IBM Storage Virtualize for SAN Volume Controller and FlashSystem Family

Command-Line Interface User's Guide



Note

Before using this information and the product it supports, read the information in “Notices ” on page 1149.

This edition applies to version 8, release 7, modification x, and to all subsequent modifications until otherwise indicated in new editions.

© Copyright International Business Machines Corporation 2003, 2025.
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Tables ..... **xvii**

About this guide ..... **xxv**
Who should use this guide xxv
Accessibility xxv
Emphasis xxv

Chapter 1. Setting up an SSH client **1**
Setting up an SSH client on a Windows host 1
Generating an SSH key pair using PuTTY 2
Configuring a PuTTY session for the CLI 3
Connecting to the CLI using PuTTY 3
Starting a PuTTY session for the CLI 5
Setting up an SSH client on an AIX or Linux host 6
Generating an SSH key pair using OpenSSH 6
Connecting to the CLI with OpenSSH 7
Managing SSH keys by using the CLI 7

Chapter 2. Copying the software update files by using PuTTY pscp or openssh scp **9**

Chapter 3. Using the dump commands to work with directories **11**

Chapter 4. CLI reference **13**
Syntax diagrams 13
Using wildcards in the CLI 14
Data types and value ranges 15
CLI flags 20
CLI messages 20
CLI commands and parameters 21
CLI special characters 22
CLI deprecated and discontinued commands 22
Understanding capacity indicators 24
Attributes of the -filtervalue parameters 24
UNIX commands available on the CLI 25

Chapter 5. Array commands **27**
charray 27
charraymember 28
expandarray 30
lsarray 31
lsarrayexpansionprogress 34
lsarrayinitprogress 36
lsarraylba 36
lsarraymember 38
lsarraymembergoals 39
lsarraymemberprogress 40
lsarrayrecommendation 41
lsarraysyncprogress 43

lscompatibledriveclasses....43
lspotentialarraysize....44
mkarray (Deprecated)....45
mkdistributedarray....46
recoverarray....49
recoverarraybycluster (Discontinued)....49
recoverarraybysystem....49
rmarray....50

Chapter 6. Audit log commands....51
cauditlog....51
dumpauditlog....53
lsauditlogdumps (Deprecated)....54

Chapter 7. Backup and restore commands....55
svcconfig....55
backup....56
clear....57
cron....58
recover....58
restore....59

Chapter 8. Cloud commands....61
cfgcloudstorage....61
cfgcloudstorage....61
querycloudstoragecandidate....62
querycloudstoragecandidate....63
chcloudaccountawss3....64
chcloudaccountswift....66
chcloudaccountazure....69
lscloudaccount....71
lscloudaccountusage....73
lscloudaccountimportcandidate....74
mkcloudaccountawss3....75
mkcloudaccountswift....77
mkcloudaccountazure....79
rmcloudaccount....80
testcloudaccount....81

Chapter 9. Clustered system commands....83
addnode....83
addiscsistorageport....86
cfgportip (Deprecated)....87
chbanner....87
chcluster (Discontinued)....89
chiogrp....89
chiscsistorageport....91
chiscsiportauth....92
chiscsiportauth....93
chnode (SVC) / chnodecanister (Storwize family products)....94
chnodebattery....96
chnodebootdrive....97
chnodecanister....98
chnodehw....100
chnodecanisterhw....101
chquorum....102
chsecurity....104

chsite ... 111
chsra ... 112
chsystem ... 115
chsystemlimits ... 121
chsystemcert ... 122
chsystemip (Deprecated) ... 125
chsystemethernet ... 125
chsystempatch ... 126
chthrottle ... 127
chcompatibilitymode ... 127
cpdumps ... 128
detectiscsistorageportcandidate ... 129
dumpconfig (Discontinued) ... 131
help ... 131
lsavailablepatch ... 131
lsclustercandidate (Discontinued) ... 132
lscluster (Discontinued) ... 132
lsclusterip (Discontinued) ... 132
lsclusterstats (Discontinued) ... 132
lsdiscoverystatus ... 132
lsfabric ... 133
lsfabricport ... 135
lsnvmefabric ... 137
lsfcportcandidate ... 138
lsiscsiportauth ... 139
lsiscsiportauth ... 139
lsiscsistorageport ... 140
lsiscsistorageportcandidate ... 143
lsiogrp ... 144
lshbaportcandidate (Deprecated) ... 147
lsiogrphost ... 147
lsiogrpcandidate ... 147
lsiostatsdumps (Deprecated) ... 148
lsiotracedumps (Deprecated) ... 148
lsnode ... 148
lsnodebattery ... 149
lsnodecandidate (SAN Volume Controller) ... 151
lsnodecanister ... 151
lsnodecanisterhw ... 154
lsnodecanisterstats ... 155
lsnodecanisterstats ... 159
lsnodecanistervpd ... 162
lsnodedependentvdisks (Deprecated) ... 163
lsnodehw ... 163
lsnodepsu ... 165
lsnodestats ... 166
lsnodevpd ... 170
lsportip (Deprecated) ... 171
lsportfc ... 172
lsportsas ... 174
lsportstats ... 176
lsportusb ... 178
lsquorum ... 180
lsroute ... 182
lstimezones ... 182
lssasportcandidate ... 182
lssecurity ... 183
lssite ... 187

lssra....188
lsthrottle....190
lssystem....191
lssystemlimits....198
lssystemcert....199
lssystemethernet....200
lssystemip....200
lssystemstats....201
lstargetportfc....205
lscompatibilitymode....207
mkquorumapp....207
mkthrottle....208
ping....209
rmiscsistorageport....210
rmnode....211
rmnodecanister....213
rmportip (Deprecated)....215
rmthrottle....215
setclustertime (Discontinued)....215
setsystemtime....215
setpwdreset....216
settimezone....217
showtimezone....217
startstats....218
stopstats (Deprecated)....219
stopcluster (Discontinued)....219
stopsystem....219
swapnode....221
chnodeserviceip....223
svcinfo lsnodestatus....224
svctask restartservice....224
svctask startsystem....225
svcinfo traceroute....225

Chapter 10. Clustered system diagnostic and service-aid commands....227
applysoftware....227
caterrlog (Deprecated)....229
caterrlogbyseqnum (Deprecated)....229
cherrstate (Deprecated)....229
chdnsserver....229
cheventlog....230
chsyslogserver....231
cleardumps....232
clearerrlog....234
cpfabricdumps (Discontinued)....234
dumperrlog....234
finderr....235
setevent (Discontinued)....236
lscimomdumps (Deprecated)....236
lscopystatus....236
lsdumps....236
lsdnsserver....237
lserrlogbyfcconsistgrp (Deprecated)....239
lserrlogbyfcmap (Deprecated)....239
lserrlogbyhost (Deprecated)....239
lserrlogbyiogrp (Deprecated)....239
lserrlogbymdisk (Deprecated)....239

lserrlogbymdiskgrp (Deprecated)..... 239
lserrlogbynode (Deprecated)..... 239
lserrlogbyrcconsistgrp (Deprecated)..... 239
lserrlogbyrcrelationship (Deprecated)..... 239
lserrlogbyvdisk (Deprecated)..... 239
lserrlogdumps (Deprecated)..... 239
lsfeaturedumps (Deprecated)..... 239
lseventlog..... 240
lssyslogserver..... 244
lssoftwaredumps (Deprecated)..... 246
lssoftwareupgradestatus (Deprecated)..... 246
lssystemsupportcenter..... 246
lsupdate..... 247
mkdnsserver..... 249
mksyslogserver..... 250
mksystemsupportcenter..... 251
rmdnsserver..... 253
rmsyslogserver..... 253
rmsystemsupportcenter..... 254
setlocale..... 254
svqueryclock..... 255
writesernum..... 256
lspatch..... 256
lssystempatches..... 257
svctask removepatch..... 258
svctask purgepatches..... 258
svc_snap..... 259

Chapter 11. Controller commands..... 261
chcontroller..... 261
lscontroller..... 262
lscontrollerports..... 265

Chapter 12. Drive commands..... 267
applydrivesoftware..... 267
chdrive..... 269
lsdrive..... 270
lsdriveclass..... 274
lsdrivelba..... 275
lsdriveprogress..... 276
lsdriveupgradeprogress..... 277
triggerdrivedump..... 278

Chapter 13. Call home and event notification commands..... 281
cfgcloudcallhome..... 281
cfgcloudcallhome..... 281
chcloudcallhome..... 282
chemail 283
chemailserver..... 286
chemailuser..... 287
chproxy..... 289
chsnmpserver..... 290
lscloudcallhome..... 292
lsemailserver..... 293
lsemailuser..... 294
lsproxy..... 295
lssnmpserver..... 295

mkemailserver ... 297
mkemailuser ... 298
mkproxy ... 300
mksnmpserver ... 301
rmemailserver ... 303
rmemailuser ... 304
rmproxy ... 304
rmsnmpserver ... 304
sendcloudcallhome ... 305
sendinventoryemail ... 306
setemail (Discontinued) ... 306
startemail ... 306
stopemail ... 307
testemail ... 308
testsnmpserver ... 308

Chapter 14. Enclosure commands ... 311
addcontrolenclosure ... 311
addcontrolenclosure ... 311
chenclosure ... 312
chenclosurecanister ... 314
chenclosuredisplaypanel ... 315
chenclosurefanmodule ... 316
chenclosuresem ... 316
chenclosureslot ... 317
sainfo chenclosurevpd (Deprecated) ... 319
lsenclosure ... 319
lsenclosurebattery ... 321
lsenclosurebattery ... 323
lscontrolenclosurecandidate ... 325
lsenclosurecanister ... 325
lsenclosuredisplaypanel ... 329
lsenclosurefanmodule ... 330
lsenclosurepsu ... 331
lsenclosuresem ... 332
lsenclosureslot ... 334
lsenclosurestats ... 336
lssasfabric ... 337
resetleds ... 338
triggerenclosuredump ... 339

Chapter 15. Encryption commands ... 341
chencryption ... 341
chkeyserver ... 343
chkeyserverisklm ... 344
chkeyserverciphertrustmanager ... 345
chkeyserverkeysecure (Deprecated) ... 346
lsencryption ... 346
lskeyserver ... 348
lskeyserverisklm ... 349
lskeyserverciphertrustmanager ... 350
lskeyserverkeysecure (Deprecated) ... 350
mkkeyserver ... 351
rmkeyserver ... 352
testkeyserver ... 352

Chapter 16. Policy-based replication commands ... 355

chvolumegroupreplication....355
lsreplicationpolicy....356
lsvolumegroupreplication....357
mkreplicationpolicy 364
rmreplicationpolicy....366

Chapter 17. Licensing and featurization commands....367
activatefeature....367
chlicense....368
deactivatefeature....369
lsfeature....369
lslicense....371

Chapter 18. Fibre Channel portset....373
lsfcportsetmember....373
addfcportsetmember....374
rmfcportsetmember....374

Chapter 19. FlashCopy commands....377
chfcconsistgrp....377
chfcmap....377
lsfcconsistgrp....379
lsfcmap....381
lsfcmapcandidate....384
lsfcmapprogress....384
lsfcmapdependentmaps....385
lsrmvdiskdependentmaps....385
mkfcconsistgrp....386
mkfcmap....387
prestartfcconsistgrp....389
prestartfcmap....390
rmfcconsistgrp....391
rmfcmap....392
startfcconsistgrp....393
startfcmap....394
stopfcconsistgrp....395
stopfcmap....396

Chapter 20. FlashSystem grid commands....399
manageflashgrid....399

Chapter 21. Host commands....401
addhostclustermember....401
addhostiogrp....402
addhostport....402
chhost....405
chhostcluster....408
lshost....409
lshostports....413
lshostcluster....414
lshostclustermember....415
lshostclustervolumemap....416
lshostiogrp....417
lshostiplogin....417
lsiscsiauth....419
mkhost....421
mkhostcluster....425

mkvolumehostclustermap.... 426
rmhost.... 427
rmhostcluster.... 428
rmhostclustermember.... 429
rmvolumehostclustermap.... 430
rmhostiogrp.... 430
rmhostport.... 431

Chapter 22. Livedump commands.... 435
cancellivedump.... 435
lslivedump.... 435
preplivedump.... 436
triggerlivedump.... 436

Chapter 23. Managed disk commands.... 439
addmdisk.... 439
applymdisksoftware (Discontinued).... 441
chmdisk.... 441
detectmdisk.... 442
dumpallmdiskbadblocks.... 443
dumpmdiskbadblocks.... 444
includemdisk.... 445
lslocaldisk.... 446
lslocaldisk.... 448
lsmdisk.... 450
lsmdiskdumps (Deprecated).... 455
lsmdisklba.... 455
lsmdiskcandidate.... 456
lsmdiskextent.... 457
lsmdiskmember.... 458
setquorum (Deprecated).... 459
triggermdiskdump (Discontinued).... 459

Chapter 24. Copy Service commands.... 461
chpartnership.... 461
lspartnership.... 463
lspartnershipcandidate.... 467
mkfcpartnership.... 467
mkippartnership.... 468
mkpartnership (Discontinued).... 470
rmpartnership.... 470

Chapter 25. Migration commands.... 471
lsmigrate.... 471
migrateexts.... 471
migratetoimage.... 473
migratevdisk.... 475

Chapter 26. Multifactor authentication commands.... 477
lsauthmultifactorverify.... 477
chauthmultifactorverify.... 478
lsauthmultifactorduo.... 479
chauthmultifactorduo.... 480

Chapter 27. IP addresses commands.... 483
chip.... 483

chportethernet....484
chportset....485
lsip....486
lsportethernet....487
lsportset....490
mkip....491
mkportset....492
rmip....493
rmportset....494

Chapter 28. Partition commands....495
chpartition....495
lspartition....498
mkpartition....502
rmpartition....503
mergepartition....505

Chapter 29. Provisioning policies commands....507
chprovisioningpolicy....507
lsprovisioningpolicy....507
mkprovisioningpolicy....508
rmprovisioningpolicy....509

Chapter 30. Recovery access key command....511
satask unlockencryption....511

Chapter 31. Legacy Safeguarded policy commands....513
chsafeguardedpolicy....513
lssafeguardedpolicy....513
lssafeguardedschedule....514
mksafeguardedpolicy....516
rmsafeguardedpolicy....517

Chapter 32. Service information commands....519
host....519
lsbootdrive....519
lscmdstatus....521
sainfo lsfcportfeature....522
lsfiles....523
lshardware....524
lsnodediscoverysubnet....525
lsnodeip (Deprecated)....525
lsnodeipconnectivity....526
lsservicenodes....527
lsservicerecommendation....528
lsservicestatus....529
traceroute....536

Chapter 33. Service mode commands (Discontinued)....537
applysoftware (Discontinued)....537
cleardumps (Discontinued)....537
dumperrlog (Discontinued)....537
exit (Discontinued)....537

Chapter 34. Service mode information commands (Discontinued)....539
ls2145dumps (Discontinued)....539

lscimomdumps (Discontinued)..... 539
lsclustervpd (Discontinued)..... 539
lserrlogdumps (Discontinued)..... 539
lsfeaturedumps (Discontinued)..... 539
lsiostatsdumps (Discontinued)..... 539
lsiotracedumps (Discontinued)..... 539
lsmdiskdumps (Discontinued)..... 539
lssoftwaredumps (Discontinued)..... 539

Chapter 35. Service task commands..... 541
addnodediscoverysubnet 541
chbootdrive..... 541
satask chfcportfeature 542
chnodeled..... 543
chnodeip (Discontinued)..... 543
chnodeusb..... 543
chnodeusb..... 544
satask chserviceip..... 545
chwwnn..... 547
chvpd..... 547
cpfiles..... 549
satask downloadsoftware..... 550
dumpinternallog (Discontinued)..... 553
installsoftware..... 553
leavecluster..... 554
locksuperuser..... 555
metadata..... 556
mkcluster..... 557
mkcluster (Deprecated)..... 558
overridequorum..... 558
satask rescuenode..... 559
resetpassword..... 560
satask restartservice..... 560
rmnodediscoverysubnet 561
setlocale..... 562
setpacedccu..... 563
settempsshkey..... 563
snap..... 564
startservice..... 565
stopnode..... 566
stopservice..... 566
supportupload..... 567
t3recovery..... 569
unlocksuperuser..... 569
satask removepatch..... 570
satask purgepatches..... 570

Chapter 36. Service node task commands..... 573
sntask applysecuritypatch..... 573
sntask chnode..... 574
sntask cleansnap..... 575
sntask initnode..... 576
sntask rmnode..... 578
sntask snap..... 578
sntask startnode..... 579
sntask startservice..... 579
sntask stopnode..... 580

sntask stopservice ... 580

Chapter 37. Single sign-on commands ... 581
lsauthsinglesignon ... 581
chauthsinglesignon ... 582

Chapter 38. System certificates commands ... 585
sainfo lsrootcertificate ... 585
satask exportrootcertificate ... 585

Chapter 39. Snapshot commands ... 587
addsnapshot ... 587
chsnapshot ... 590
lsvolumegrouppopulation ... 591
lsvolumepopulation ... 592
lsvolumegroupsnapshot ... 593
lsvolumesnapshot ... 595
refreshfromsnapshot ... 597
restorefromsnapshot ... 598
rmsnapshot ... 599

Chapter 40. Snapshot policy commands ... 601
addsnapshotpolicyschedule ... 601
chsnapshotpolicy ... 602
chvolumegroupsnapshotpolicy ... 602
lssnapshotpolicy ... 603
lssnapshotschedule ... 604
lsvolumegroupsnapshotpolicy ... 605
mksnapshotpolicy ... 606
rmsnapshotpolicy ... 607
lsvolumegroupsnapshotschedule ... 608

Chapter 41. Storage pool commands ... 611
chmdiskgrp ... 611
lsfreeextents ... 614
lsmdiskgrp ... 614
mkmdiskgrp ... 619
rmmdisk ... 625
rmmdiskgrp ... 626

Chapter 42. Truststore management commands ... 629
chtruststore ... 629
lstruststore ... 630
mktruststore ... 631
rmtruststore ... 632

Chapter 43. Two person integrity commands ... 633
chtwopersonintegrityrequest ... 633
lstwopersonintegrityrequest ... 634
mktwopersonintegrityrequest ... 635

Chapter 44. User management and access control commands ... 637
chauthservice ... 637
chcurrentuser ... 639
chldap ... 640
chldapserver ... 644

chownershipgroup ... 645
chuser ... 646
chusergrp ... 648
lscurrentuser ... 649
lsldap ... 650
lsldapserver ... 651
lsownershipgroup ... 652
lsuser ... 653
lsusergrp ... 654
mkldapserver ... 655
mkownershipgroup ... 657
mkuser ... 657
mkusergrp ... 659
rmldapserver ... 660
rmownershipgroup ... 661
rmuser ... 661
rmusergrp ... 662
testldapserver ... 663

Chapter 45. Volume commands ... 665
addvolumecopy ... 665
addvdiskcopy ... 667
addvdiskaccess ... 675
analyzevdisk ... 676
analyzevdiskbysystem ... 677
backupvolume ... 677
backupvolumegroup ... 678
chvdisk ... 680
chvolumegroup ... 686
converttoclone ... 689
expandvdisksize ... 690
expandvolume - Deprecated ... 691
lsdependentvdisks ... 692
lshostvdiskmap ... 693
lsmetadatavdisk ... 694
lsrepairsevdiskcopyprogress ... 695
lsrepairvdiskcopyprogress ... 696
lssevdiskcopy ... 697
lsvdisk ... 700
lsvdiskaccess ... 707
lsvdiskanalysis ... 707
lsvdiskanalysisprogress ... 709
lsvdiskcopy ... 710
lsvdiskdependentmaps ... 713
lsvdiskextent ... 714
lsvdiskfcmapcopies ... 715
lsvdiskfcmappings ... 715
lsvdiskhostmap ... 716
lsvdisklba ... 717
lsvdiskmember ... 718
lsvdiskprogress ... 719
lsvdisksyncprogress ... 720
lsvolumebackup ... 720
lsvolumebackupgeneration ... 722
lsvolumebackupprogress ... 723
lsvolumegroup ... 724
lsvolumerestoreprogress ... 726

mkmetadatavdisk....728
mkvdisk....728
mkvdiskhostmap....737
mkvolume....739
mkvolumegroup....749
mkimagevolume....755
movevdisk....757
recovervdisk....758
recovervdiskbycluster (Discontinued)....759
recovervdiskbyiogrp....759
recovervdiskbysystem....760
repairsevdiskcopy....761
repairvdiskcopy....761
restorevolume....763
rmmetadatavdisk....764
rmvdisk....764
rmvdiskcopy....766
rmvdiskaccess....768
rmvdiskhostmap....768
rmvolume....769
rmvolumecopy....770
rmvolumegroup....771
rmvolumebackupgeneration....772
shrinkvdisksize....773
splitvdiskcopy....775
lsvdisktiers....776

