 240

I

- includemdisk command 136
- information
 - center xxxii
- information commands 290, 310
 - caterrlog 195
 - caterrlogbyseqnum 197
 - ls2145dumps 198
 - lscimomdumps 199
 - lscluster 201
 - lsclustercandidate 204
 - lsclusterip 205
 - lscontroller 207
 - lscopystatus 200
 - lscurrentuser 210
 - lsdiscoverystatus 211
 - lsemailserver 212
 - lserrlogbyfccoconsistgrp 214
 - lserrlogbyfcmap 216
 - lserrlogbyhost 217
 - lserrlogbyiogrp 218
 - lserrlogbymdisk 219
 - lserrlogbymdiskgrp 220
 - lserrlogbynode 222
 - lserrlogbyrccoconsistgrp 223
 - lserrlogbyrccrelationship 224
 - lserrlogbyvdisk 226
 - lserrlogdumps 227
 - lsfabric 228
 - lsfccoconsistgrp 230

information commands (continued)

- lsfcmap 232
 - lsfcmapcandidate 234
 - lsfcmapdependentmaps 236
 - lsfcmapprogress 235
 - lsfeaturedumps 237
 - lsfreeextents 238
 - lshbaportcandidate 239
 - lshost 240
 - lshostiogrp 242
 - lshostvdiskmap 243
 - lsiogrp 245
 - lsiogrpcandidate 248
 - lsiogrpghost 247
 - lsiostatdumps 248
 - lsiotracedumps 249
 - lsiscsiauth 250
 - lslicense 252
 - lsmdisk 253
 - lsmdiskcandidate 259
 - lsmdiskdumps 257
 - lsmdiskextent 261
 - lsmdiskgrp 262
 - lsmdisklba 258
 - lsmdiskmember 265
 - lsmigrate 267
 - lsnode 267
 - lsnodecandidate 270
 - lsnodependentvdisks 271
 - lsnodevpd 273
 - lsportip 276
 - lsquorum 279
 - lsrccoconsistgrp 280
 - lsrccrelationship 283
 - lsrccrelationshipcandidate 286
 - lsrccrelationshipprogress 287
 - lsroute 291
 - lssnmpserver 294
 - lssoftwaredumps 295
 - lssshkeys 296
 - lssyslogserver 297
 - lstimezones 298
 - lsuser 299
 - lsusergrp 300
 - lsvdisk 301
 - lsvdiskdependentmaps 308
 - lsvdiskextent 309
 - lsvdiskfmappings 311
 - lsvdiskhostmap 312
 - lsvdisklba 313
 - lsvdiskmember 314
 - lsvdiskprogress 316
 - overview 195
 - showtimezone 318
- inventory commands
 - chcluster 51
 - chemail 30
 - mkemailuser 37
 - rmemailuser 41
 - sendinventoryemail 43
 - startemail 43
 - stopemail 44
 - testemail 44

K

- keyboard 545

keys
 adding
 for other hosts 7

L

legal notices 547
license
 changing settings 323
 viewing 252
licensing commands
 chlicense 323
 dumpinternallog 325
list dumps commands 193
ls2145dumps command 198, 333
lsauiditlogdumps command 14
lsauth command 29
lscimomdumps command 199, 334
lscluster command 201
lsclustercandidate command 204
lsclusterip command 205
lsclustervpd command 335
lscontroller command 207
lscontrollerdependentvdisks
 command 210
lscopystatus command 200
lscurrentuser command 210
lsdiscoverystatus command 211
lsemailer command 212
lsemaileruser command 213
lserrlogbyfconsistentgrp command 214
lserrlogbyfcmapp command 216
lserrlogbyhost command 217
lserrlogbyiogrp command 218
lserrlogbymdisk command 219
lserrlogbymdiskgp command 220
lserrlogbynode command 222
lserrlogbyrconsistentgrp command 223
lserrlogbyrrelationship command 224
lserrlogbyvdisk command 226
lserrlogdumps command 227, 336
lsfabric command 228
lsfcconsistentgrp command 230
lsfcmap command 232
lsfcmapcandidate command 234
lsfcmapdependentmaps command 236
lsfcmapprogress command 235
lsfeaturedumps command 237
lsfeaturedumps commands 337
lsfreeextents command 238
lshbaportcandidate command 239
lshost command 240
lshostiogrp command 242
lshostvdiskmap command 243
lsiogrp command 245
lsiogrpcandidate command 248
lsiogrpghost command 247
lsiostatsdumps command 248, 338
lsiotracedumps command 249
lsiotracedumps commands 339
lsiscsiauth command 250
lslicense command 252
lsmdisk command 253
lsmdiskcandidate command 259
lsmdiskdumps command 257, 340
lsmdiskextent command 261
lsmdiskgrp command 262