Chapter 46. Command-line interface messages....779

Appendix A. Accessibility features for the system....1147

Notices....1149
Trademarks....1150

Tables

1. Syntax diagrams 13

2. Data types 15

3. Capacity indicators 24

4. UNIX commands for interactive SSH sessions 25

5. charraymember combination options 29

6. Array output 32

7. Output 35

8. lsarrayinitprogress output 36

9. lsarraylba output 37

10. lsarraymember output 38

11. lsarraymembergoals output 39

12. lsarraymemberprogress output 41

13. lsarrayrecommendation output 42

14. lsarraysyncprogress output 43

15. Output 44

16. lspotentialarraysize output 45

17. catauditlog output 52

18. querycloudstoragecandidate output 62

19. querycloudstoragecandidate output 63

20. lscloudaccount output 72

21. lscloudaccountusage output 74

22. lscloudaccountimportcandidate output 75

23. Memory that is required for copy services and volume mirroring 90

24. Number of extents reserved by extent size..... 103

25. lsavailablepatch output..... 132

26. lsfabric output..... 134

27. lsfabricport output..... 136

28. lsnvmefabric output..... 137

29. lsfcportcandidate output..... 139

30. lsiscsiportauth output..... 139

31. lsiscsiportauth output..... 140

32. lsiscsistorageport output..... 141

33. lsiscsistorageportcandidate output..... 143

34. lsnodebattery attribute values..... 149

35. lsnodecandidate outputs..... 151

36. lsnodecanister attribute values..... 152

37. lsnodecanisterhw attribute values..... 154

38. Node canister attribute values..... 156

39. Stat_name field values..... 156

40. Node canister attribute values..... 160

41. Stat_name field values..... 160

42. Attribute values for lsnodehw and lsnodecanisterhw..... 164

43. Attribute values for lsnodepsu..... 165

44. Attribute values for lsnodestats or lsnodecanister 167

45. Stat_name field values..... 167

46. Attribute values for lsnodevpd and lsnodecanistervpd..... 171

47. lsportfc output..... 172

48. lsportsas output..... 174

49. lsportstats output ... 176

50. lsportusb output ... 179

51. lsquorum output ... 181

52. lssasportcandidate output ... 183

53. lssecurity attribute values ... 184

54. lssite attribute values ... 188

55. lssra output ... 188

56. lssra output ... 190

57. lsthrottle output ... 191

58. lssystem output ... 192

59. lssystemlimits output ... 198

60. lssystemcert output ... 199

61. lssystemethernet output ... 200

62. lssystemip output ... 201

63. lssystemstats attribute values ... 202

64. Stat_name field values ... 203

65. lstargetportfc output ... 206

66. lscompatibilitymode output ... 207

67. lsdnsserver output ... 238

68. lseventlog output ... 241

69. lssyslogserver output ... 245

70. lssystemsupportcenter output ... 246

71. lspatch output ... 257

72. lssystempatches output ... 258

73. lscontroller output ... 263

74. lscontrollerports attribute values....265
75. lsdrive output....271
76. lsdriveclass output....274
77. lsdrivelba output....276
78. lscloudcallhome output....292
79. lsemailserver output....293
80. lsproxy command output values....295
81. lssnmpserver output....296
82. lsenclosure output....319
83. lsenclosurebattery outputs....321
84. lsenclosurebattery outputs....323
85. lscontrolenclosurecandidate attribute values....325
86. lsenclosurecanister output....326
87. lsenclosuredisplaypanel output....329
88. lsenclosurefanmodule attribute values....330
89. lsenclosurepsu output....331
90. lsenclosuresem output....333
91. lsenclosureslot output....335
92. lsenclosurestats outputs....337
93. Stat_name field values....337
94. lssasfabric output....338
95. lsencryption output....347
96. lskeyserver output....348
97. lskeyserverisklm output....349
98. lskeyserverciphertrustmanager Output....350

99. lsreplicationpolicy output 357
100. lsvolumegroupreplication output 358
101. lsfeature outputs 370
102. lslicense output 371
103. lsfcportsetmember output 373
104. Relationship between the rate, data rate, and grains per second values 379
105. Relationship between the rate, data rate, and grains per second values 389
106. lshost output 412
107. lshostports attribute values 413
108. lshostcluster output 414
109. lshostclustermember output 415
110. lshostclustervolumemap output 416
111. svcinfo lshostiplogin output 418
112. lsiscsiauth output 420
113. lslivedump outputs 436
114. lslocaldisk output 447
115. lslocaldisk output 449
116. MDisk output 451
117. lsmdisklba command output 456
118. lspartnership attribute values 464
119. lspartnershipcandidate output 467
120. lsauthmultifactorverify output 477
121. lsauthmultifactorduo output 480
122. lsip output 486
123. lsportethernet output 488

124. lsportset output 490
125. lspartition output 498
126. lsprovisioningpolicy output 508
127. lssafeguardedpolicy output 514
128. lssafeguardedschedule output 515
129. lsbootdrive attribute values 520
130. lscmdstatus output 521
131. lsfcportfeature attribute values 523
132. lshardware attribute values 524
133. lsnodediscoverysubnet attribute values 525
134. sainfo lsnodeipconnectivity output 526
135. sainfo lsnodeipconnectivity error_data 527
136. lsservicenodes outputs 528
137. lsservicestatus output 529
138. lsauthsinglesignon output 581
139. lsvolumegrouppopulation output 591
140. lsvolumepopulation output 593
141. lsvolumegroupsnapshot output 594
142. lsvolumesnapshot output 596
143. lssnapshotpolicy output 604
144. lssnapshotschedule output 605
145. lsvolumegroupsnapshotpolicy output 606
146. lsvolumegroupsnapshotschedule output 608
147. Parameter differences for child pools and parent pools 613
148. Parameter differences for child pools and storage pools 623

149. lstruststore output....630
150. lstwopersonintegrityrequest output....634
151. lscurrentuser output....650
152. lsldap attribute values....651
153. lsldapserver attribute values....652
154. lsownershipgroup output....653
155. lsuser attribute values....653
156. lsusergrp output....655
157. testldapserver attribute values....664
158. Relationship between the syncrate value and the amount of data copied per second....673
159. Relationship between the syncrate value and the amount of data copied per second....685
160. lsmetadatavdisk output....694
161. Output from the lsvdiskanalysis command....708
162. lsvdiskanalysisprogress output....709
163. lsvdisklba command output scenarios....717
164. lsvolumebackup output....721
165. lsvolumebackupgeneration output....722
166. lsvolumebackupprogress output....724
167. lsvolumegroup output....725
168. lsvolumerestoreprogress output....727
169. Relationship between the syncrate value and the amount of data copied per second....736
170. lsvdisktiers attribute values....777
171. Accepted IP address formats....1020

About this guide

This publication provides information that helps you configure and use the IBM Storage Virtualize family. IBM Storage Virtualize family includes IBM SAN Volume Controller and FlashSystem family that is built with IBM Storage Virtualize software.

Who should use this guide

This guide is intended for system administrators or others who install and use the IBM Storage Virtualize family. IBM Storage Virtualize family includes IBM SAN Volume Controller and FlashSystem family that is built with IBM Storage Virtualize software.

Before you use the IBM Storage Virtualize family of products, you should have an understanding of storage area networks (SANs), the storage requirements of your enterprise, and the capabilities of your storage units.

Accessibility

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

Accessibility features

The following list includes the major accessibility features of the system:

- The IBM Documentation is provided in XHTML 1.0 format, which is viewable in most web browsers. With XHTML, you can view documentation according to the display preferences that are set in your browser. XHTML supports screen readers and other assistive technologies.
- All images in the IBM Documentation are provided with alternative text, so that visually impaired users can understand the contents of the images.
- This product uses standard Windows navigation keys.
- Interfaces are commonly used by screen readers.
- Industry-standard devices, ports, and connectors.

Keyboard navigation

You can use keys or key combinations for operations and to initiate menu actions that can also be done through mouse actions. You can go to the system online documentation from the keyboard by using the keyboard shortcuts for your browser or screen-reader software. See your browser or screen-reader software Help for a list of keyboard shortcuts that it supports.

IBM and accessibility

See the IBM® Human Ability and Accessibility Center for more information about the commitment that IBM has to accessibility

Emphasis

Different typefaces are used in this guide to show emphasis.

The following typefaces are used to show emphasis.

Emphasis	Meaning
Boldface	Text in **boldface** represents menu items.

© Copyright IBM Corp. 2003, 2025

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

Chapter 1. Setting up an SSH client

Secure Shell (SSH) is a client/server network application. It is used as a communication vehicle between the host system (for example, a laptop computer) and the system command-line interface (CLI).

Overview

The system acts as the SSH server in this relationship. If you require command-line access without entering a password, it uses the principles of public and private keys for authentication.

Authenticating SSH logins

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

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

Note: The authentication process for host systems that are not AIX is similar.

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

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

When SSH client (A) attempts to connect to SSH server (B), the SSH password (if you require command-line access without entering a password, the key pair) authenticates the connection. The key consists of two halves: the public keys and private keys. The SSH client public key is put onto SSH Server (B) using some means outside of the SSH session. When SSH client (A) tries to connect, the private key on SSH client (A) is able to authenticate with its public half on SSH server (B).

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

Note: After an SSH interactive session times out, session gets automatically closed. Session timeout limit is set to 15 minutes, by default. The limit value can be changed by using the CLI command. .

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

For multifactor authentication, IBM Security Verify communicates with the system and uses a PAM module to handle second factor authentication for SSH logins.

You can connect to the system by using the same username with which you log in to the system.

Setting up an SSH client on a Windows host

You can prepare the SSH client on a Windows host.

The workstation for the system include the PuTTY client program, which is a Microsoft Windows SSH client program. The PuTTY client program can be installed on your workstation in one of these ways:

- If you purchased the workstation hardware option from IBM, the PuTTY client program has been preinstalled on the hardware.

© Copyright IBM Corp. 2003, 2025

- You can use the workstation software installation CD to install the PuTTY client program.
- You can use the separate PuTTY client program-installation wizard, `putty-`*`version`*`-installer.exe`. You can download the PuTTY client program from this website:

  Download Putty

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

You can connect to the system by using the same username with which you log in to the system.

Generating an SSH key pair using PuTTY

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

About this task

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

Procedure

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

   Note: Leave the number of bits in a generated key value at 1024.
3. Click **Generate** and then move the cursor around the blank area of the Key section to generate the random characters that create a unique key.

   When the key has been completely generated, the information about the new key is displayed in the Key section.

   Attention: Do not modify the Key fingerprint or the Key comment fields; this can cause your key to no longer be valid.
4. Optional: Enter a passphrase in the **Key passphrase** and **Confirm passphrase** fields. The passphrase encrypts the key on the disk; therefore, it is not possible to use the key without first entering the passphrase.
5. Save the public key by:
   a) Click **Save public key**. You are prompted for the name and location of the public key.
   b) Type `icat.pub` as the name of the public key and specify the location where you want to save the public key.

      For example, you can create a directory on your computer called *keys* to store both the public and private keys.
   c) Click **Save**.
6. Save the private key by:
   a) Click **Save private key**. The PuTTYgen Warning panel is displayed.
   b) Click **Yes** to save the private key without a passphrase.
   c) Type `icat` as the name of the private key, and specify the location where you want to save the private key.

      For example, you can create a directory on your computer called *keys* to store both the public and private keys. It is recommended that you save your public and private keys in the same location.
   d) Click **Save**.
7. Close the PuTTY Key Generator window.

Configuring a PuTTY session for the CLI

You must configure a PuTTY session using the Secure Shell (SSH) password. If you require command line access without entering a password, use the SSH key pair that you created for the command-line interface (CLI).

About this task



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

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

Procedure

1. Select **Start** > **Programs** > **PuTTY** > **PuTTY**. The PuTTY Configuration window opens.
2. Click **Session** in the Category navigation tree. The Basic options for your PuTTY session are displayed.
3. Click **SSH** as the Protocol option.
4. Click **Only on clean exit** as the Close window on exit option. This ensures that connection errors are displayed.
5. Click **Connection** > **SSH** in the Category navigation tree. The options controlling SSH connections are displayed.
6. Click **2** as the Preferred SSH protocol version.
7. Click **Connection** > **SSH** > **Auth** in the Category navigation tree. The Options controller SSH authentication are displayed.
8. Click **Browse** or type the fully qualified file name and location of the SSH client and password. If no password is used, the private key in the **Private key file for authentication** field.
9. Click **Connection** > **Data** in the Category navigation tree.
10. Type the user name that you want to use on the system in the **Auto-login username** field.
11. Click **Session** in the Category navigation tree. The Basic options for your PuTTY session are displayed.
12. In the **Host Name (or IP Address)** field, type the name or Internet Protocol (IP) address of one of the system IP addresses or host names.
13. Type 22 26 in the **Port** field. The system uses the standard SSH port.
14. Type the name that you want to use to associate with this session in the **Saved Sessions** field. For example, you can name the session "System 1".
15. Click **Save**.

Results

You have now configured a PuTTY session for the CLI.

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

Connecting to the CLI using PuTTY

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

Note: Windows users can download PuTTY from the following website: Download Putty.

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

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

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

Note: The SSH server keys on the system are regenerated when a microcode load is performed on the clustered system. As a result, a challenge is sent because the fingerprint of the SSH server has changed.

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

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

Interactive mode

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

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

Note: After an SSH interactive session times out, session gets automatically closed. Session timeout limit is set to 15 minutes, by default. The limit value can be changed by using the CLI command. .

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

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

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

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

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

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

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

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

Single-line command

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

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

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

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

```
C:\Program Files\IBM\svcconsole\cimom>
 plink superuser@9.43.225.208 lsuser
 The server's host key is not cached in the registry. You
have no guarantee that the server is the computer you
think it is.
The server's key fingerprint is:
ssh-rsa 1024 e4:c9:51:50:61:63:e9:cd:73:2a:60:6b:f0:be:25:bf
If you trust this host, enter "y" to add the key to
PuTTY's cache and carry on connecting.
If you want to carry on connecting just once, without
adding the key to the cache, enter "n".
If you do not trust this host, press Return to abandon the
connection.
Store key in cache? (y/n) y
Authenticating with public key "imported-openssh-key"
id name      password ssh_key remote usergrp_id usergrp_name
0  superuser yes      yes     no     0          SecurityAdmin
1  smith     no       yes     no     4          Monitor
2  jones     no       yes     no     2          CopyOperator
```

Starting a PuTTY session for the CLI

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

Before you begin

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

About this task

Start a PuTTY session:

Procedure

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

   Note: If you are using the PuTTY application first time since you generated and uploaded the SSH password or key pair, a PuTTY Security Alert window is displayed. Click **Yes** to accept the change and trust the new key.
4. Type the *username* in the **login as** field and press **Enter**.

Setting up an SSH client on an AIX or Linux host

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

Before you begin

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

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

Linux operating systems
The OpenSSH client is installed by default on most Linux distributions.

About this task

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

Results

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

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

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

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

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

For more information, see “Connecting to the CLI with OpenSSH” on page 7.

Generating an SSH key pair using OpenSSH

You can generate an SSH key pair using OpenSSH.

About this task

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

Procedure

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

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

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

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

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

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

2. Associate the public key with a user on the clustered system using the management GUI.

Connecting to the CLI with OpenSSH

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

To connect to a clustered system by using a *username* and SSH password, issue the following command:

```
ssh username@my_system
```

To use an SSH key, issue the following command:

```
-i full_path_to_key
```

Where *my_system* is the name of the system IP, *full_path_to_key* is the full path to the generated key file, and *username* is the user name to connect to the system. Use port 26 to connect to the system.

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

Managing SSH keys by using the CLI

You can use the command-line interface (CLI) to work with Secure Shell (SSH) keys.

About this task

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

Note: After an SSH interactive session times out, session gets automatically closed. Session timeout limit is set to 15 minutes, by default. The limit value can be changed by using the CLI command.

Procedure

To work with SSH keys, complete these steps.

1. Issue the **lsuser** CLI command to list the SSH keys that are available on the system.
2. Create an SSH key on a host system and then copy the file to the system.
   For example, use `ssh-keygen` on a Linux server or use the PuTTY tool on a Windows server.
3. Issue the **mkuser** or **chuser** CLI command to install a new SSH key on the system.
   When you run the command, the SSH key is copied into system state and activated for the user, and the input file is deleted. Each key must be unique and can only be associated with one user on the system.
4. Use the **chuser** command or the **chcurrentuser** command to remove an SSH key from the system.

Example

To create an SSH key on a Linux host, run the ssh command.

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

Use secure copy to copy the file to the system.

```
scp /tmp/OU812.pub user_name@system_ip:/tmp
```

Run one of the following commands.

- To create a user ID and assign the SSH key to the user, run the mkuser command.

```
mkuser -name jane -usergrp Service -keyfile /tmp/OU812.pub
```

- To assign the SSH key to an existing user ID, run the chuser command.

```
chuser -keyfile /tmp/OU812.pub user_id_or_name
```

Chapter 2. Copying the software update files by using PuTTY pscp or openssh scp

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

Before you begin

Note: PuTTY pscp is useful in systems earlier than Microsoft Windows 10. Windows 10 includes openssh as standard so that scp works in a workstation or power shell window.

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

About this task

The pscp or scp application is available when you have an SSH client on your host system. You can access the pscp application through a Microsoft Windows command prompt. For Linux and Windows 10 users, the scp is part of the openssh packages.

Complete these steps to use the pscp application. The scp process is similar from step "6" on page 9.

Procedure

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

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

   Where *C:\Program Files\putty* is the directory where PuTTY is installed.
6. Issue this command to copy the package onto the node where the CLI runs:

   ```
   pscp -load saved_putty_configuration
   directory_software_upgrade_files/software_upgrade_file_name
   username@cluster_ip_address:/update
   ```

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

   Note: Saving the PuTTY configuration session in step "3" on page 9 and then loading the PuTTY configuration session in step "6" on page 9 is optional. To copy without loading a PuTTY configuration session, use the following syntax:

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

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

   a) Use pscp to copy data that you want to preserve from the /update directory.

   b) Use the following command to delete dump files in the /update directory:

© Copyright IBM Corp. 2003, 2025

```
cleardumps -prefix /home/admin/upgrade
```

c) Repeat step “6” on page 9.

Chapter 3. Using the dump commands to work with directories

The **lsdumps** command returns a list of dumps in a particular directory.

Dumps are contained in the following directory structure:

- /dumps
- /dumps/audit
- /dumps/cimom
- /dumps/elogs
- /dumps/feature
- /dumps/iostats
- /dumps/iotrace
- /home/admin/update
- /dumps/drive
- /dumps/enclosure
- /dumps/configs
- /dumps/debug
- /dumps/mdisk
- /dumps/syslogs
- /dumps/tejas

Use the **lsdumps** command with the optional **prefix** parameter to specify a directory. If you do not specify a directory, /dumps is used as the default. Use the optional *node_id_or_name* parameter to specify the node to list the available dumps. If you do not specify a node, the available dumps on the configuration node are listed.

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

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

An audit log tracks the action commands that are issued through an SSH session or from the management GUI. To list a specified number of the most recently audited commands, issue the **catauditlog** command. To dump the contents of the audit log to a file on the current configuration node, issue the **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 system. These files are produced during normal operations of the CIMOM.

Dumps that are contained in the /dumps/elogs directory are dumps of the contents of the error and event log at the time that the dump was taken. An error or event log dump is created by using the **dumperrlog** command. The command dumps the contents of the error or event log to the /dumps/elogs directory. If no file name prefix is supplied, the default errlog_ is used. The full default file name is errlog_*NNNNNN_YYMMDD_HHMMSS*, where *NNNNNN* is the node front panel name. If the command is used with the **-prefix** parameter, the prefix value is used instead of errlog.

Dumps that are contained in the /dumps/iostats directory are dumps of the per-node I/O statistics for disks on the system. An I/O statistics dump is created by using the **startstats** command. As

© Copyright IBM Corp. 2003, 2025

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 were collected are written to a file in the `/dumps/iostats` directory. The file names that are used for storing I/O statistics dumps are `Nm_stats_`*NNNNNN*`_YYMMDD_HHMMSS`, `Nv_stats_`*NNNNNN*`_YYMMDD_HHMMSS`, `Nn_stats_`*NNNNNN*`_YYMMDD_HHMMSS`, and `Nd_stats_`*NNNNNN*`_YYMMDD_HHMMSS`, where *NNNNNN* is the node name for the MDisk, volume, node, or drive.

Dumps that are contained in the `/dumps/iotrace` directory are dumps of I/O trace data. The type of data that is traced depends on the options that are specified by the **settrace** command. The collection of the I/O trace data is started by using the **starttrace** command. The I/O trace data collection is stopped when the **stoptrace** command is used. It is when the trace is stopped that the data is written to the file. The file name is *prefix*_*NNNNNN*_*YYMMDD*_*HHMMSS*, where *prefix* is the value that is entered for the **filename** parameter in the **settrace** command, and *NNNNNN* is the node name.

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

Dumps of support data from a disk drive are contained in the `/dumps/drive` directory. This data can help to identify problems with the drive, and does not contain any data that applications might have written to the drive.

Dumps from an enclosure or enclosures are contained in the `/dumps/enclosure` directory.

Dumps that are contained in the `/dumps` directory result from application abends. Such dumps are written to the `/dumps` directory. The default file names are dump.*NNNNNN.YYMMDD.HHMMSS*, where *NNNNNN* is the node front panel name. In addition to the dump file, there might be some trace files written to this directory that are named *NNNNNN*`.trc`.

Because files can only be copied from the current configuration node (by using secure copy), you can issue the **cpdumps** command to copy the files from a nonconfiguration node to the current configuration node.

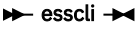
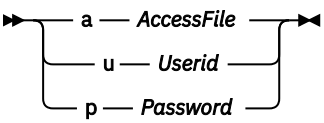
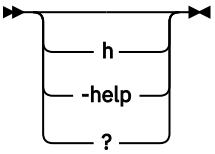
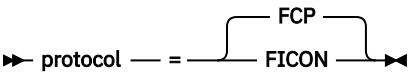
Chapter 4. CLI reference

Syntax diagrams

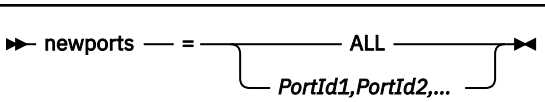
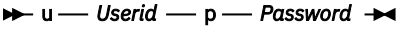
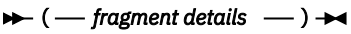
A syntax diagram uses symbols to represent the elements of a command and to specify the rules for using these elements.

Table 1 on page 13 explains how to read the syntax diagrams that represent the command-line interface (CLI) commands. In doing so, it defines the symbols that represent the CLI command elements.

Table 1. Syntax diagrams

Element	Syntax	Description
Main path line	>>><>() () ()	The main path line begins on the left with double arrowheads: >>. The main path line 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 must be written 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 must be written below the main path line. Mutually exclusive optional keywords are stacked vertically.
Default value		The default value must be written above the main path line.

© Copyright IBM Corp. 2003, 2025

Table 1. Syntax diagrams (continued)		
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.
Equalsign operator		Separates a name from its value in a name-value pair.
Syntax fragment	 **Fragment name** 	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.

Using wildcards in the CLI

You can use wildcards in the system command-line interface (CLI).

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

1. Running the command while logged onto the node.

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

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

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

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

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

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

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

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

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

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

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

Data types and value ranges

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

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

This table lists the data types and the value ranges for each.

Table 2. Data types

Data types	Value ranges
`filename_arg`	This is a (optionally fully qualified) file name, containing a maximum of 169 characters. Valid characters are: • . (period; the field must not start with, end with, or contain two consecutive periods) • / (forward slash) • - (hyphen) • _ (underscore) • a–z (lowercase letters, A through Z) • A–Z (uppercase letters, A through Z) • 0–9 (numerals 0 through 9)

Table 2. Data types (continued)

Data types	Value ranges
directory_or_file_filter	Specifies a directory, file name filter, or both, within the specified directory. Valid directory values are: • /dumps • /dumps/audit • /dumps/configs • /dumps/elogs • /dumps/feature • /dumps/iostats • /dumps/iotrace • /dumps/software The file name filter can be any valid file name, containing a maximum of 128 characters, with or without the "*" (wildcard), and appended to the end of a directory value. Valid characters are: • * (asterisk/wildcard) • . (the field must not start with, end with, or contain two consecutive periods) • / • - • _ • a–z • A–Z • 0–9
filename_prefix	The prefix of a file name, containing a maximum of 128 characters. Valid characters are: • - • _ • a–z • A–Z • 0–9

Table 2. Data types (continued)

Data types	Value ranges
name_arg	Names can be specified or changed using the **create** and **modify** functions. The view commands provide the name and ID of an object. **Note:** The system name is set when the system is created. The first character of a *name_arg* must be nonnumeric. The first character of an object name cannot be a – (dash) because the CLI (command-line interface) interprets it as being the next parameter. Valid characters are: • . (a period - the field must not start with, end with, or contain two consecutive periods) • / • - • _ • space • a through z • A through Z • 0 through 9
password	This is a user-defined password containing 6 - 64 printable ASCII characters.
serial_number	The serial number is 7 digits, the first two of which define the manufacturing location, leaving 5 digits for the product.
ip_address_arg	The argument follows the standard rules for dotted decimal notation. The following Internet Protocol 4 (IPv4) and Internet Protocol 6 (IPv6) address formats are supported: **IPv4 no port** 1.2.3.4 **IPv4 with specific port** 1.2.3.4:22 **Full IPv6, default port** 1234:1234:0001:0123:1234:1234:1234:1234 **Full IPv6, default port, leading zeros suppressed** 1234:1234:1:123:1234:1234:1234:1234 **Full IPv6 with port** [2002:914:fc12:848:209:6bff:fe8c:4ff6]:23 **Zero-compressed IPv6, default port** 2002::4ff6 **Zero-compressed IPv6 with port** [2002::4ff6]:23
dns_name	This is the dotted domain name for the system subnet.
hostname	The host name assigned to the system. This name can be different from the system name, and is modifiable. A combination of the host name and the *dns_name* is used to access the system, for example: https://hostname.

Table 2. Data types (continued)

Data types	Value ranges
capacity_value	The capacity expressed within a range of 512 bytes to 2 petabytes (PB). **Tip:** Specify the capacity as megabytes (MB), kilobytes (KB), gigabytes (GB), or PB. When using MB, specify the value in multiples of 512 bytes. A capacity of 0 is valid for a striped or sequential volume. The smallest number of supported bytes is 512.
node_id	A node ID differs from other IDs in that it is a unique ID assigned when a node is used to create a system, or when a node is added to a system. A *node_id* value is never reused in a system. Node IDs are internally represented as 64-bit numbers, and like other IDs, cannot be modified by user commands.
xxx_id	All objects are referred to by unique integer IDs, assigned by the system when the objects are created. All IDs are represented internally as 32-bit integers; node IDs are an exception. IDs in the following ranges identify the various types of objects: • 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 • fcmap_id: 0–4095 • rcrel_id: 0–8191 • controller_id: 0–63 **Note:** The io_group 4 exists but is used only in certain error recovery procedures. 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 working with objects.
xxx_list	A colon-delimited list of values of type *xxx*.
wwpn_arg	The Fibre Channel worldwide port name (WWPN), expressed as a 64-bit hexadecimal number and consisting of the characters 0–9, a–f, and A–F; for example: 1A2B30C67AFFE47B. **Note:** Entering `WWPN 0` in the command string causes a command failure.
panel_name	This is a string of up to six characters corresponding to the number on the printed label below the display on the front panel of a node in the system.
sequence_number	A 32-bit unsigned integer, expressed in decimal format.
csi_num_arg	A 32-bit unsigned integer, expressed in decimal format.

Table 2. Data types (continued)	
Data types	**Value ranges**
`percentage_arg`	An 8-bit unsigned integer, expressed in decimal 0–100 format.
`extent_arg`	A 32-bit unsigned integer, expressed in decimal format.
`num_extents_arg`	A 32-bit unsigned integer, expressed in decimal format.
`threads_arg`	An 8-bit unsigned integer, expressed in decimal format. Valid values are 1, 2, 3, or 4.
`velocity_arg`	The fabric speed in gigabytes per second (GBps). Valid values are 1 or 2.
`timezone_arg`	The ID as detailed in the output of the **`lstimezones`** command.
`timeout_arg`	The command timeout period. An integer from 0 to 600 (seconds).
`stats_time_arg`	The frequency at which statistics are gathered. Valid values are 1 to 60 minutes in increments of 1 minute.
`directory_arg`	Specifies a directory, file name filter, or both, within the specified directory. Valid directory values are: • `/dumps` • `/dumps/audit` • `/dumps/cimom` • `/dumps/configs` • `/dumps/elogs` • `/dumps/feature` • `/dumps/iostats` • `/dumps/iotrace` • `/home/admin/upgrade` The file name filter can be any valid file name, containing a maximum of 128 characters, with or without the wildcard (*, an asterisk), and appended to the end of a directory value. Valid characters are: • * • . (the field must not start with, end with, or contain two consecutive periods) • / • - • _ • a–z • A–Z • 0–9

Table 2. Data types (continued)	
Data types	**Value ranges**
`locale_arg`	The system locale setting. Valid values are: • 0 en_US: US English (default) • 1 zh_CN: Simplified Chinese • 2 zh_TW: Traditional Chinese • 3 ja_JP: Japanese • 4 fr_FR: French • 5 de_DE: German • 6 it_IT: Italian • 7 es_ES: Spanish
`key_arg`	A user-defined identifier for a secure shell (SSH) key, containing a maximum of 30 characters.
`user_arg`	Specifies the user: admin or service.
`copy_rate`	A numeric value of 0–100.

The maximum number of values entered into a colon-separated list is 128; exceeding this maximum number returns an error.

CLI flags

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

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

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

```
mkmdiskgrp -ext 16
```

it displays:

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

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

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

the following information is displayed:

```
6
```

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

CLI messages

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

When some commands complete successfully, textual output is normally provided. However, some commands do not provide any output. The phrase `No feedback` is used to indicate that no output is

provided. If the command does not complete successfully, an error is generated. For example, if the command has failed as a result of the cluster being unstable, the following output is provided:

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

CLI commands and parameters

Command-line interface (CLI) commands and parameters are represented in the syntax diagram.

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

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

- When a command name is specified, the first argument given must be the action that you want to be performed.
- Where you are specifying a command against a specific object, the object ID or name must be the last argument in the line.

A valid parameter meets the following requirements:

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

```
chuser -usergrp=-usergrp -- -password
```

- The symbol '--' is valid as the final word on the command line.

Examples that are valid

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

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

Examples that are invalid

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

CLI special characters

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

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

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

delimiters (: or , or !)
Delimiters are used to delimit items listed after issuing an information command.

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

CLI deprecated and discontinued commands

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

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

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

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

The following commands are deprecated:

- These commands are replaced by **`lseventlog`**:
  - **`caterrlog`**
  - **`caterrlogbyseqnum`**
  - **`lserrlogbyfcconsistgrp`**
  - **`lserrlogbyfcmap`**
  - **`lserrlogbyhost`**
  - **`lserrlogbyiogrp`**
  - **`lserrlogbymdisk`**
  - **`lserrlogbymdiskgrp`**
  - **`lserrlogbynode`**
  - **`lserrlogbyrcconsistgrp`**
  - **`lserrlogbyrcrelationship`**
  - **`lserrlogbyvdisk`**
- These commands are replaced by **`lsdumps`**:
  - **`ls2145dumps`**

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

The following commands are discontinued:

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

- **recoverdiskbycluster** is replaced by **recoverdiskbysystem**
- **svcservicemodeinfo lsclustervpd** is replaced by **satask lsservicestatus**
- **svcservicemodetask applysoftware** is replaced by **satask installsoftware**
- **svcservicemodetask cleardumps** is replaced by **cleardumps**
- **svcservicemodetask dumperrlog** is replaced by **dumperrlog**
- **svcservicemodetask exit** is replaced by **stopservice**
- **setclustertime** is replaced by **setsystemtime**
- **stopcluster** is replaced by **stopsystem**
- **triggermdiskdump** is replaced by **triggerdrivedump**
- **setevent** is replaced by **chsnmpserver**, **lssnmpserver**, **mksnmpserver**, and **rmsnmpserver**
- **setemail** is replaced by **chemail**, **chemailserver**, **lsemailserver**, **mkemailserver**, and **rmemailserver**
- **svctask chkeyserverkeysecure** is replaced by **svctask chkeyserverciphertrustmanager**.
- **svcinfo lskeyserverkeysecure** is replaced by **svcinfo lskeyserverciphertrustmanager**.

Understanding capacity indicators

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

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

Table 3. Capacity indicators. This table displays the differences in how capacity indicators are displayed in the management GUI and the CLI.

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

Attributes of the -filtervalue parameters

The **-filtervalue** parameter filters a view that is based on specific attribute values that relate to each object type. You can combine multiple filters to create specific searches, for example, **-filtervalue name=fred:status=online**. The help (**-filtervalue**) specifies the attributes that are available for each object type.

The **-filtervalue** parameter must be specified with *attrib=value*. The **-filtervalue?** and **-filtervalue** parameters cannot be specified together.

Note: The qualifier characters less than (<) and greater than (>) must be enclosed within double quotation marks (" "). For example, **-filtervalue vdisk_count "<"4**. It is also valid to include the entire expression within double quotation marks. For example, **-filtervalue "vdisk_count<4"** .

When an attribute requires the **-unit** parameter, it is specified after the attribute. For example, **-filtervalue capacity=24 -unit mb**. The following input options are valid for the **-unit** parameter:

- b (bytes)
- kb (Kilobytes)
- mb (Megabytes)
- gb (Gigabytes)
- tb (Terabytes)
- pb (Petabytes)

Capacity values displayed in units other than bytes might be rounded. When filtering on capacity, use a unit of bytes, **-unit b**, for exact filtering.

You can use the asterisk (*) character as a wildcard character when names are used. The asterisk character can be used either at the beginning or the end of a text string, but not both. Only one asterisk character can be used in a **-filtervalue** parameter.

UNIX commands available on the CLI

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

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

Note: After an SSH interactive session times out, session gets automatically closed. Session timeout limit is set to 15 minutes, by default. The limit value can be changed by using the CLI command. .

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

Table 4. UNIX commands for interactive SSH sessions

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

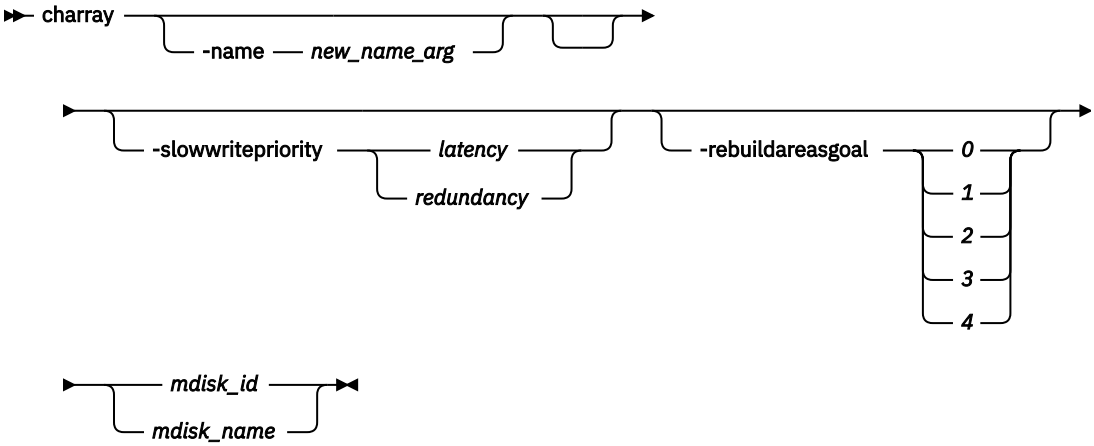
Chapter 5. Array commands

Use the array commands to manage arrays and their properties.

charray

Use the **charray** command to change array attributes.

Syntax



Parameters

-name *new_name_arg*
(Optional) Specifies the new name to apply to the array MDisk.

-slowwritepriority latency | redundancy
(Optional) Controls array ability to complete write operations that take too long, even if it temporarily compromises redundancy.

The value can be either `latency` or `redundancy`:

- `latency` implies the feature is enabled for normal I/O operations
- `redundancy` implies the feature is not enabled for normal I/O operations

The default value is `latency` mode for existing arrays.

Important: An array can cause member drives to become unsynchronized (to preserve response time) if the value is `latency`. If the value is `redundancy`, the array cannot cause member drives to become unsynchronized (to preserve time) and I/O performance is impacted.

-rebuildareasgoal *0 | 1 | 2 | 3 | 4*
(Optional) Specifies the rebuild areas threshold. The array logs an error when the available rebuild areas drop below this specified threshold. The values are `0`, 1, 2, 3, or 4. (If you specify `0`, an error is not logged if the system runs out of rebuild areas.)

Note: This parameter is only applicable for distributed arrays.

mdisk_id | mdisk_name
(Required) Identifies (by ID or user-defined name) which array the MDisk command applies to.

Description

This command changes an array's attributes.

© Copyright IBM Corp. 2003, 2025

An invocation example to change the name of an array

```
charray -name raid6 mdisk0 0
```

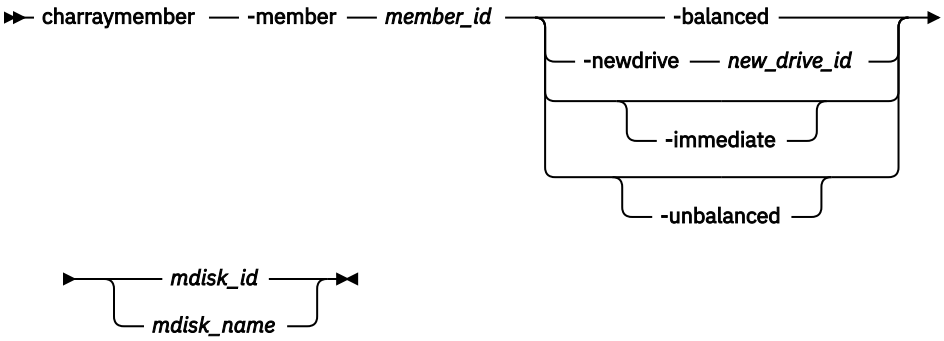
The resulting output:

```
No feedback
```

charraymember

Use the **charraymember** command to modify an array member's attributes, or to swap (exchange) a member of a RAID array with that of another drive.

Syntax



Parameters

-member *member_id*
Identifies the array member index.

-balanced
(Optional) Forces the array member spare goals to be set to the:

- Present array member goals
- Existing exchange goals
- The newDrive goals

Note: If **-balanced** is specified and the goal of the associated array MDisk changes, the tier of the array MDisk is updated to match the new goal.

-newdrive *new_drive_id*
(Optional) Identifies the drive to add to the array.

For distributed arrays:

- If the **-newdrive** parameter is specified with the **-immediate** parameter, and the member is not associated with a rebuild area, the command begins a distributed rebuild to a rebuild area in such a way that a copyback begins immediately when the rebuild finishes.
- If the **-newdrive** parameter is specified and the member is already associated with a rebuild area, the array configures itself to use the new member, and might begin a copyback. It occurs whether the **-immediate** parameter is specified or not.
- If the **-immediate** parameter is not set and the **-newdrive** parameter is set (but the array member is not allocated to a rebuild area), the command fails. If the **-balanced** parameter is set, the command fails.
- The **-newdrive** parameter is not allowed when an array expansion is in progress.

-immediate
(Optional) Specifies that the old disk is to be immediately removed from the array, and the new disk rebuilt. If you do not choose this option, exchange is used; it preserves redundancy during the rebuild.

-unbalanced
(Optional) Forces the array member to change if the newDrive does not meet array member goals.

mdisk_id
(Either the ID or the name is required) Identifies which ID array the MDisk command applies to.

mdisk_name
(Either the ID or the name is required) Identifies which name array the MDisk command applies to.

Description

This command modifies an array member's attributes, or swaps a member of a RAID array with a member of another drive.

Remember: You cannot create an array that contains an NVMe drive and another drive of a different technology type or transport protocol. If the member is an online compressed drive, then upon completion of the exchange, it is formatted and then becomes a candidate drive.

This table shows the command combination options.

*Table 5. **charraymember** combination options*

Option	Description
-balanced	• Member goals are set to the properties of the existing member or exchange drive. • The command fails if the member is not populated with a drive. • Member goals are set to the properties of the current member drives being exchanged into the array count as members. • If no exchange exists, the existing member drive goals are used.
-newdrive drive_id	• The command processes the exchange, and does NOT update the member goals. • You must specify a new drive that is an exact match for the member goals. • The command fails if the drive is not an exact match.
-newdrive drive_id -balanced	The command processes the exchange and updates the member goals to the properties of the new drive.
-newdrive drive_id -unbalanced	• The command processes the exchange and does NOT update the member goals. • It is only permitted when the array is degraded and the member is empty. • It means **-immediate** is not needed. The exchange is always immediate. • Later, if drives are a sufficient member goal match, the array rebalance selects those drives. • A balancing exchange restarts the member goals.

An invocation example to swap a spare or candidate drive for a member 0 drive by using exchange

```
charraymember -member 0 -newdrive 4 mdisk2
```

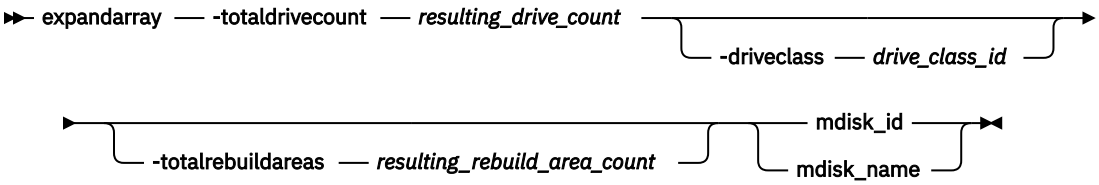
The resulting output:

```
No feedback
```

expandarray

Use the **expandarray** command to add new drive capacity to an existing distributed array MDisk.

Syntax



Parameters

totaldrivecount
(Required) The resulting drive count after expansion is completed. This value must be an integer in the range of 5 and 128. For distributed RAID 1 arrays, this value must be an integer in the range of 3 and 16.

- Up to 48 NVMe drives in a distributed RAID 6 array.
- For (any) distributed RAID 6 arrays it is not allowed to add more than 42 members drives in one **expandarray** invocation.
- For (any) distributed RAID 1 arrays it is not allowed to add more than 12 member drives in one **expandarray** invocation.

driveclass
(Optional) The drive class (as specified by **lsdriveclass**) to be used for selecting new member drives. If you do not specify a drive class, the default drive class is the one quoted in **lsarray** for the distributed array.

totalrebuildareas
(Optional) The resulting rebuild area count after expansion is completed. The additional rebuild area value must be included in **totaldrivecount**. This value must be an integer in the range of 1 and 4. For distributed RAID 1 arrays, this value must be 1. Values 2, 3, and 4 are not supported. For distributed RAID 1 arrays, this value must be 1. Values 2, 3, and 4 are not supported. If you do not specify a value, the rebuild area count for the array remains unchanged, even if the total drive count exceeds a threshold where the default array creation behavior would result in creating an extra rebuild area.