lsmdisklba command 258
lsmdiskmember command 265
lsmigrate command 267
lsnode command 267
lsnodecandidate command 270
lsnodecandidatevdisks command 271
lsnodevpd command 273, 341
lsportip command 276
lsquorum command 279
lsrconsistentgrp command 280
lsrrelationship command 283
lsrrelationshipcandidate command 286
lsrrelationshipprogress command 287
lsrepairvdiskcopyprogress
 command 288
lsrepairvdiskcopyprogress
 command 289
lsrmvdependentmaps command 290
lsroute command 291
lssevdiskcopy command 292
lssnmpserver command 294
lssoftwareumps command 295, 344
lssoftwareupgradestatus command 296
lsshkeys command 296
lssyslogserver command 297
lstimezones command 298
lsuser command 299
lsusergrp command 300
lsvdisk command 301
lsvdiskcopy command 306
lsvdiskdependentmaps command 308
lsvdiskextent command 309
lsvdiskfcmappcopies command 310
lsvdiskfcmappings command 311
lsvdiskhostmap command 312
lsvdisklba command 313
lsvdiskmember command 314
lsvdiskprogress command 316
lsvdisksyncprogress command 317

M

managed disk commands
 applydisksoftware 135
 chmdisk 136
 includemdisk 136
 lsquorum 279
 overview 135
 setquorum 137
 triggermdiskdump 138
managed disk group commands
 addmdisk 127
 chmdiskgrp 128
 mkmdiskgrp 129
 overview 127
 rmmdisk 130
 rmmdiskgrp 132
managed disks
 viewing disks 253, 258
 viewing groups 262
master console
 configuration 4
MDisks *See* managed disks 127, 135
Metro Mirror commands
 chpartnership 155
 chrconsistentgrp 156
 chrrelationship 156

Metro Mirror commands (*continued*)

mkpartnership 158
mkrconsistentgrp 159
mkrrelationship 160
overview 155
rmpartnership 163
rmrconsistentgrp 163
rmrrelationship 164
startcrconsistentgrp 165
startcrrelationship 167
stopcrconsistentgrp 169
stopcrrelationship 171
switchcrconsistentgrp 172
switchcrrelationship 173
migrateexts command 175
migratetoimage command 176
migratevdisk command 178
migration 175
migration commands
 migrateexts 175
 migratetoimage 176
 migratevdisk 178
 overview 175
mkauth command 27
mkemailserver command 36
mkemailuser command 37
mkfconsistentgrp command 143
mkfcmapp command 144
mkhost commands 90
mkmdiskgrp command 129
mkpartnership command 158
mkrconsistentgrp command 159
mkrrelationship command 160
mksnmpserver command 38
mksyslogserver command 39
mkuser command 22
mkusergrp command 24
mkvdisk commands 106
mkvdiskhostmap command 112

N

nodes
 adding 48
 addnode command 48
 changing 59
 chnode command 59
 deleting 65
 rmnode command 65
 viewing 267

O

overview
 audit log commands 11
 backup and restore commands 75
 cluster commands 47
 cluster diagnostic and service-aid
 commands 79
 controller commands 347
 e-mail commands 29
 error log commands 321
 event notification commands 29
 FlashCopy commands 141
 host commands 87
 information commands 195

- overview (*continued*)
 - licensing commands 323
 - list dumps commands 193
 - managed disk commands 135
 - managed disk group commands 127
 - migration commands 175
 - role-based security commands 27
 - secure shell 3
 - secure shell key commands 327
 - service mode commands 329
 - service mode information commands 333
 - tracing commands 179
 - user management commands 17

P

- ping command 64
- prestartfcconsistgrp command 147
- prestartfcmap command 148
- PuTTY
 - configuring 5
 - generating an SSH key pair 4
 - installing 1
 - scp (pscp) 9

R

- recovervdisk command 114
- recovervdiskbycluster command 115
- recovervdiskbyiogrp command 115
- related information xxxii
- repairsevdiskcopy command 116
- repairdiskcopy command 117
- restore command 77
- restore commands
 - clear 76
 - help 77
 - restore 77
- rmallsshkeys command 327
- rmauth command 27
- rmemailserver command 41
- rmemailuser command 41
- rmfcconsistgrp command 149
- rmfcmap command 149
- rmhost command 92
- rmhostiogrp command 92
- rmhostport command 93
- rmmdisk command 130
- rmmdiskgrp command 132
- rmnode command 65
- rmpartnership command 163
- rmportip command 67
- rmrconsistgrp command 163
- rmrrelationship command 164
- rmsnmpserver command 42
- rmsshkey command 329
- rmsyslogserver command 42
- rmuser command 25
- rmusergrp command 25
- rmvdisk command 118
- rmvdiskcopy command 120
- rmvdiskhostmap command 121
- Role based security commands
 - lsauth 29
 - mkauth 27

- Role based security commands (*continued*)
 - rmauth 27
- role-based security commands
 - overview 27