Note: You cannot expand a distributed RAID array of compressing drives to have more than one rebuild area.

mdisk_id
(Optional) The identity of the array MDisk. Required if **mdisk_name** is not specified.

mdisk_name
(Optional) The name of the array MDisk. Required if **mdisk_id** is not specified.

Description

This command starts the array expansion task.

Important: The **-totaldrivecount** value includes any extra drives that are added as rebuild areas.

This list describes some items to be aware of when you expand an array:

- Multiple expansions are not configured and queued. The maximum number of active expansions in the system is four and there can be only one in a pool at one time.
- The command uses the **totaldrivecount/totalrebuildareas** to ensure that if a mistake occurs and the current size of the array is entered then the command has no effect.

- If no drive class is specified, then the drive class of the array you are expanding is used. If a drive class is specified, then it is validated to ensure it is equivalent or superior to the drive class of the array you are expanding.
- The array must have no used rebuild areas or background tasks in progress or the command fails.
- The progress of the expansion can be monitored by using the **lsarrayexpansionprogress** command.
- An expansion cannot be canceled.
- Volume copies with sequential (seq) policy that are allocated from the expanding MDisk change to striped as part of the expansion migration.
- Expansion can have an impact on host performance and latency when performed on an array that has over 50% load.

An expansion is suspended if any of the following conditions exist:

- The expanding MDisk goes offline.
- There is a read error during migration activity.
- It is necessary (but not possible) to record a virtual medium error.
- An out of space condition occurs on an MDisk that is used as an expansion migration target.

Invocation examples

```
svctask expandarray -totaldrivecount 8 mdisk3                       // Expand mdisk3 to 8 drives
from the array's driveclass
svctask expandarray -totaldrivecount 6 -totalrebuildareas 2 32      // Expand mdisk 32 to 6
drives, 2 of which being rebuild areas
svctask expandarray -driveclass 0 -totaldrivecount 16 48            // Expand mdisk 48 to 16
drives from drive class 0
```

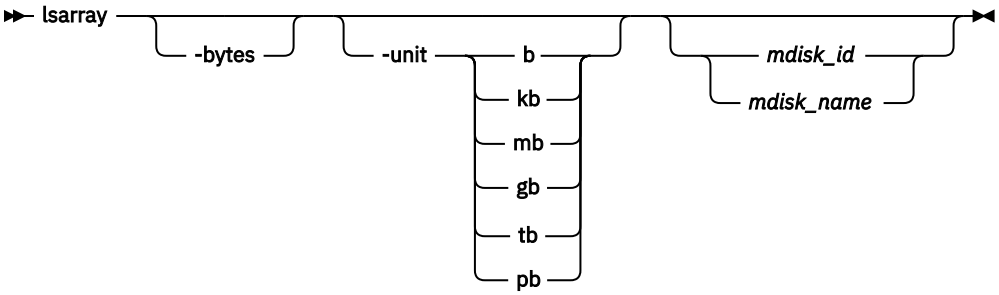
The resulting output:

```
None
```

lsarray

Use the **lsarray** command to list the array MDisks.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-bytes
(Optional) Requests output of capacities in bytes (instead of rounded values).

-filtervalue *capacity=value*
(Optional) Specifies a list of one or more filter capacities (size) values matching the specified values for the **unit** parameter. Use the **unit** parameter to interpret the value for size or capacity.

-unit *b | kb | mb | gb | tb | pb*
(Optional) The units that are used when you specify the **-filtervalue** capacity, where:

- *b* = bytes
- *kb* = 1,024 bytes
- *mb* = 1,048,576 bytes
- *gb* = 1,073,741,824 bytes
- *tb* = 1,099,511,627,776 bytes
- *pb* = 1,125,899,906,842,624 bytes

mdisk_id
(Optional) The identity of the array MDisk.

mdisk_name
(Optional) The name of the array MDisk.

Description

This command returns a concise list or a detailed view of array MDisks visible to the system. This table provides the attribute values that can be displayed as output view data.

Table 6. Array output

Attribute	Values
capacity	Indicates the value for the capacity you specify by using the **-unit** parameter.
status	Indicates the array status. The values are: • online • offline • excluded • degraded (applies only to internal MDisks)
mode	Indicates the mode. The values are: • `unmanaged` • `managed` • `image` • `array`
block_size	Indicates the block size. The value is 512 bytes (or blank) in each block of storage.

Table 6. Array output (continued)	
Attribute	**Values**
raid_status	Indicates the RAID status. The values are: **offline** The array is offline on all nodes. **degraded** The array has deconfigured or offline members; the array is not fully redundant. **syncing** The array members are all online. The array is synchronizing parity or mirrors to achieve redundancy. **initializing** The array members are all online. The array is initializing; the array is fully redundant. **online** The array members are all online and the array is fully redundant. **expanding** The array members are all online and the array is expanding.
raid_level	Indicates the RAID level of the array.
redundancy	Indicates the number of member disks that can fail concurrently without causing the array to fail.
rebuild_ahead	Indicates the `rebuild_ahead` specification: • `off` indicates that the array cannot rebuild ahead. • `minimum` indicates that the array can rebuild ahead of a limited number of RAID stripes. • `medium` indicates that the array can rebuild ahead of a moderate number of RAID stripes. • `maximum` indicates that the array can rebuild ahead of many RAID stripes.
rebuild_priority	Indicates the rebuild priority: • `hostmax` indicates that the maximum priority is given to the host operations to minimize the impact of the rebuild operation. • `hostweighted` indicates that host I/O operations are given higher priority than rebuild operations. • `rebuildweighted` indicates that rebuild operations are given higher priority than host I/O operations. • `rebuildmax` indicates that maximum priority is given to the rebuild operations (over host I/O operations).
scrub_rate	Indicates the speed of the background array sweeper. The value must be a number in the range `0` - `100`: • `0` turns off the sweeper. • `100` consumes the most bandwidth.
scrub_progress	Indicates the progress (percentage) of the background array sweeper. The value is reset to `0` after it reaches `100`, and another sweep begins.
flashsystem	Indicates the type of array for this system. The values are `yes` and no.

Table 6. Array output (continued)	
Attribute	**Values**
spare_goal (Deprecated)	The value of this attribute is always `blank`.
spare_protection_min (Deprecated)	The value of this attribute is always `blank`.
distributed (Deprecated)	Indicates whether the array is distributed. The value of this attribute is always `yes`.
effective_used_capacity	The amount of logical data that is written to the MDisk. If this MDisk is not `over_provisioned`, then this field is blank.
allocated_capacity	The amount of logical capacity that is allocated (by a pool) from this array for volume data and quorum.
fips	Indicates if the array is FIPs only mode for drives. The values are `yes` and `no`.

This list defines the status fields:

online
The MDisk is online and available.

degraded
(Internal MDisks only) The array has members that are degraded, or the `raid_status` is degraded.

degraded_ports
There are one or more MDisk port errors.

degraded_paths
One or more paths to the MDisk are lost; the MDisk is not online to every node in the system.

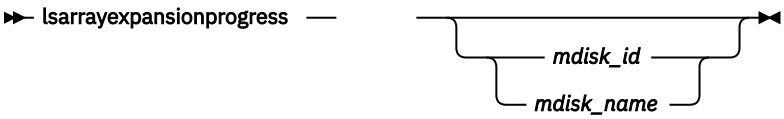
offline
All paths to the MDisk are lost.

excluded
The MDisk is excluded from use by the system; the MDisk port error count exceeded the threshold.

lsarrayexpansionprogress

Use the **lsarrayexpansionprogress** command to display the array expansion task status for one or more arrays.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

mdisk_id
(Optional) The identity of the array MDisk.

mdisk_name
(Optional) The name of the array MDisk.

Description

This command returns a concise list of the array expansion task status for one or more arrays. If the ID or name parameter is not set, then all array MDisks currently expanding are displayed. If the parameter is set, then the view is displayed for that MDisk regardless if it is expanding or not. This table provides the attribute values that can be displayed as output view data.

Table 7. Output

Attribute	Values
mdisk_id	Indicates the ID of the MDisk.
mdisk_name	Indicates the name of the MDisk.
progress	Indicates the progress of the expansion. This value is a number in the range of 0 - 100 or blank if no expansion is in progress.
estimated_completion_time	Indicates the estimated completion time for the expansion. This value is in the format of YYMMDDHHMMSS or blank if no expansion is in progress or the expansion is suspended.
status	Indicates the current status of the expansion.
target_capacity	Indicates the capacity of the array after the expansion process is completed. This value is in GB with two decimal places.
additional_capacity_remaining	Indicates the capacity that is added to the array (target capacity - current capacity). This value is in GB with two decimal places.
drives_added	Indicates the number of drives added.

This list defines the status fields:

blank
No expansion in progress.

inactive
No expansion in progress.

running
The expansion activity is proceeding normally.

offline
The expansion activity is suspended because the MDisk is offline.

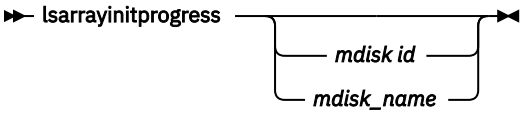
suspended
The expansion activity is suspended because of an error. Fix the migration errors against the MDisk.

syncing
The expansion is suspended because the array is syncing.

lsarrayinitprogress

Use the **lsarrayinitprogress** command to view the progress of array background initialization that occurs after creation.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

mdisk_id
(Optional) The identity of the array MDisk.

mdisk_name
(Optional) The user-defined MDisk name.

Description

This command shows the progress of array background initialization. This table shows the possible outputs.

Note: Non-volatile memory express (NVMe) arrays can handle faster initialization. Thus, for NVMe arrays, initialization occurs only at drive recovery and expansion, not at the time of array creation.

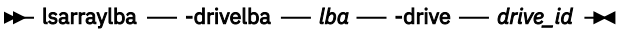
Table 8. lsarrayinitprogress output

Attribute	Value
progress	The percentage of initialization task that is completed.
estimated_completion_time	The expected initialization task completion time, in YYMMDDHHMMSS format.

lsarraylba

Use the **lsarraylba** command to permit an array logical block addressing (LBA) to be found from a drive and LBA.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-drivelba ***lba***
The LBA on the drive to convert to the array LBA. The *lba* value must be specified in hex, with a `0x` prefix.

-drive ***drive_id***
The ID of the drive to view.

Description

This command permits an array LBA to be found on a drive and LBA.

The system provides volumes that have LBAs for 512-byte block sizes; however, back-end disks that have a block size of either 512 or 4096 bytes can also be used. Drives are listed in their physical size.

Use the **lsdrive** command to display the drive block size, and use the **lsdrive** or `lsarray` command to list each object (the drive and the MDisk).

During an expansion, if **lsarraylba** is used on a drive that is added to the array and the LBA specified on the command is below the crossover between the old and new geometry regions of the array then this LBA is not allocated to the array. The command output lists this LBA as `unallocated`. Continued expansion activity eventually allocates this LBA.

This table shows possible outputs.

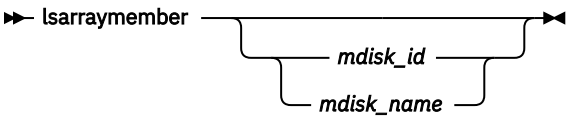
Table 9. ***lsarraylba*** *output*

Attribute	Value
type	The type of MDisk extent allocation: • `allocated` • `unallocated` For distributed arrays only: • If the LBA is an unused rebuild area, this value displays `rebuild_area`. • If the LBA is a used rebuild area, this value displays `allocated`.
mdisk_lba	The LBA on the array MDisk (blank if none).
mdisk_start	The start of range of LBAs (strip) on the array MDisk (blank if none).
mdisk_end	The end of range of LBAs (strip) on the array MDisk (blank if none).
drive_start	The start of range of LBAs (strip) on the drive (blank if none).
drive_end	The end of range of LBAs (strip) on the drive (blank if none).

lsarraymember

Use the **lsarraymember** command to list the member drives of one or more array MDisks.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

mdisk_id
(Optional) The identity of the array MDisk.

mdisk_name
(Optional) The MDisk name that you provided.

Description

This command lists the member drives of one or more array MDisks. It describes positions within an array occupied by a drive. The positions determine how mirroring the RAIDs takes place.

This table shows the potential output for this command.

Table 10. lsarraymember output

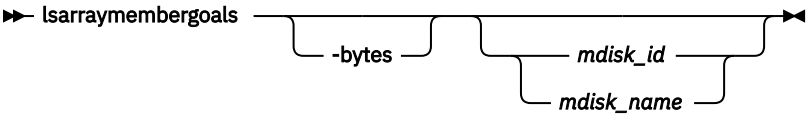
Attribute	Value
member_id	Specifies the identity of the array member. It represents drive order in RAID array
drive_id	Specifies the identity of the drive for member ID, or the source drive if an exchange is in progress. It is blank if there is no drive that is configured.
new_drive_id	Specifies the ID of the drive that is exchanged with this member ID. It is blank if there is no ID.
spare_protection	Specifies the number of non-degrading spares for the member. This includes spare drives with different attributes from the array member goals that perform equally or better than the array member goals. For distributed array members this field is blank.

Table 10. lsarraymember output (continued)	
Attribute	**Value**
balanced	For distributed arrays, this value indicates whether a superior drive class is being used for the array: • `exact` indicates that the same drive class is in use. • `yes` indicates that the drive exceeds the array drive class. For distributed arrays, the array must also be balanced to its rebuild area goals.
slow_write_count	Indicates the number of times this member becomes unsynchronized because of high response time on write I/O operations.
slow_write_time_last	Creates a timestamp of when the component last became unsynchronized. The time format is `YYMMDDhhmmss` in clustered system time. No time is indicated if the value for `slow_write_count` is `0`.

lsarraymembergoals

Use the **lsarraymembergoals** command to list the spare goals for member drives of one or more array MDisks.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

Description

This command lists the spare goals for member drives of one or more array MDisks. This table provides the potential output for this command.

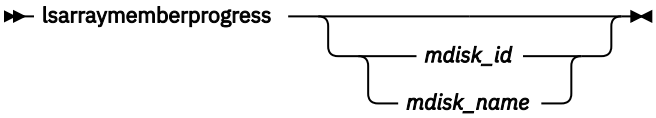
Table 11. ***lsarraymembergoals*** *output*	
Attribute	**Values**
member_id	Indicates the ID of the array member that represents the drive order in the RAID array.
drive_id	Indicates the ID of the drive for the member ID (it is blank if none are configured).

Table 11. **lsarraymembergoals** output (continued)	
Attribute	**Values**
capacity_goal	Indicates the capacity goal for the array member (it is the same for all members in the array).
tech_type_goal	Indicates the technology goal for the array member: • tier0_flash • tier1_flash • tier_enterprise • tier_nearline • tier_scm
RPM_goal	Indicates the drive RPM characteristic that the array member must have (it is blank for flash drives).
enclosure_id_goal	Indicates the ID of the enclosure that must contain the array member (it is blank if no enclosure is selected).
slot_id_goal	Indicates the ID of the slot in the enclosure that must contain the array member.
node_id_goal	Indicates the ID of the node that must contain this array member.
block_size_goal	Indicates the array member block size. The value is either 512 or 4096. This value is the same for all member drives in the array, and is the smallest value for the block size of one of the original drives or the set of drives in the array when it is set to be balanced.
drive_class_id_goal	Indicates the preferred drive class for this array member.

lsarraymemberprogress

Use the **lsarraymemberprogress** command to display array member background process status.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

mdisk_id
(Optional) The identity of the array MDisk.

mdisk_name
(Optional) The MDisk name that you provided.

Description

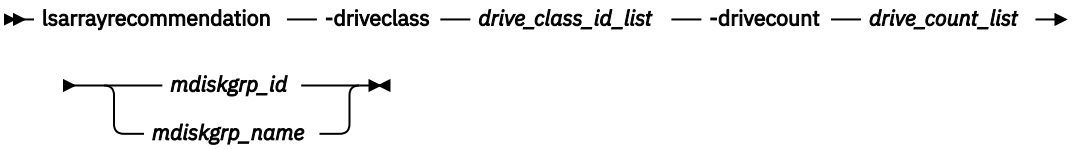
This command displays array member background process status. Exchange cannot start on a rebuilding member because both component rebuild and exchange are shown in the same view. This table provides the potential output for this command.

*Table 12. **lsarraymemberprogress** output*	
Attribute	**Value**
member_id	Indicates the array member index.
drive_id	Indicates the ID of the drive.
task	Indicates the identity of task that is being performed by the array member: • `rebuild` indicates that the array is recovering all the data on the component (after it was removed). • `exchange` indicates that the component is copying data to another drive. • `resync` indicates that this member is unsynchronized and is performing write operations that were completed early.
new_drive_id	The identity of drive that is being exchanged.
progress	Indicates the task percentage completion.
estimated_completion_time	Indicates the expected task completion time in the format YYMMDDHHMMSS. It is blank if completion time is unknown.

lsarrayrecommendation

Use the **lsarrayrecommendation** command to view a recommended configuration for the specified drive class and number of drives.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-driveclass *drive_class_id_list*
(Required) Specifies the drive class, or classes, for which the array recommendation is made. You must specify at least 1 *drive_class_id_list* value. You can specify a total of 32 *drive_class_id_list* values on a single command; however, you must separate each value with a colon character (:).

-drivecount *drive_count_list*
(Required) Specifies the number of drives for which to make recommendation. You must specify at least 1 *drive_count* value. You can specify a total of 32 *drive_count* values on a single command; however, you must separate each value with a colon character (:).

Remember: Each drive_count value must be a numerical value between 2 and 128 (only redundant arrays are considered).

-mdiskgrp *mdiskgrp_id* | *mdiskgrp_name*
(Required) Specifies the ID or name of the pool for which to make the recommendation.

Description

This command displays the system-recommended array configuration for a specific drive class and number of drives.

This table provides the attribute values that can be displayed as output view data.

Table 13. ***lsarrayrecommendation*** *output*

Attribute	Possible Values
mdiskgrp_id	Indicates the MDisk group ID.
mdiskgrp_name	Indicates the MDisk group name.
drive_class_id	Indicates the drive class ID for the recommendation.
raid_level	Indicates the recommended RAID level.
distributed (Deprecated)	Indicates whether the array is a distributed array recommendation. The value of this attribute is always yes.
min_stripe_width	Indicates the stripe width minimum.
max_stripe_width	Indicates the stripe width maximum.
stripe_width	Indicates the recommended stripe size for this storage pool.
rebuild_areas	Indicates the recommended number of rebuild areas.
strip_size	Indicates the recommended strip size for this storage pool.
drive_count	Indicates the number of drives to configure per array.
array_count	Indicates the number of arrays that can be created at the recommended drive count.
capacity	Indicates the expected usable capacity for virtualization for the array or arrays.

Note: The **lsarrayrecommendation** command will not recommend a distributed RAID array of compressing drives, such as FlashCore Modules, if another distributed RAID array of compressing drives is already in the same storage pool. To add more storage, you can expand the existing distributed RAID array or add a new distributed RAID array to another storage pool.

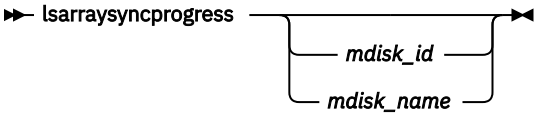
Note: The default recommendation for each drive class is the first row that is listed in the output for that drive class. All other rows for each drive class are not sorted by order of recommendation.

Note: The **lsarrayrecommendation** command will not recommend an array if the total effective capacity of the arrays in the system would exceed the supported effective capacity limit for the system.

lsarraysyncprogress

Use the `lsarraysyncprogress` command to display how synchronized a RAID array is.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

mdisk_id
(Optional) The ID of the MDisk you want to view.

mdisk_name
(Optional) The user-defined name of the MDisk you want to view.

Description

This command shows you how synchronized a RAID array is. It includes internal activity that is working toward a fully synchronized array. This table provides the potential output.

Table 14. ***lsarraysyncprogress*** *output*

Attribute	Value
progress	The percentage of the array that is synchronized.
estimated_completion_time	The expected synchronization completion time (YYMMDDHHMMSS; blank if completion time unknown).

lscompatibledriveclasses

Use the **lscompatibledriveclasses** command to display all compatible drive class IDs for a given existing drive class.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

driveclass_id
(Optional) The identity of the drive class.

Description

This command lists the IDs of the drive classes that are compatible with the given drive class. This information is used in conjunction with the **lsdriveclass** command to ascertain additional details. This table provides the attribute values that can be displayed as output view data.

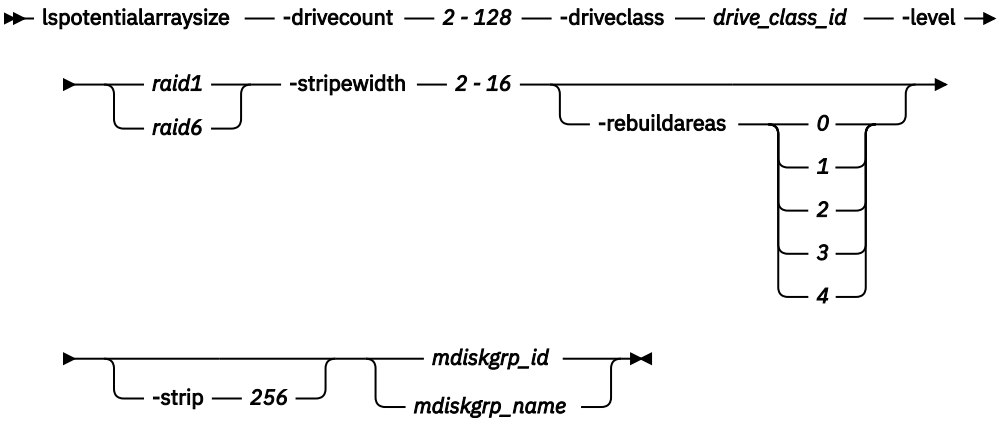
Table 15. Output

Attribute	Values
id	Indicates the drive class IDs that are compatible with the specified ID.

lspotentialarraysize

Use the **lspotentialarraysize** command to display the size of a potential array for a specified drive count, drive class, and RAID level in the specified MDisk group.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-drivecount *2 - 128*
(Required) Specifies the number of the drives. The value must be a number in the range 2 - 128. The value for distributed RAID-1 arrays must be in the range 2 - 16.

-driveclass *drive_class_id*
(Required) Specifies the drive class. The *driveclass_id* value must be a number.

-level ***raid1***

| raid6

(Required) Specifies the RAID level for the array that is being created. The possible values are:

- *raid1*
- *raid6*
-

-stripewidth ***2 - 16***
(Required) Indicates the width of a single unit of redundancy within a distributed set of drives. The value must be:

- RAID-1: 2
- RAID-6: 5 - 16

-rebuildareas ***0 | 1 | 2 | 3 | 4***
(Optional) Specifies the number of rebuild areas in the array. This value must be in the range 1 - 4 for RAID-6 arrays.

For distributed RAID-1 arrays, the value must be either 0 or 1 with these restrictions:

- 0 is valid only if **drivecount** is 2.
- 1 is valid only for **drivecount** values in the range 3 - 16.

-strip ***256***
(Optional) Specifies sets the strip size in KiB for the array that is being configured. The value is 256.

mdiskgrp_id | mdiskgrp_name
(Required) Indicates the MDisk array ID or name.

Description

This command displays the size of a potential array for a specified drive count, class, and RAID level.

You can use this command to define potential sizes for distributed arrays. (Distributed array descriptions are triggered by using **-rebuildareas**.) This command assists with the configuration options that are provided during array creation, and estimates the array capacity if it were to be configured

This table provides the attribute values that can be displayed as output view data.

Table 16. ***lspotentialarraysize*** *output*

Attribute	Possible Values
capacity	Indicates the expected usable capacity for virtualization for the array or arrays.

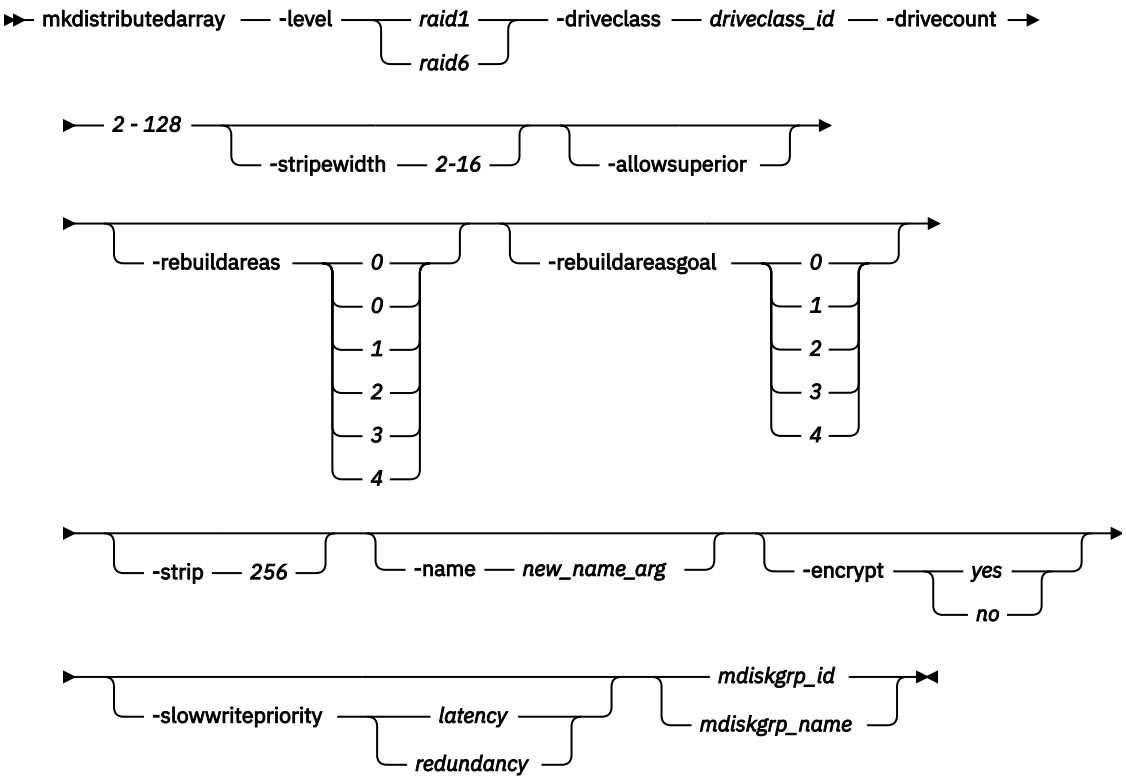
mkarray (Deprecated)

The **mkarray** command has been deprecated. Use the **mkdistributedarray** command to create distributed arrays.

mkdistributedarray

Use the **mkdistributedarray** command to create a distributed array and add it to a storage pool.

Syntax



Parameters

-level *raid1 | raid6*
(Required) Specifies the RAID level for the array that is being created. The possible values are:

- *raid1*
- *raid6*

-driveclass *driveclass_id*
(Required) Specifies the class that is being used to create the array. The *driveclass_id* must be a numeric value (specified with the **lsdriveclass** command).

-drivecount *2 - 128*
(Required) Specifies the number of the drives. The value must be a number in the range 2 - 128. The value for distributed RAID 1 arrays must be in the range 2 - 16 for SSDs and SCMs. The value for distributed RAID 1 arrays must be in the range 3 - 16 for HDDs.

-stripewidth *2-16*
(Optional) Indicates the width of a single unit of redundancy within a distributed set of drives. The value must be:

- RAID 1: 2
- RAID 6: 5 - 16

The default width for RAID 6 is 12 . The width plus the number of rebuild areas must be less than or equal to the drive count.

If **-stripewidth** is not specified and **drivecount** - **rebuildareas** is less than the default stripewidth (12 for RAID 6), then **stripewidth** is set to **drivecount** - **rebuildareas**. If **drivecount** - **rebuildareas** is less than the minimum allowed (5 for RAID 6), the command fails.

-allowsuperior
(Optional) Specifies that you can use drives that are not an exact match to the drive class used when creating the array (such as drives that use different capacity or technology). The system attempts to select the closest match to the class when satisfying the drive count. You can select higher capacity members of the same technology type before you select higher technology members.

Note: For a drive A to be considered superior to drive B, these situations must be true:

1. Drives A and B are `use=candidate`
2. Drives A and B are in the same I/O group.
3. Drive A's speed is equal to or greater than drive B's. Solid-state drives (SSDs) are higher speed than all hard disk drives (HDDs). For SSDs, **tech_type** *tier0_flash* is considered to have better performance than *tier1_flash*.
4. Drive A's capacity is equal to or greater than drive B's.
5. Drive A has a block size that is smaller than or equal to drive B.
6. Drives A and B have the same transport protocol.
7. Drives A and B are both not compressing drives or Drives A and B are both compressing drives of the same physical capacity and logical capacity.

-rebuildareas ***0 /1 | 2 | 3 | 4***
(Optional) Specifies the overhead capacity that is distributed across all drives available to an array. This capacity restores data after a drive failure. The values are:

- `0`
- `0`
- `1`
- `2`
- `3`
- `4`

The value is 1 - 4 (inclusive) for RAID 6 arrays.

For RAID 1 arrays, only the values 0 (when total drive count is 2) and 1 (when total drive count is in the range 3 - 16) are valid. The value is 1 - 4 (inclusive) for RAID 6 arrays. For RAID 1 arrays, only the values 0 (when total drive count is 2) and 1 (when total drive count is in the range 3 - 16) are valid.

The default number of rebuild areas increases as the drive count increases.

Note: The number of rebuild areas plus the stripe width must be less than or equal to the total drive count.

Note: You cannot create a distributed RAID array of compressing drives with more than one rebuild area.

-rebuildareasgoal ***0 /1 | 2 | 3 | 4***
(Optional) Specifies the number of rebuild areas that the array can target to keep available. If the number available in the array falls below this number, a system alert is raised.

Note: The goal value should not exceed the number of rebuild areas that are specified for the array.

-strip ***256***
(Optional) Specifies the strip size in KiB for the array that is being configured. The value is 256.

-name ***new_name_arg***
(Optional) Specifies the name of the array.

-encrypt *yes | no*
(Optional) Specifies the array to encrypt. The values are yes and no. This parameter defaults to yes when **lsencryption** has its status set to enabled and all nodes in the I/O group that the array is being defined on are encryption-capable.

Note: The value can be yes only if encryption is enabled on the array's I/O group.

If you specify -encrypt yes when the I/O group does not support encryption, the command fails.

-slowwritepriority latency | redundancy
(Optional) Controls array ability to complete write operations that take too long, even if it temporarily compromises redundancy.

The value can be either latency or redundancy:

- latency implies that the feature is enabled for normal I/O operations
- redundancy implies that the feature is not enabled for normal I/O operations

The default value is latency mode for existing arrays).

Important: An array can cause member drives to become unsynchronized (to preserve response time) if the value is latency. If the value is redundancy, the array cannot cause member drives to become unsynchronized (to preserve time) and I/O performance is impacted.

mdiskgrp_id | mdiskgrp_name
(Required) Indicates the MDisk array ID or name.

Description

This command creates distributed arrays.

If you are adding an MDisk to a pool that uses a provisioning policy, the capabilities of the MDisk must match the capacity savings method that the provisioning policy defines. For example, an MDisk that uses self-compressing drive technology cannot be added to a data reduction pool that defines a provisioning policy that uses compression and deduplication.

Remember: You cannot create an unencrypted array to add to an encrypted storage pool.

Each distributed array occupies 16 slots, which start at an MDisk ID that is divisible by 16. See the **lsmdisk** command for more information.

An invocation example to create an array that uses 40 drives of class 3 with 3 rebuild areas

```
 mkdistributedarray -level raid6 -driveclass 3 -drivecount 40 -stripewidth 10 -rebuildareas 3
mdiskgrp5
```

The detailed resulting output:

```
MDisk, id [16], sucessfully created
```

Related reference

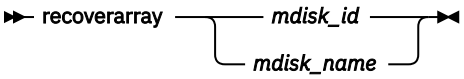
“lsmdisk” on page 450

Use the **lsmdisk** command to display a concise list or a detailed view of managed disks (MDisks) visible to the system. It can also list detailed information about a single MDisk.

recoverarray

Use the **recoverarray** command to recover a specific corrupted array when you recover from either a T3 or when both nodes in an I/O group are removed.

Syntax



Parameters

mdisk_id
(Optional) Identifies (by ID) the specific array to recover.

mdisk_name
(Optional) Identifies (by user-assigned name) the specific array to recover.

Description

This command recovers a specific corrupt array. An array has metadata representing ongoing or pending product writes, which are lost when the domain nodes are lost.

It is not necessary to run this command when both nodes go offline and return to an online state.

An invocation example

```
recoverarray mdisk_1
```

The resulting output:

```
There is no output if the command is successful.
```

recoverarraybycluster (Discontinued)

Attention: The **recoverarraybycluster** command has been discontinued. Use the **recoverarraybysystem** command instead.

recoverarraybysystem

Use the **recoverarraybysystem** command to recover corrupted arrays from either a T3 or when both nodes in an I/O group are removed.

Syntax



Parameters

None.

Description

Use the **recoverarraybysystem** command to recover corrupted arrays. An array has metadata representing ongoing or pending product writes, which are lost when the domain nodes are lost.

It is not necessary to run this command when both nodes go offline and return to an online state.

An invocation example

```
recoverarraybysystem
```

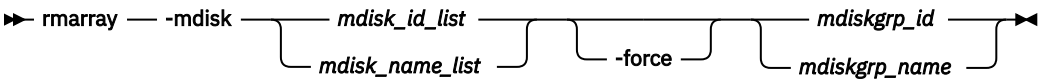
The resulting output:

```
No feedback.
```

rmarray

Use the **rmarray** command to remove an array MDisk from the configuration.

Syntax



Parameters

-mdisk *mdisk_id_list | mdisk_name_list*
(Required) Identifies the array MDisk or a colon-delimited list of MDisks to remove from the storage pool.

-force
(Optional) Forces a removal when the MDisk has allocated extents by migrating the used extents to free extents in the storage pool.

mdiskgrp_id | mdiskgrp_name
(Required) Identifies (by name or ID) the storage pool to remove the created array MDisk from.

Description

This command removes an array MDisk from the configuration. Each array is divided into candidate drives.

Remember: This command cannot be used to remove an array MDisk from a child pool.

Note: If **rmarray** is used to remove an array that contains compressed member drives, it causes all compressed drive members to perform a format. The format must complete before the drives come online as candidates available for use.

An invocation example

```
rmarray -mdisk 6 mdiskgrp_10
```

The resulting output:

```
No feedback
```

Chapter 6. Audit log commands

Use the audit log commands to track command specifications and related data. An audit log keeps track of successful action commands that are issued by users through all supported interfaces.

The audit log entries provide the following information:

- Identity of the user who issued the action command
- The name of the actionable command
- The timestamp of when the actionable command was issued on the configuration node
- The parameters which were issued with the actionable command
- The interface where the command originated: GUI, CLI, and REST

The following commands are not documented in the audit log:

- **dumpconfig**
- **cpdumps**
- **finderr**
- **dumperrlog**

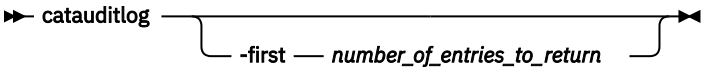
The following items are also not documented in the audit log:

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

cataudituplog

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

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-first *number_of_entries_to_return*
(Optional) Specifies the number of most recent entries to display.

Description

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

Use this command to display the in-memory audit log. Use the **dumpauditlog** command to manually dump the contents of the in-memory audit log to a file on the current configuration node and clear the contents of the in-memory audit log.

The in-memory portion of the audit log is limited to 500 entries.

© Copyright IBM Corp. 2003, 2025

Once the in-memory audit log reaches maximum capacity, the log is written to a local file on the configuration node in the `/dumps/audit` directory. The **catauditlog** command only displays the in-memory part of the audit log; the on-disk part of the audit log is in readable text format and does not require any special command to decode it.

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

The **lsdumps** command with **-prefix** parameter (and the `/dumps/audit` file) can be used to list the files on the disk.

As commands are ran,, audit log file, they are recorded in the in-memory audit log. When the in-memory audit log becomes full, it is automatically dumped to an and the in-memory audit log is cleared.

The origin field displays the command that is derived from GUI, CLI, VASA, or REST. The field only populates for the cluster commands, and will be blank for all the service commands.

The following attribute values can be displayed in the output.

*Table 17. **catauditlog** output*

Attribute	Possible Values
two_person_integrity_promoted	Indicates that the executed command shown in the audit log was executed by the user shown in the audit log while that user had an approved role elevation.
duration_ms	Indicates the time taken for CLI command to complete.

An invocation example

```
catauditlog
```

The resulting output:

```
audit_seq_no timestamp    cluster_user challenge source_panel target_panel
ssh_ip_address result res_obj_id action_cmd                                origin
two_person_integrity_promoted duration_ms
0            230727122854 superuser             01-1         01-1
0                 satask cpfiles -prefix /dumps/audit/auditlog.json -source 01-2
no                     0
1            230727123210 superuser                                        x.xx.xx.xxx
0                 svctask chenclosure -managed yes 1                            CLI
no                     71
2            230727125742 superuser                                        x.xx.xx.xxx
0      0          svctask mkmdiskgrp -ext 256                                   CLI
no                     7
3            230727133555 superuser                                        x.xx.xx.xxx
0      0          svctask mkarray -level raid0 -drive 0:1:2:3:4:5 0             CLI
no                     41
4            230727133617 superuser                                        x.xx.xx.xxx
0      0          svctask mkvdisk -mdiskgrp 0 -size 512 -unit mb                CLI
no                     40
5            230727133617 superuser                                        x.xx.xx.xxx
0      1          svctask mkvdisk -mdiskgrp 0 -size 512 -unit mb                CLI
no                     62
6            230727133617 superuser                                        x.xx.xx.xxx
0      2          svctask mkvdisk -mdiskgrp 0 -size 512 -unit mb                CLI
no                     45
```

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



Parameters

There are no parameters.

Description

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

Use this command to manually dump the contents of the in-memory audit log to a file on the current configuration node and clear the contents of the in-memory audit log. Use the **cauditlog** command to display the in-memory audit log. The dumped audit log includes all audit log entries. This includes entries that are hidden by default in the **cauditlog** view, such as configuration changes performed by the system to configure policy-based replication.

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

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

```
auditlog_firstseq_lastseq_timestamp_clusterid
```

where

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

The audit log dump file names cannot be changed.

The **dumpauditlog** command displays the information with one field per line. The **lsdumps** command displays a list of the audit log dumps that are available on the nodes in the clustered system.

A sample audit log entry:

```
Auditlog Entry 0
    Sequence Num        : 0
    Timestamp        : Fri Aug 19 01:26:04 2022
             : Epoch + 1660872364
    Cluster User        : superuser
    Challenge        :
    SSH IP Address        : x.xxx.xx.xxxx
    Result Code        : 0
    Result Obj ID        :
    Action Cmd        : svctask setlocale -locale 10
    Source_Panel        :
    Target_Panel        :
    Initiator        : user
```

```
    Initiator ClusterID    :
    Initiator Parent Seq No    :
    Parent Seq No        :
    Origin            : CLI
```

An invocation example

```
dumpauditlog
```

The resulting output:

```
No feedback
```

lsauditlogdumps (Deprecated)

The **lsauditlogdumps** command is deprecated. Use the **lsdumps** command to display a list of files in a particular dumps directory.

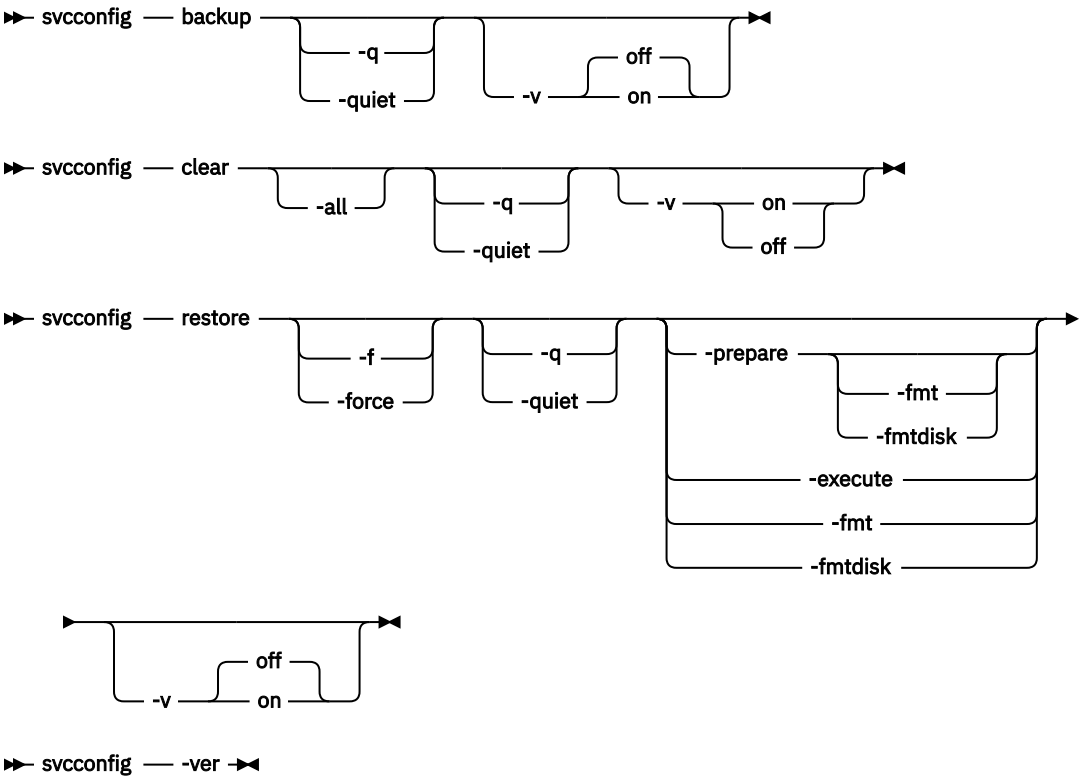
Chapter 7. Backup and restore commands

Use the backup and restore commands to back up and restore configuration information about the system.

svcconfig

Use the **svcconfig** command help option to obtain summary information about the syntax of the **svcconfig** command and actions. You can enter this command any time after a system is created.

Syntax



Parameters

backup
(Optional) Saves the current system configuration in the /tmp directory.

-quiet
(Optional) Suppresses standard output (STDOUT) messages from the console.

clear
(Optional) Erases the files in the /tmp directory.

-all
(Optional) Erases all configuration files.

-f | force
(Optional) Forces continued processing where possible.