S

- SAN Volume Controller library
 - related publications xxxii
- scp
 - PuTTY application 9
- secure shell
 - client system
 - preparing for CLI 1
 - PuTTY 5
 - viewing keys 296
- secure shell (SSH)
 - adding keys 7
 - authenticating logins 3
 - client system 4
 - configuring (not PuTTY) 4
 - creating keys 4
 - overview 3
- secure shell client
 - preparing for CLI on AIX 1
 - preparing for CLI on Linux 1
 - preparing for CLI on Windows 1
- secure shell key commands
 - addsshkey 327
 - overview 327
 - rmallsshkeys 327
 - rmsshkey 329
- security 3
- sendinventoryemail command 43
- service mode
 - commands 329
 - information commands 333
- service mode commands
 - applysoftware 329
 - cleardumps 330
 - dumperrlog 331
 - exit 332
 - overview 329
- service mode information commands
 - ls2145dumps 333
 - lscimomdumps 334
 - lsclustervpd 335
 - lserrlogdumps 336
 - lsfeaturedumps 337
 - lsiosatsdumps 338
 - lsiotracedumps 339
 - lsmdiskdumps 340
 - lsnodevpd 341
 - lssoftwareumps 344
 - overview 333
- setclustertime command 68
- setdisktrace command 179
- setemail command 43
- setevent command 83, 323
- setlocale command 83
- setpwdreset command 69
- setquorum command 137
- settimezone command 70
- settrace command 180
- shortcut keys 545
- showtimezone command 318

- shrinkvdisksize command 121
- software packages
 - listing 344
 - viewing 295
- splitvdiskcopy command 124
- SSH *See* secure shell 1, 3
- SSH *See* SSH client 2
- SSH keys
 - creating 4
- startemail command 43
- startfcconsistgrp command 150
- startfcmap command 151
- starttrconsistgrp command 165
- starttrrelationship command 167
- startstats command 70
- starttrace command 182
- stopcluster command 72
- stopemail command 44
- stopfcconsistgrp command 152
- stopfcmap command 154
- stoprconsistgrp command 169
- stoprrelationship command 171
- stopstats command 74
- stoptrace command 182
- svqueryclock command 84
- switchrconsistgrp command 172
- switchrrelationship command 173

T

- testemail command 44
- time zones 298
- tracing commands
 - overview 179
 - setdisktrace 179
 - settrace 180
 - starttrace 182
 - stoptrace 182
- trademarks 549
- triggermdiskdump command 138

U

- user management commands
 - chauthservice 17
 - chcurrentuser 20
 - chuser 20
 - chusergrp 22
 - mkuser 22
 - mkusergrp 24
 - overview 17
 - rmuser 25
 - rmusergrp 25
- users
 - creating 6

V

- VDisks (virtual disks)
 - viewing 301
- VDisks *See* virtual disks 97
- viewing
 - clusters 201
 - Global Mirror
 - consistency groups 280
 - relationships 283

- viewing (*continued*)
 - I/O groups 245
 - Metro Mirror
 - consistency groups 280
 - relationships 283
- virtual disk
 - viewing FlashCopy mappings 311
- virtual disk (VDisks) extent
 - viewing 309
- virtual disk commands
 - addvdiskcopy 97
 - chvdisk 100
 - expandvdisksize 104
 - lscontrollerdependentvdisks 210
 - lsrepairsevdiskcopyprogress 288
 - lsrepairvdiskcopyprogress 289
 - lssevdiskcopy 292
 - lsvdiskcopy 306
 - lsdisksyncprogress 317
 - mkvdisk 106
 - mkvdiskhostmap 112
 - recovervdisk 114
 - recovervdiskbycluster 115
 - recovervdiskbyiogrp 115
 - repairsevdiskcopy 116
 - repairvdiskcopy 117
 - rmvdisk 118
 - rmvdiskcopy 120
 - rmvdiskhostmap 121
 - shrinkvdisksize 121
 - splitvdiskcopy 124
- virtual disks
 - copying 97
 - creating 106
 - removing 120
 - viewing disks 313
- virtual disks commands
 - overview 97
- vital product data (VPD)
 - listing 335
 - viewing 273

W

- writesernum command 84

Readers' Comments — We'd Like to Hear from You

IBM System Storage SAN Volume Controller
Command-Line Interface User's Guide
Version 5.1.0

Publication No. SC26-7903-06

We appreciate your comments about this publication. Please comment on specific errors or omissions, accuracy, organization, subject matter, or completeness of this book. The comments you send should pertain to only the information in this manual or product and the way in which the information is presented.

For technical questions and information about products and prices, please contact your IBM branch office, your IBM business partner, or your authorized remarketer.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you. IBM or any other organizations will only use the personal information that you supply to contact you about the issues that you state on this form.

Comments:

Thank you for your support.

Send your comments to the address on the reverse side of this form.

If you would like a response from IBM, please fill in the following information:

Name

Address

Company or Organization

Phone No.

E-mail address



Fold and Tape

Please do not staple

Fold and Tape



NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Information Development
Department 61C
9032 South Rita Road
Tucson, Arizona 85755-4401



Fold and Tape

Please do not staple

Fold and Tape



Printed in USA

SC26-7903-06



Spine information:



IBM System Storage SAN Volume
Controller

SAN Volume Controller Command-Line Interface
User's Guide

Version 5.1.0