-q | quiet
(Optional) Suppresses console output (STDOUT).

© Copyright IBM Corp. 2003, 2025

restore
(Optional) Checks the current configuration against the backup configuration in the `/tmp` directory.

-prepare ***-fmt | fmtdisk***
(Optional) Verifies the current configuration against the information in `svc.config.backup.xml`; then prepares commands for processing in `svc.config.restore.sh`, and then produces a log of events in `svc.config.restore.prepare`.

-execute
(Optional) Runs the command script `svc.config.restore.sh`, and produces a log of events in `svc.config.restore.execute.log`.

-fmt
(Optional) Specifies that the volume should be formatted before use. Includes the **-fmtdisk** option on all **mkvdisk** commands to be issued. You cannot specify **-fmt** with **-execute**.

-fmtdisk
(Optional) Specifies that the volume should be formatted before use. You cannot specify **-fmtdisk** with **-execute**.

-v on | off
Produces verbose output (on); the default is regular output (`off`).

-ver
(Required) Returns the version number for the **svcconfig** command.

Description

This command provides syntax help for `svcconfig`.

After you use the **svcconfig restore** command to restore the system configuration from backup, the two person integrity (TPI) setting must be manually re-enabled on the system.

The **svcconfig restore** command cannot restore the recovery key automatically if the only form of encryption is recovery access key. User need to configure encryption using either USB or key server before proceeding.

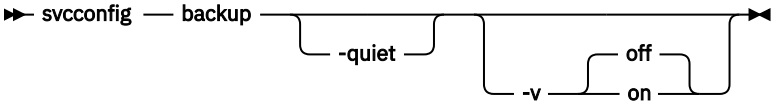
An invocation example

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

backup

Use the **backup** command to back up the configuration. Enter this command any time after creating the system.

Syntax



Parameters

-quiet
(Optional) Suppresses standard output (STDOUT) messages from the console.

-v on | off
(Optional) Displays normal (`off`, the default state) or verbose (on) command messages.

Description

The **backup** command extracts and stores configuration information from the system. The **backup** command produces the `svc.config.backup.xml`, `svc.config.backup.sh`, and `svc.config.backup.log` files, and saves them in the `/tmp` folder. The **.xml** file contains the extracted configuration information; the **.sh** file contains a script of the commands used to determine the configuration information; and the **.log** file contains details about command usage.

Note: If a previous `svc.config.backup.xml` file exists in the `/tmp` folder, it is archived as `svc.config.backup.bak`; only one archive file is stored in the `/tmp` folder.

The underscore character (_) prefix is reserved for backup and restore command usage; do not use the underscore character in any object names.

An invocation example

```
svcconfig backup
```

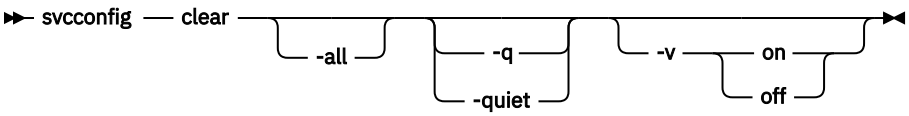
The resulting output:

```
No feedback
```

clear

Use the **clear** command to erase files in the `/tmp` directory that were previously produced by other **svcconfig** commands. You can enter this command any time after a system has been created.

Syntax



Parameters

-all
Erases all configuration files.

-q | quiet
(Optional) Suppresses console output (STDOUT).

-v on | off
(Optional) Produces verbose output (on); the default is regular output (`off`).

Description

This command erases configuration files on the current configured node.

An invocation example

```
svcconfig clear -all
```

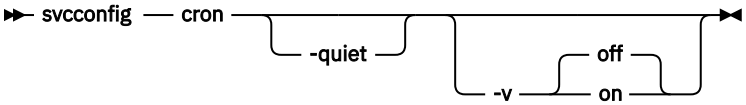
The resulting output:

```
No feedback
```

cron

Use the **cron** command to back up the configuration. Enter this command any time after creating the system.

Syntax



Parameters

-q, -quiet
Suppresses standard output (STDOUT) messages from the console.

-v on, -v off
Displays normal (`off`, the default state) or verbose (on) command messages.

Description

This command generates configuration files and places them in the configuration files directory. The file `svc.config.cron.xml_`*(node)* contains configuration detail. The file `svc.config.cron.log_`*(node)* contains a log of events. The file `svc.config.cron.sh_`*(node)* contains a script of the commands used to determine the configuration.

Any pre-existing file `svc.config.cron.xml_`*(node)* is archived as `svc.config.cron.bak_`*(node)*. Only one such archive is kept.

The configuration files directory is `/dumps`.

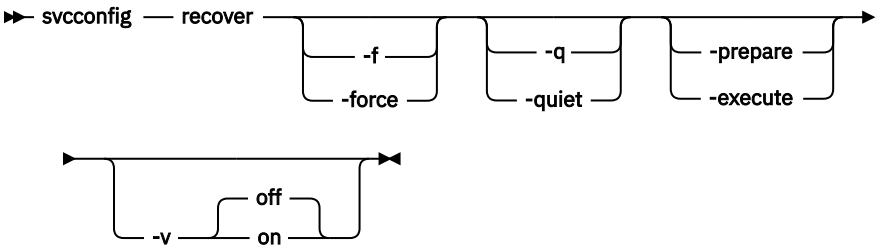
An invocation example

```
svcconfig cron
svcconfig cron -q
svcconfig cron -v on
```

recover

Use the **recover** command to recover the clustered system configuration in two phases, the preparation phase and the execution phase. This is a component of T3 Recovery.

Syntax



Parameters

-execute
(Optional) Runs the command script `svc.config.recover.sh` and produces a log of events in `svc.config.recover.execute.log`.

-f, -force
(Optional) Forces continued processing where possible.

-prepare
(Optional) Verifies the current configuration against the information in `svc.config.backup.xml` on the configuration to be recovered. Prepares commands for processing in `svc.config.recover.sh`, and produces a log of events in `svc.config.recover.prepare.log`.

-q, -quiet
(Optional) Suppresses console output (STDOUT).

-v on, -v off
(Optional) Produces verbose output (on); the default is regular output (off).

Description

The **recover** command recovers the target system configuration from the `svc.config.backup.xml` file, and associated `.key` files (if present) in the configuration files folder.

The recover operation is performed in two phases: prepare and execute. If neither the **-prepare** nor the **-execute** option is specified, the command performs both phases in sequence, producing only a single event log: `svc.config.recover.log`.

The configuration files directory is `/tmp`.

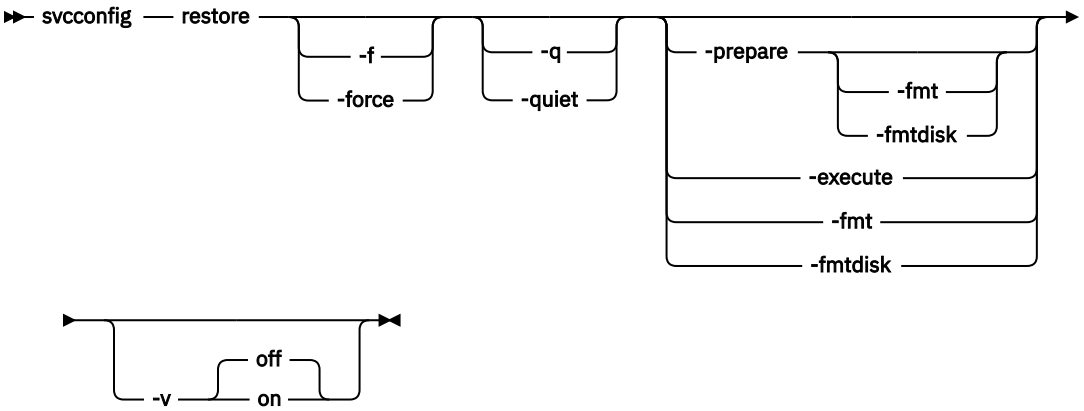
An invocation example

```
svcconfig recover -prepare
svcconfig recover -execute
```

restore

Use the **restore** command to restore the system to its previous configuration. This command uses the configuration files in the `/tmp` folder.

Syntax



Parameters

-f | force
(Optional) Forces continued processing where possible.

-q | quiet
(Optional) Suppresses console output (STDOUT).

-prepare *-fmt | fmtdisk*
(Optional) Verifies the current configuration against the information in `svc.config.backup.xml`, prepares commands for processing in `svc.config.restore.sh`, and produces a log of events in `svc.config.restore.prepare`.

-execute
(Optional) Runs the command script `svc.config.restore.sh`, and produces a log of events in `svc.config.restore.execute.log`.

-fmt
(Optional) Specifies that the volume must be formatted before use. Includes the **-fmtdisk** option on all **mkvdisk** commands to be issued. You cannot specify **-fmt** with **-execute**.

-fmtdisk
(Optional) Specifies that the volume must be formatted before use. You cannot specify **-fmtdisk** with **-execute**.

-v on | off
(Optional) Produces verbose output (on); the default is regular output (off).

Description

The **restore** command restores the target system configuration from the `svc.config.backup.xml` file in the `/tmp` folder. If neither the **-prepare** nor the **-execute** option is specified, the command performs both phases in sequence, producing only a single event log: `svc.config.restore.log`.

The restore operation is also known as a T4 (Tier 4) Recovery, and can only be used on a system having just been started. The restore operation can not be used on a system having any nonautomatic objects configured, such as storage pools or volumes.

The restore operation is performed in two phases: prepare and execute.

The command pauses for eight minutes if any nodes are added during this process, informing the user of this at run-time.

An invocation example

```
svcconfig restore
```

The resulting output:

```
No feedback
```

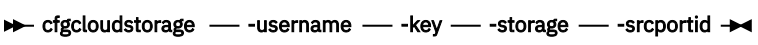
Chapter 8. Cloud commands

Use the cloud commands to create, change, or list details about cloud-related objects.

cfgcloudstorage

Use the **cfgcloudstorage** command to configure IBM Cloud® storage.

Syntax



Parameters

-username
Specifies the IBM Cloud application programming interface (API) user name.

-key
Specifies the IBM Cloud API key.

-storage
Specifies the IBM Cloud storage name.

-srcportid
Specifies the node port ID.

Description

This command configures the IBM Cloud backend storage.

An invocation example

```
cfgcloudstorage
```

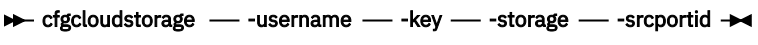
The following output is displayed:

```
No feedback
```

cfgcloudstorage

Use the **cfgcloudstorage** command to configure IBM Cloud storage.

Syntax



Parameters

-username
Specifies the IBM Cloud application programming interface (API) user name.

-key
Specifies the IBM Cloud API key.

-storage
Specifies the IBM Cloud storage name.

© Copyright IBM Corp. 2003, 2025

-srcportid
Specifies the node port ID.

Description

This command configures the IBM Cloud backend storage.

An invocation example

```
cfgcloudstorage
```

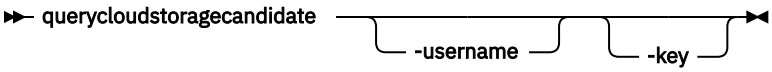
The following output is displayed:

```
No feedback
```

querycloudstoragecandidate

Use the **querycloudstoragecandidate** command to query the IBM Cloud storage candidate that is mapped to the system.

Syntax



Parameters

-username
(Optional) IBM Cloud API user name.

-key
(Optional) IBM Storage Virtualize for Public Cloud API key.

Description

This command queries the IBM Cloud storage candidate that is mapped to the clustered system.

This table describes attribute values that can be displayed as output view data.

*Table 18. **querycloudstoragecandidate** output*

Attribute	Description
storage	Indicates the candidate storage name.
datacenter	Indicates the data center to which the storage belongs.
storage_type	Indicates either Endurance or Performance type. Storage type depends on the storage type when purchased.
iops	Indicates IOPS in total
capacity_gb	Indicates the capacity of storage when purchased.

An invocation example for querycloudstoragecandidate

```
$ querycloudstoragecandidate -usr qingyuanhou -key xxxxx
```

The following output is displayed:

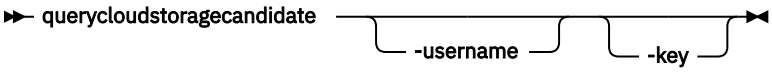
```
storage                datacenter         storage_type       iops    capacity_gb

IBM01SEL571877-10      lon02              ENDURANCE          4000    100
IBM01SEL571877-11      lon02              ENDURANCE          4000    100
IBM01SEL571877-12      lon02              PERFORMANCE        40000   40
```

querycloudstoragecandidate

Use the **querycloudstoragecandidate** command to query the IBM Cloud storage candidate that is mapped to the system.

Syntax



Parameters

-username
(Optional) IBM Cloud API user name.

-key
(Optional) IBM Storage Virtualize for Public Cloud API key.

Description

This command queries the IBM Cloud storage candidate that is mapped to the clustered system.

This table describes attribute values that can be displayed as output view data.

*Table 19. **querycloudstoragecandidate** output*

Attribute	Description
storage	Indicates the candidate storage name.
datacenter	Indicates the data center to which the storage belongs.
storage_type	Indicates either Endurance or Performance type. Storage type depends on the storage type when purchased.
iops	Indicates IOPS in total
capacity_gb	Indicates the capacity of storage when purchased.

An invocation example for querycloudstoragecandidate

```
$ querycloudstoragecandidate -usr qingyuanhou -key xxxxx
```

The following output is displayed:

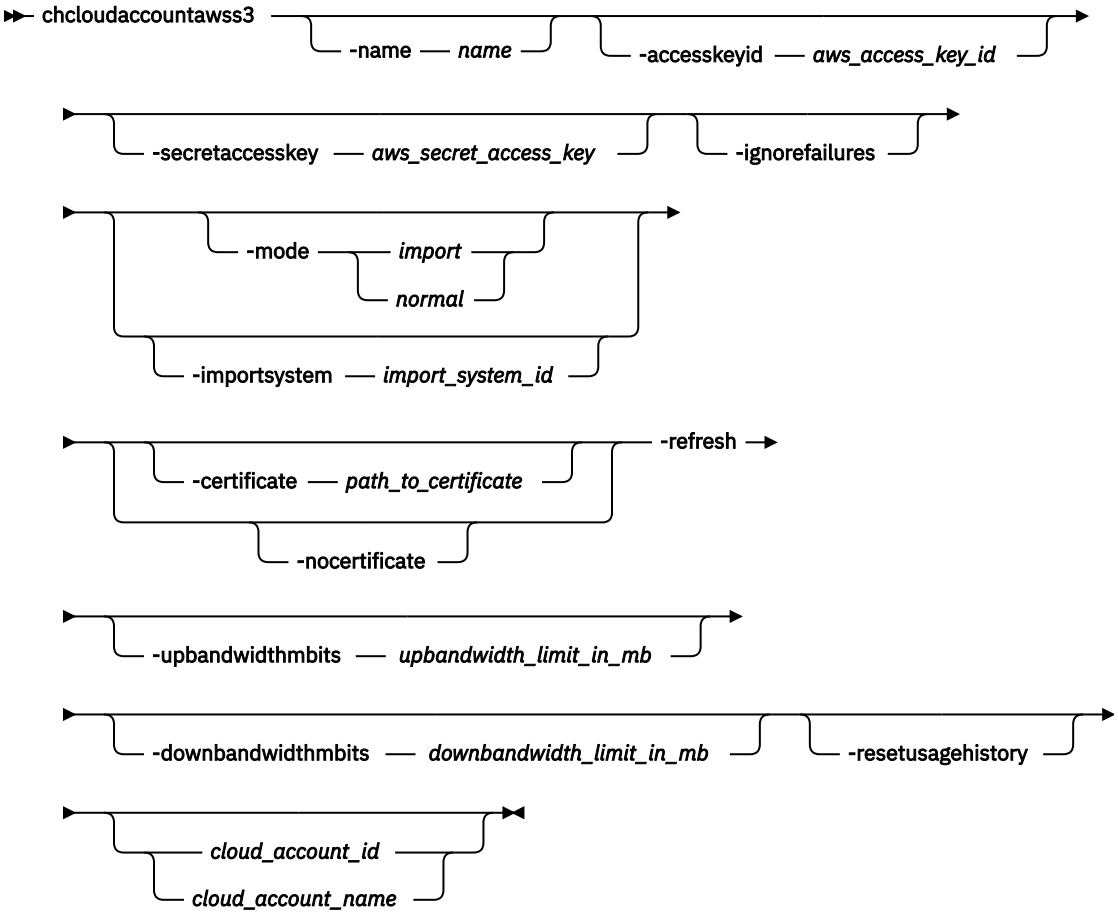
```
storage                datacenter         storage_type       iops    capacity_gb

IBM01SEL571877-10      lon02              ENDURANCE          4000    100
IBM01SEL571877-11      lon02              ENDURANCE          4000    100
IBM01SEL571877-12      lon02              PERFORMANCE        40000   40
```

chcloudaccountawss3

Use the **chcloudaccountawss3** command to modify the cloud account (that uses Amazon S3 storage) parameters or mode.

Syntax



Parameters

-name *name*
(Optional) Specifies the new or modified cloud account name. The value must be an alphanumeric value.

-accesskeyid *aws_access_key_id*
(Optional) Specifies the value for the public part of the Amazon Web Services (AWS) access key. Use this access key to access cloud storage.

-secretaccesskey *aws_secret_access_key*
(Optional) Specifies the value for the private part of the Amazon Web Services (AWS) access key. This access key is for the AWS user that the system uses to access cloud storage.

-ignorefailures
(Optional) Changes the access key whether the new access key works.

-mode *import | normal*
(Optional) Specifies the new or modified cloud account mode. The values are `import` or `normal`.

-importsystem *import_system_id*
(Optional) Specifies that the system's data be imported.

Note: You must specify `-mode` *`import`* first.

-certificate ***path_to_certificate***
(Optional) Specifies the path for the SSL certificate to use when you authenticate to the new or modified cloud account storage. The value must be an alphanumeric string of 1 - 255 characters (in base64-encoded PEM format).

-nocertificate
(Optional) Specifies that the custom SSL certificate that was used to authenticate to the new or modified cloud account storage is used to stop the system.

-refresh
(Optional) Specifies a refresh of the system import candidates. If the account is in `import` mode, this parameter specifies a refresh of the data available for import.

-downbandwidthmbits ***downbandwidth_limit_in_mb***
(Optional) Specifies the download bandwidth limit in megabits per second (Mbps). The value must be a number 1 - 10240.

Note: No download bandwidth limit is set unless you specify this keyword.

-upbandwidthmbits ***upbandwidth_limit_in_mb***
(Optional) Specifies the upload bandwidth limit in megabits per second (Mbps). The value must be a number 1 - 10240.

Note: No upload bandwidth limit is set unless you specify this keyword.

Note: Setting *0* to *upbandwidthmbits* and *downbandwidthmbits* means no limit has been configured and is for consistency with **chthrottle** CLI.

-resetusagehistory
(Optional) Resets the usage history (to `0`). Storage consumption that reflects the space that is consumed on the cloud account is cumulative, which means that it remains in the current day row (the `0th` row).

cloud_account_id | cloud_account_name
(Required) Specifies the cloud account ID or name to modify. The value for the ID must be a number and the value for the name must be an alphanumeric string.

Description

This command modifies the parameters for the cloud account (created by using **mklcloudaccountawss3**) that uses Amazon S3 storage.

The **-mode** parameter, the **-refresh** parameter, and any of the user credentials parameters groups are mutually exclusive.

This command fails and no changes are made if the supplied credentials do not provide authentication. Credentials include:

- **-accesskeyid**
- **-secretaccesskey**
- **-certificate** or **nocertificate**

For example, if the network is down then the system cannot confirm that a new secret access key is valid, the command fails. Specify **-ignorefailures** to overwrite this feature. If you specify invalid credentials but you do specify **-ignorefailures**, an online account becomes offline and an error is generated in the error log that describes the authentication failure.

The secret access key is sensitive system information and is stored in encrypted form. It is not available in the system dumps and in the audit log it is replaced with six hash ("#") symbols.

If you specify this command against an offline account and these new details enable the account to start working (for example, you enter an expired password) the account becomes online.

You must change the mode if the account is not being used by any system volumes. A mode change requires the account to be online and the system be able to communicate with the cloud server.

Note: You can have a maximum of:

- One cloud account per system.
- 1024 volumes with cloud snapshots enabled.
- 256 cloud snapshots per volume.
- 512 volume groups.

An invocation example

```
chcloudaccountawss3 -name myamazon cloudaccount0
```

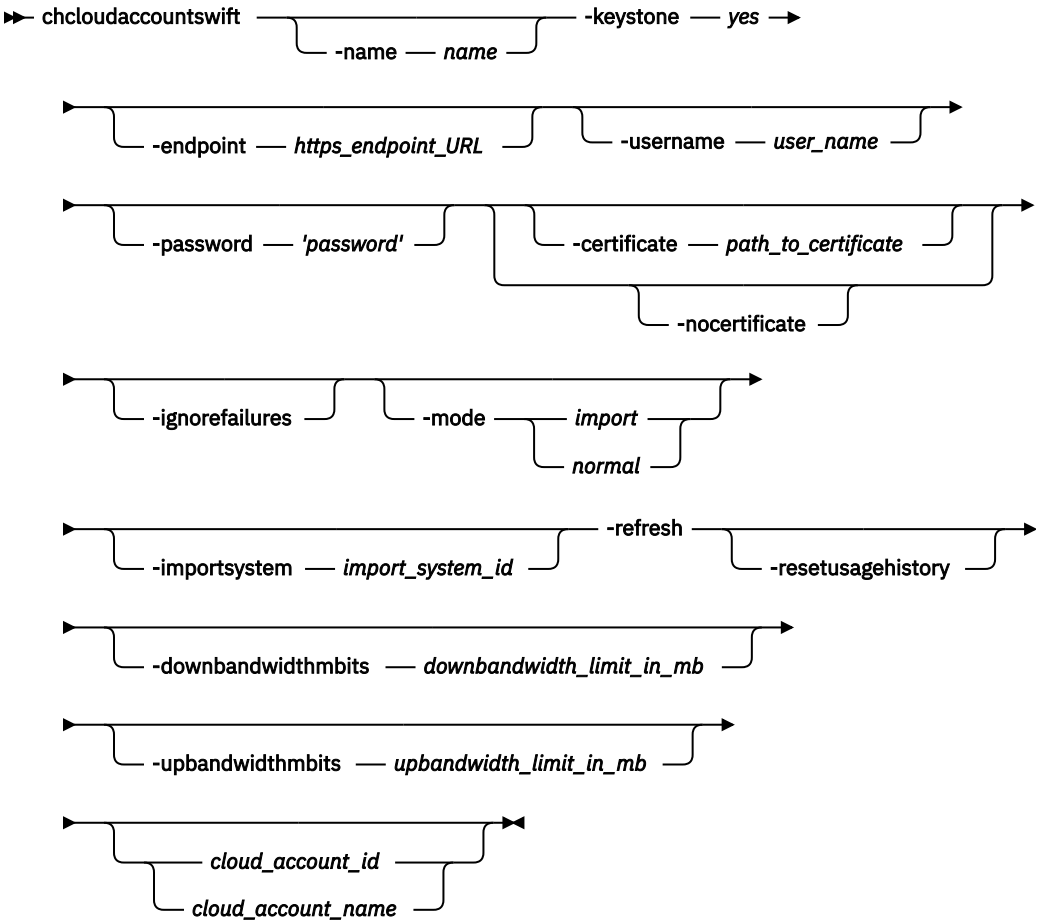
The resulting output:

```
No feedback
```

chcloudaccountswift

Use the **chcloudaccountswift** command to modify the cloud account (that uses OpenStack Swift storage) parameters or mode.

Syntax



Parameters

-name *name*
(Optional) Specifies the new or modified OpenStack name you must use to access cloud account storage. The value must be an alphanumeric value.

-keystone ***yes | no***
(Required) Specifies that keystone authentication is used. Default value is *yes*.

-endpoint ***https_endpoint_URL***
(Optional) Specifies the URL (that the system uses to access object storage) to change for the cloud account. If OpenStack Keystone authentication is used, the URL specified must be the URL for the Keystone authentication. If Keystone authentication is not used, the URL specified must be the URL for the Swift account. The value must be 8 - 128 characters and must be a valid URL address.

-username ***user_name***
(Optional) Specifies the OpenStack user name that the system must use to access cloud account storage.

-password ***'password'***
(Optional) Specifies the password value to use to authenticate to cloud storage. For IBM Cloud accounts, this password is the application programming interface (API) key. The value must be 1 - 64 alphanumeric characters and it must not begin or end with a space. You must enclose the password in single quotation marks.

-certificate ***path_to_certificate***
(Optional) Specifies the path for the SSL certificate to use when you authenticate to new or modified cloud account storage. The value must be an alphanumeric string of 1 - 255 characters (in base64-encoded PEM format).

-nocertificate
(Optional) Specifies that the custom SSL certificate that was used to authenticate to the new or modified cloud account storage is used to stop the system.

-ignorefailures
(Optional) Specifies that the access key is changed whether or not the new access key works.

-mode ***import | normal***
(Optional) Specifies the new or modified cloud account mode. The values are `import` or `normal`.

-importsystem ***import_system_id***
(Optional) Specifies that the system's data be imported.

Note: You must specify `-mode` *`import`* first.

-refresh
(Optional) Specifies a refresh of the system import candidates. If the account is in `import` mode, this parameter specifies a refresh of the data available for import.

-downbandwidthmbits ***downbandwidth_limit_in_mb***
(Optional) Specifies the download bandwidth limit in megabits per second (Mbps). The value must be a number 1 - 10240.

Note: No download bandwidth limit is set unless you specify this keyword.

-upbandwidthmbits ***upbandwidth_limit_in_mb***
(Optional) Specifies the upload bandwidth limit in megabits per second (Mbps). The value must be a number 1 - 10240.

Note: No upload bandwidth limit is set unless you specify this keyword.

Note: Setting *0* to *upbandwidthmbits* and *downbandwidthmbits* means no limit has been configured and is for consistency with **chthrottle** CLI.

-resetusagehistory
(Optional) Resets the usage history (to `0`). Storage consumption that reflects the space that is consumed on the cloud account is cumulative, which means that it remains in the current day row (the `0th` row).

cloud_account_id | cloud_account_name
(Required) Specifies the cloud account ID or name to modify. The value for the ID must be a number and the value for the name must be an alphanumeric string.

Description

This command modifies the parameters for the cloud account (created by using **mklcloudaccountswift**) that uses OpenStack Swift storage.

At least one parameter must be set.

The **-mode** parameter, the **-refresh** parameter, and any of the user credentials parameters groups are mutually exclusive. Credentials include:

- **-keystone**
- **-endpoint**
- **-username**
- **-password**
- **-certificate** or **nocertificate**

The command fails if the supplied authentication credentials are unsuccessful. For example, if the network is down the system cannot confirm that the `secretaccesskey` is valid (and the command fails). Specify **-ignorefailures** to override this feature. If you specify incorrect credentials and the **-ignorefailures** parameter, an online account becomes offline and an error is generated in the log describing the authentication failure.

The password is treated as sensitive system information. It is stored in an encrypted form and not available in system dumps. In the audit log, it is replaced with six hash ("#") symbols.

If a certificate is supplied and the command succeeds, the certificate file is deleted from the local file system.

If you specify this command against an offline account and these new details enable the account to start working (for example, you enter a new password against an expired password) the account becomes online.

You can change the mode if the account is not being used by any system volumes. A mode change requires the account to be online and the system be able to communicate with the cloud server.

Note: You can have a maximum of:

- One cloud account per system
- 1024 volumes with cloud snapshots enabled
- 256 cloud snapshots per volume
- 512 volume groups

An invocation example

```
chcloudaccountswift -certificate /tmp/new-cert.pem -ignorefailures myswift
```

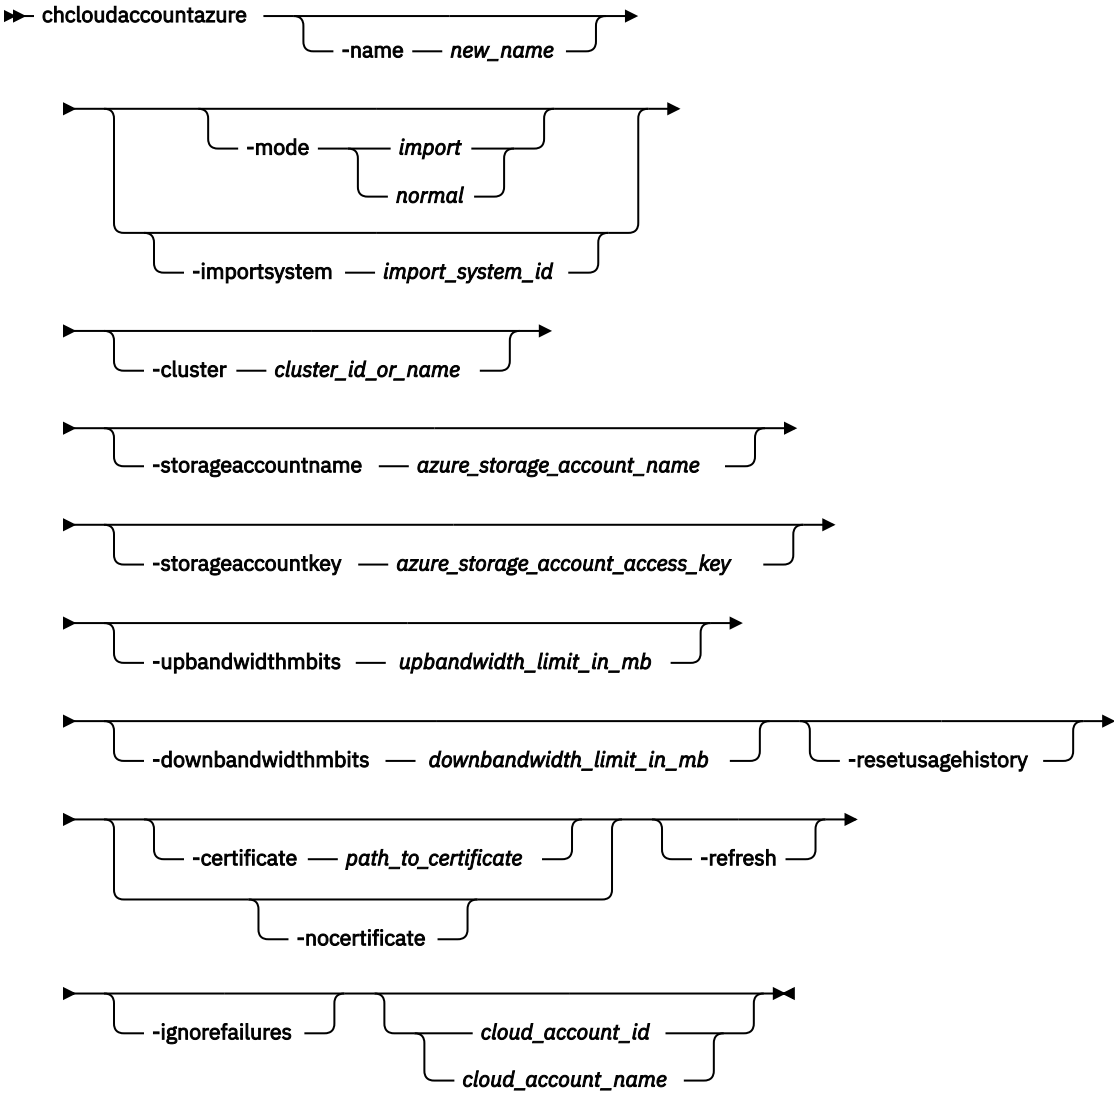
The resulting output:

```
No feedback
```

chcloudaccountazure

Use the **chcloudaccountazure** command to modify the parameters or mode of an existing cloud account that uses Microsoft Azure cloud storage.

Syntax



Parameters

-name *new_name*
(Optional) Specifies the new or modified cloud account name.

Note: Before you change the name of the storage account, make a copy of existing data that uses the Azure specific method in the target storage account.

-mode *import | normal*
(Optional) Specifies the new or modified cloud account mode. The values are `import` or `normal`.

Note:

- The **-mode** parameter can only be changed if the account is not used by any system volumes.
- To change the mode, the account must be online and the system must be able to communicate with the cloud.

-importsystem *import_system_id*
(Optional) Specifies that the system's data be imported.

Note: You must first specify `-mode import`.

-cluster *cluster_id_or_name*
(Optional) Specifies the name or ID of the Microsoft Azure cluster.

-storageaccountname *azure_storage_account_name*
(Optional) Specifies the name of the Microsoft Azure storage account that the system uses to access cloud storage. The value must be 3 - 24 characters long and can contain lowercase letters and numbers only. The storage account name is the first part (public) of the Microsoft Azure storage credential system, which provides a unique namespace in Microsoft Azure for your data objects. Every object that is stored in Microsoft Azure storage has an address that includes your unique storage account name.

-storageaccountkey *azure_storage_account_key*
(Optional) Specifies the account access key that the system uses to access the cloud storage. The value must be 88 characters long, base64-encoded, and can contain uppercase letters, lowercase letters, numbers, or slashes only. The storage account key is the second part (secret) of the Microsoft Azure storage credential system that comes with two 512-bit storage account access keys for that storage account.

-upbandwidthmbits *upbandwidth_limit_in_mb*
(Optional) Specifies the upload bandwidth limit in megabits per second (Mbps). The value must be a number 1 - 10240.

Note: No upload bandwidth limit is set unless you specify this keyword.

-downbandwidthmbits *downbandwidth_limit_in_mb*
(Optional) Specifies the download bandwidth limit in megabits per second (Mbps). The value must be a number 1 - 10240.

Note: No download bandwidth limit is set unless you specify this keyword.

-resetusagehistory
(Optional) Specifies to reset the usage history. This can reset the usage history to zero including the current day usage.

-certificate *path_to_certificate*
(Optional) Specifies the file path to an SSL CA certificate for Microsoft Azure storage. This certificate must be in base64-encoded PEM format. The value must be 1 - 255 characters long and can contain alphanumeric and special characters (. / - _) only. If periods are used, they must not be placed adjacent to each other or at the start or end of the prefix.

Note: The -certificate parameter is only required if Microsoft Azure storage updates its certificates to be incompatible with the IBM Java 8 CA certificate bundle or if you don't want to use that certificate bundle.

-nocertificate
(Optional) Specifies to stop using the custom SSL certificate to authenticate the cloud storage.

-refresh
(Optional) Specifies a refresh of the system import candidates. If the account is in `import` mode, this parameter specifies a refresh of the data available for import.

-ignorefailures
(Optional) Specify changes to the access key even if the specified access key does not work. If invalid credentials are provided with the **-ignorefailures** parameter then an online account becomes offline. An error in the log describes the authentication failure.

Description

The **mklcloudaccountazure** command modifies the parameters for the cloud account that uses Microsoft Azure cloud storage.

Note: At least one optional parameter must be specified for the command to be executed successfully.

The following parameters can be set at any time:

- **-name**
- **-upbandwidthmbits**
- **-downbandwidthmbits**
- **-resetusagehistory**

This command fails and no changes are made if the supplied credentials do not provide authentication. The Credentials group consists of following parameters:

- **-storageaccountname**
- **-storageaccountkey**
- **-ignorefailures**
- **-certificate** or **-nocertificate**

The secret access key is sensitive system information and is stored in encrypted form. It is not available in the system dumps and in the audit log it is replaced with six hash ("#") symbols. If you specify this command against an offline account and these new details enable the account to start working (for example, you enter an expired password), the account becomes online.

Setting 0 to **upbandwidthmbits** and **downbandwidthmbits** means that no limit was configured to be consistent with **chthrottle** CLI.

For example, if the network is down then the system cannot confirm that a new secret access key is valid, the command fails. Specify **-ignorefailures** to overwrite this feature. If you specify invalid credentials but you do specify **-ignorefailures**, an online account becomes offline and an error is generated in the error log that describes the authentication failure.

An invocation example

```
chcloudaccountazure -name myazure0 cloudaccount0
```

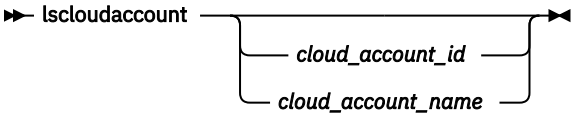
The resulting output:

```
No feedback.
```

lscloudaccount

Use the **lscloudaccount** command to display information about the configured cloud accounts.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

cloud_account_id | cloud_account_name
(Optional) Specifies the name or ID for the detailed view of the account.

Description

This command displays information about configured cloud accounts.

This table provides the attribute values that can be displayed as output view data.

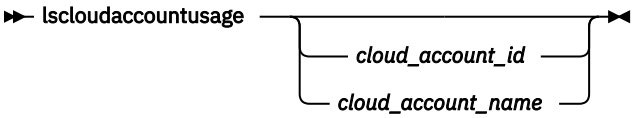
*Table 20. **lscloudaccount** output*	
Attribute	**Description**
id	Indicates the cloud account ID. The value is a number.
name	Indicates the cloud account name. The value is an alphanumeric string.
type	Indicates the cloud account provider. The values are: • awss3 • swift • azure
status	Indicates the cloud account status. The values are `online` or `offline`.
mode	Indicates the cloud account mode. The values are `normal` or `import`.
active_volume_count	Indicates the number of volumes in the system that use the account. The value must be a number.
backup_volume_count	Indicates the number of volumes that are backed up to the cloud account. The value must be a number.
import_system_id	Indicates the system ID for the system from where the data is being imported. The value must be a 16-character uppercase hexadecimal number (or blank).
import_system_name	Indicates the system name from where the data is being imported. The value must be an alphanumeric string (or blank).
error_sequence_number	Indicates an error (for offline accounts). The value must be a number (or blank).
refreshing	Indicates whether the system is refreshing its cloud storage view (for `import` mode accounts). The values are `yes` or `no`.
backup_timestamp	Indicates the timestamp for the most recent backup. The value must be in the format YYMMDDHHMMSS (or blank).
certificate	Indicates whether SSL is configured for an account that uses certificates. The values are `yes` or `no`.
certificate_expiry	Indicates the time and date that a certificate expires. The value must be blank or be in this format: `Dec 7 10:07:59 2015 GMT`
endpoint	Indicates the endpoint URL for `swift` accounts. The value must be a valid URL (or blank).
awss3_bucket_prefix	Indicates the bucket prefix that is being used for S3 accounts. The value must be a valid bucket prefix (or blank).
awss3_access_key_id	Indicates the user access key ID for S3 accounts. The value must be a valid access key ID (or blank).
awss3_region	Indicates the region that is chosen for cloud storage for S3 accounts. The value must be for a valid AWS region (or blank).
swift_keystone	Indicates whether keystone authentication is in use. The value must be `yes` or `no`.

Table 20. **lscloudaccount** output (continued)	
Attribute	**Description**
swift_container_prefix	Indicates the container prefix for Swift accounts. The value must be a valid container prefix or blank.
swift_tenant_name	Indicates the tenant name that is used for authentication for `swift` accounts. The value must be a valid tenant name (or blank).
swift_user_name	Indicates the user name that is used for authentication for `swift` accounts. The value must be a valid user name (or blank).
encrypt	Indicates the encryption status for the cloud account. The values are `yes` and no.
up_bandwidth_mbits	Specifies the upload bandwidth limit in megabits per second (Mbps). The value must be a number *0 - 10240*. If no limit is configured then *0* will be displayed.
down_bandwidth_mbits	Specifies the download bandwidth limit in megabits per second (Mbps). The value must be a number *0 - 10240*. If no limit is configured then *0* will be displayed.
azure_container_prefix	Indicates the container prefix of the Microsoft Azure storage account. The value must be 3-58 characters long and can only contain lowercase letters, numbers, and hyphens.
azure_storage_account_name	Indicates the account name of Microsoft Azure storage account.

lscloudaccountusage

Use the **lscloudaccountusage** command to list usage information about configured cloud storage accounts.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

cloud_account_id | cloud_account_name
(Optional) Specifies the cloud account to list details for. The ID value must be a number and the name value must be an alphanumeric string.

Description

This command displays usage information about configured cloud storage accounts. The information involves chargeable resource usage.

This table provides the attribute values that can be displayed as output view data.

Table 21. lscloudaccountusage output	
Attribute	**Description**
id	Indicates the ID for the cloud account. The value must be a number 0 - 4294967295.
name	Indicates the name for the cloud account. The value must be an alphanumeric string.
date	Indicates the date for the system data that is displayed. Each row shows usage of one day. The value must be in YYYYMMDD format. This value is computed relative to the current configured system date. The date in the very first entry must equate to the current date. If you manually change the system date, the changes are not reflected in the output for the date field until midnight (the time `00:00`). If you change your system date got to accommodate time zone changes, it is reflected in the output instantly. After midnight, any subsequent entry is for the next 24-hour period.
upload_data_mb	Indicates the uploaded data for one day. The value must be a number 0 - 18446744073709551615.
download_data_mb	Indicates the downloaded data for one day. The value must be a number 0 - 18446744073709551615.
storage_consumed_gb	Indicates the volume of data that is stored in this cloud account. The value must be a number 0 - 18446744073709551615.

Note: For a detailed view, there are 180 rows. Each row has information corresponding to one full day, and every field reflects activity for that day except for the `storage_consumed_gb`, which is cumulative. The most recent entry reflects the current day.

lscloudaccountimportcandidate

Use the **lscloudaccountimportcandidate** command to list information about systems that have data that is stored in the cloud accounts that are defined on this system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

Description

This command lists information about systems that store data in cloud accounts that are configured on this system. It provides information about the valid options that can be specified for import mode by various command-line interfaces based on the type of cloud account. Use the **chcloudaccountawss3** command for Amazon S3 storage, **chcloudaccountswift** for OpenStack Swift storage, and **chcloudaccountazure** for Microsoft Azure storage to set the import mode. Use the **-refresh** parameter to refresh the view by reloading what is on the specific cloud server. For example, use the **chcloudaccountawss3 -refresh** for Amazon S3 storage.

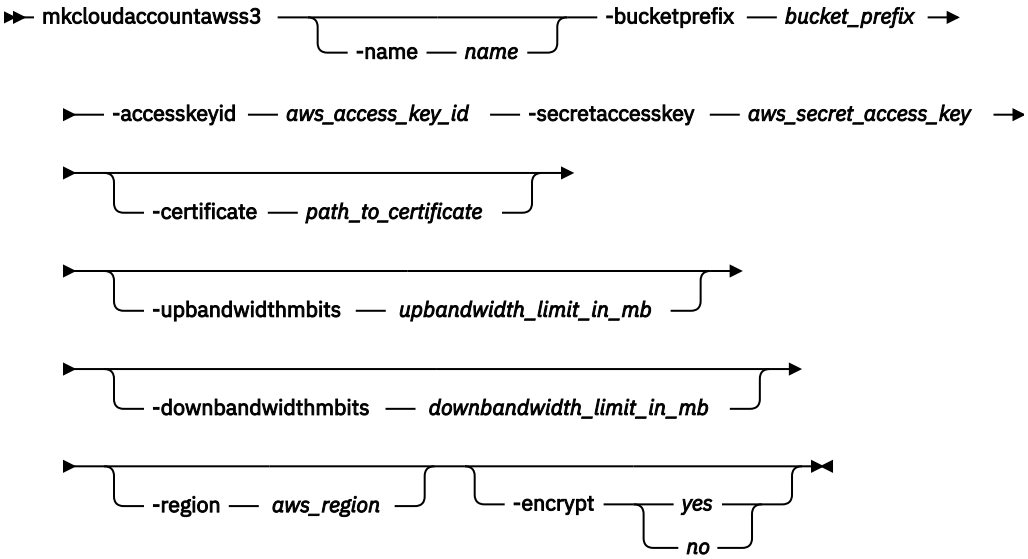
This table provides the attribute values that can be displayed as output view data.

Table 22. lscloudaccountimportcandidate output	
Attribute	**Description**
cloud_account_id	Indicates the ID for the cloud account that contains data from another system. The value must be a numeric string.
cloud_account_name	Indicates the name for the cloud account that contains data from another system. The value must be an alphanumeric string.
import_system_id	Indicates the system ID of the system that has data on the cloud account. The value must be a 16-character string in hexadecimal uppercase.
import_system_name	Indicates the system name of the system that has data on the cloud account. The value must be an alphanumeric string.
backup_volume_count	Indicates the number of volumes that are backed up by the imported system. The value must be a numeric string.
backup_size	Indicates the approximate amount of cloud storage that is in use by snapshots from the imported system.
backup_timestamp	Indicates the timestamp of the most recent volume backup (by the other system). The value must be in the YYMMDDHHMMSS format or be blank. This value is displayed in UNIX time.

mkcloudaccountawss3

Use the **mkcloudaccountawss3** command to configure a new cloud account that uses Amazon Web Services (AWS) Amazon Simple Storage Service (Amazon S3) object storage.

Syntax



Parameters

-name ***name***
(Optional) Specifies the name for the cloud account. The value must be an alphanumeric string.

-bucketprefix ***bucket_prefix***
(Required) Specifies the prefix for the Amazon S3 bucket names that the system uses. The value must be a lower-case alphabetic string 3 - 58 characters long (with no dot or period at the start or end of

the string, and no dot or period next to another dot or period). Prefix name must begin with a letter or number and can not be started with xn--.

-accesskeyid *aws_access_key_id*
(Required) Specifies the public part of the access key credential of the Amazon S3 account that the system use to access the cloud storage. The value must be a 20-character alphanumeric string of uppercase letters and numbers.

-secretaccesskey *aws_secret_access_key*
(Required) Specifies the non-public part of the access key credential of the Amazon S3 account that the system use to access the cloud storage. The value must be a 40-character alphanumeric string (That can contain slashes, or "/").

-certificate *path_to_certificate*
(Optional) Specifies the path to an SSL certificate authority (CA) certificate for Amazon S3 storage. The value must be an alphanumeric string of 1 - 255 characters (in base64-encoded PEM format).

-upbandwidthmbits *upbandwidth_limit_in_mb*
(Optional) Specifies the upload bandwidth limit in megabits per second (Mbps). The value must be a number 1 - 10240.

Note: No upload bandwidth limit is set unless you specify this keyword.

-downbandwidthmbits *downbandwidth_limit_in_mb*
(Optional) Specifies the download bandwidth limit in megabits per second (Mbps). The value must be a number 1 - 10240.

Note: No download bandwidth limit is set unless you specify this keyword.

-region *aws_region*
(Optional) Specifies the AWS region to use to access the cloud account and store data.

-encrypt *yes | no*
(Optional) Specifies whether to encrypt the data in the cloud account. By default, encryption is enabled unless you specify `-encrypt no`.

Description

This command configures a new cloud account that uses Amazon S3 object storage.

Note: If the system contains an encrypted cloud account that uses USB encryption, a USB flash drive with the system master key must be present in the configuration node before the cloud account can move to the online state. This requirement is necessary when the system is powered down, and then restarted.

An invocation example

```
mkcloudaccountawss3 -name myamazon
                    -bucketprefix svc_backups
                    -accesskeyid AKIAIOSFODNN7EXAMPLE
                    -secretaccesskey wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY
                    -upbandwidthmbits 100
                    -downbandwidthmbits 100
```

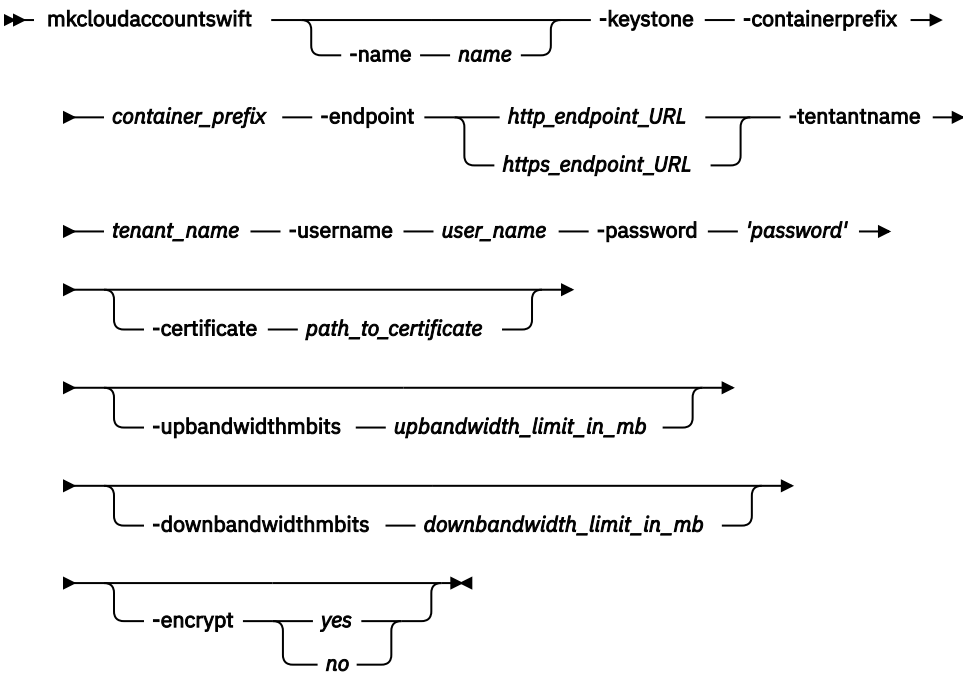
The resulting output:

```
Cloud Account, id [0], successfully created
```

mkcloudaccountswift

Use the **mkcloudaccountswift** command to configure a new cloud account that uses OpenStack Swift object storage.

Syntax



Parameters

-name *name*
(Optional) Specifies the account identifier. The value must be an alphanumeric string.

-keystone
(Mandatory) Specifies that keystone auth system is being used. TempAuth authentication is not supported.

-containerprefix *container_prefix*
(Required) Specifies the Swift container names the system uses or creates. The value must be 1 - 12 characters and contain no spaces or slashes.

-endpoint *http_endpoint_URL | https_endpoint_URL*
(Required) Specifies the URL that the system uses to access object storage.

If Keystone authentication is used, it is the URL of the Keystone service, probably ending with `v2.0`. Otherwise, it is the URL of the Swift service.

-tentantname *tenant_name*
(Required) Specifies the OpenStack tenant the system uses to access cloud storage. The value must be a 1 - 64 alphanumeric characters that contain no spaces.

-username *user_name*
(Required) Specifies the OpenStack user name the system uses to access cloud storage. The value must be 1 - 255 alphanumeric characters with no spaces.

-password *'password'*
(Required) Specifies the password that the system uses to access cloud storage. For IBM Cloud accounts, this password is the application programming interface (API) key. The value must be 1 - 64 alphanumeric characters and it must not begin or end with a space. You must enclose the password in single quotation marks.

-certificate *path_to_certificate*
(Optional) Specifies the file path for the object storage server SSL certificate. The value must be:

- 1 - 255 alphanumeric characters with no period or dot next to another period or dot, and no period or dot at the start or end of the specified value
- In base64-encoded PEM format

-upbandwidthmbits *upbandwidth_limit_in_mb*
(Optional) Specifies the upload bandwidth limit in megabits per second (Mbps). The value must be a number 1 - 10240.

Note: No upload bandwidth limit is set unless you specify this keyword.

-downbandwidthmbits *downbandwidth_limit_in_mb*
(Optional) Specifies the download bandwidth limit in megabits per second (Mbps). The value must be a number 1 - 10240.

Note: No download bandwidth limit is set unless you specify this keyword.

-encrypt *yes | no*
(Optional) Specifies whether to encrypt the data in the cloud account. By default, encryption is enabled unless you specify `-encrypt no`.

Description

This command configures a new cloud account that uses OpenStack Swift object storage.

An invocation example

```
mkcloudaccountswift -containerprefix svc_backups
                    -keystone
                    -endpoint https://lon02.objectstorage.cloud.net/auth/v1.0
                    -tenantname mytenant
                    -username jamivard
                    -password 'WKF84FAQRKLOICDF53LANBWKF84FAQRKLOICDF53LANBEXAMPLEEXAMPLEEXAMPL'
                    -upbandwidthmbits 100
                    -downbandwidthmbits 100
```

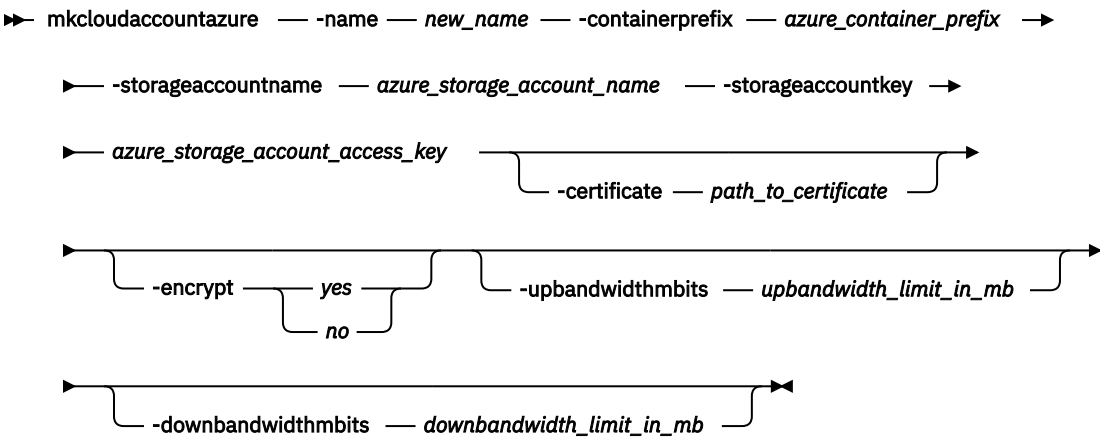
The resulting output:

```
Cloud Account, id [0], successfully created
```

mkcloudaccountazure

Use the **mkcloudaccountazure** command to configure a new cloud account that uses Microsoft Azure cloud storage.

Syntax



Parameters

-name *new_name*
(Required) Specifies the account identifier. The value must be an alphanumeric string.

-containerprefix *azure_container_prefix*
(Required) Specifies the container prefix to identify the objects that the system creates and uses on the cloud storage. The value must be 3-58 characters long and can only contain lowercase letters, numbers, and hyphens. If hypens are used, they must not be placed adjacent to each other or at the start or end of the prefix.

-storageaccountname *azure_storage_account_name*
(Required) Specifies the name of the Microsoft Azure storage account that the system uses to access cloud storage. The value must be 3-24 characters long and can only contain lowercase letters and numbers. The storage account name is the first part (public) of the Microsoft Azure storage credential system, which provides a unique namespace in Azure for your data objects. Every object that is stored in Microsoft Azure storage has an address that includes your unique storage account name.

-storageaccountkey *azure_storage_account_key*
(Required) Specifies the account access key that the system uses to access the cloud storage. The value must be 88 characters long, base64-encoded, and can only contain upper-case letters, lower-case letters, numbers, or slashes. The storage account key is the second part (secret) of the Microsoft Azure storage credential system which comes with two 512-bit storage account access keys for that storage account.

-certificate *path_to_certificate*
(Optional) Specifies the file path to an SSL CA certificate for Microsoft Azure storage. This certificate must be in base64-encoded PEM format. The value must be 1-255 characters long and can only contain alphanumeric and special characters . / - _. If periods are used, they must not be placed adjacent to each other or at the start or end of the prefix.

Note: The -certificate parameter is only required if Microsoft Azure storage updates its certificates to be incompatible with the IBM Java 8 CA certificate bundle or if you don't want to use that certificate bundle.

-encrypt *yes* | *no*
(Optional) Specifies whether to encrypt the data in the cloud account. By default, encryption is enabled unless you specify `-encrypt no`.

-upbandwidthmbits *upbandwidth_limit_in_mb*
(Optional) Specifies the upload bandwidth limit in megabits per second (Mbps). The value must be in the range of 1 to 10240.

Note: No upload bandwidth limit is set unless you specify this keyword.

-downbandwidthmbits *downbandwidth_limit_in_mb*
(Optional) Specifies the download bandwidth limit in megabits per second (Mbps). The value must be in the range of 1 to 10240.

Note:

- No download bandwidth limit is set unless you specify this keyword.

Description

The **mklcloudaccountazure** command configures a new cloud account that uses Microsoft Azure cloud storage.

Note:

- The DNS server must be configured to use the command.
- If there are already 16 other systems using the same cloud storage, then the command fails.

An invocation example

```
mkcloudaccountazure -name myazure -containerprefix sample-az
-storageaccountname sampleazureacc
-storageaccountkey  sVdi0lZQxm/BdKE9u2bVQtPkNotTgAtsVf9JMVTbBqb5zxIfRSKs9ra0pez3up/
VfX5vIFybkUZX+AStsK2n4w==
-encrypt no -upbandwidthmbits 100 -downbandwidthmbits 100
```

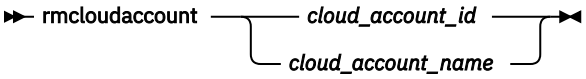
The resulting output:

```
Cloud Account, id [0], successfully created
```

rmcloudaccount

Use the **rmcloudaccount** command to delete a cloud account from the system.

Syntax



Parameters

cloud_account_id | cloud_account_name
(Required) Specifies the cloud account to remove. The ID value must be a number and the name value must be an alphanumeric string.

Description

This command deletes a cloud account from the system. If no systems have volume data that is stored in the account, the containers are deleted from the cloud storage.

As long as there are no volumes on this system that are using the cloud account, the command deletes the account. If there are no volumes that are left in the account, the system attempts to delete its containers. If you cannot connect to the cloud server, the containers are not deleted. If the command times out, the deletion proceeds asynchronously and the account object is removed.

An invocation example

```
rmcloudaccount VardyAmazAcct
```

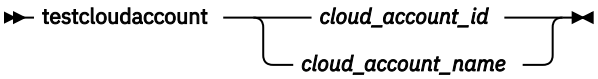
The resulting output:

```
No feedback
```

testcloudaccount

Use the **testcloudaccount** command to run diagnostics against the cloud account and report status on the results.

Syntax



Parameters

cloud_account_id | cloud_account_name
(Required) Specifies the cloud account to test. The ID value must be a number and the name value must be an alphanumeric string.

Description

This command runs diagnostics against the cloud account and reports status, which includes network connectivity, authentication, and cloud storage usage.

This command can be run against an online or offline account.

- If the command is run successfully against an offline account, the account becomes online.
- If the command is run unsuccessfully against an online account, the account becomes offline.

An invocation example

```
testcloudaccount MyVardyAccount
```

The resulting output:

```
Cloud Account, id [0], successfully tested
```

Chapter 9. Clustered system commands

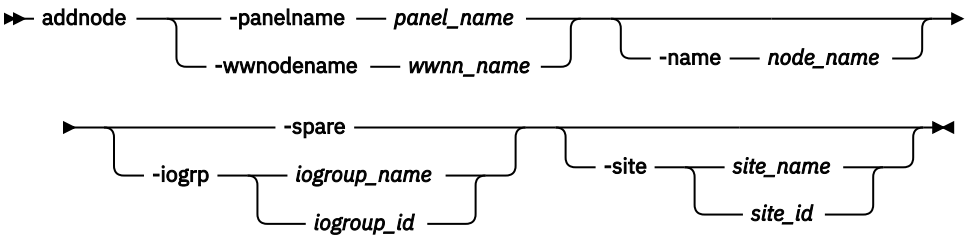
Use the system commands to monitor and modify systems and their properties.

addnode

Use the **addnode** command to add a new (candidate) node to an existing system. Enter this command any time after a system is created. If you are adding a node to a system, make sure that the model type of the new node is supported by the system code (code) version of the existing system. If the model type is not supported by the code, upgrade the system to a code version that supports the model type of the new node.

Note: The current version does not support systems with multiple I/O groups.

Syntax



Parameters

-panelname *panel_name*
(Required if you do not specify the **-wwnodename** parameter) Specifies the node that you want to add to a system by the name that is displayed in the management GUI, the service assistant, or displayed by specifying **lsnodecandidate**. You cannot use this parameter with the **-wwnodename** parameter.

Note: If `panel_name` is not supplied, it applies to the node on which the command is running.

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

-name *node_name*
(Optional) Specifies a name for the node that you want to add to the system. You can use this name in subsequent commands to refer to the node, instead of using the node ID.

Note: Node names that are supplied with the **-name** parameter on the **addnode** and **chnode** commands must not already be in use as node names or as node `failover_names`.

If you assign a name, this name is displayed as the node name from then on. If you do not assign a name, a default name is used. The default name that is used depends on whether the node is replacing one that was previously deleted. When a node is deleted, its name is retained in the I/O group as the failover name of its partner node. If no nodes remain in an I/O group, no failover names are retained. Only one failover name can be stored for each node. If you add a node into an I/O group that has a retained failover name and do not specify a node name, the retained failover name is assigned to this node. If you do not specify a name and there is no retained failover name, the name that is assigned has the format node*X*.

Important: The iSCSI Qualified Name (IQN) for each node is generated using the system and node names. If you are using the iSCSI protocol and the target name for this node is already active on its partner node, and iSCSI hosts are attached to it, adding the node with a different name changes the IQN of this node in the system and might require reconfiguration of all iSCSI-attached hosts.

© Copyright IBM Corp. 2003, 2025

-spare
(Optional) Specifies that the node that is being added is a spare node and not an I/O node group member. You cannot specify this parameter with **-iogrp**.

-iogrp *iogroup_name | iogroup_id*
(Required) Specifies the I/O group to which you want to add this node.

-site *site_name | site_id*
(Optional) Specifies the numeric site value or site name of the new node.

Description

This command adds a node to the system. You can obtain a list of candidate nodes (nodes that are not already assigned to a system) by typing `lsnodecandidate`. You cannot add a node with less memory than any potential partner nodes that are in the I/O group.

Note: The **addnode** command is supported only if the multiple I/O groups compatibility feature is enabled while adding a second I/O group.

You can create thin-provisioned volumes in a data reduction storage pool on all node types. Compressed volumes in a data reduction storage pool must be created in an I/O group with node types that support compression. Nodes that support compression can be added to an I/O group that contains compressed volumes.

If encryption is enabled on the system, then a new encryption licence for each new MTM serial must be installed by using the management GUI before the new node or enclosure can be added to the system.

Note: This command is successful only if the node-enclosure system ID matches the system, or is blank.

When the first thin or compressed volume in a data reduction pool is created for an IO group, the IO group sets CPU parameters based on the lowest number of CPU resources available based on the nodes in the I/O group. A new node with less CPU resources cannot be added to the I/O group.

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

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

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

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

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

1. Ensure that the model type of the new node is supported by the code level for the system. If the model type is not supported by the system code, you must upgrade the system to a version of code that supports the model type of the new node.
2. Record the node serial number, the WWNN, all WWPNs, and the I/O group to which the node is added. You might need to use this information later. Having it available can prevent possible data corruption if the node must be removed from and re-added to the clustered system.

Note: Transparent cloud tiering can be enabled on a system if every node on the system supports it. If a system supports transparent cloud tiering, you cannot add nodes that do not support it to the system.

Other considerations when you add a node to a system:

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

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

When a node is added to a system, it displays a state of adding. It can take 30 minutes for the node to be added to the system, particularly if the version of code that is associated with the node changed.



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

When a node is deleted, its name is retained in an I/O group as the failover name of its partner node. If no nodes remain in an I/O group, no failover names are retained.

The **addnode** command fails if you specify a name that is either an existing node name or a retained failover name, or if the system has a configuration that exceeds the limits for the node that is being added. Specify a different name for the node that is being added.

Compressed or thin deduplicated volumes can be added only to systems in which all nodes support deduplicated volumes. You can only add nodes that support deduplicated volumes to a system that contains compressed or thin deduplicated volumes. Nodes can only be added to a system that contains compressed or thin deduplicated volumes if that new node can support the amount of memory that is allocated for data deduplication in the target I/O group.

Note: A system with a storage partition is only allowed to have a single I/O group.

An invocation example

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

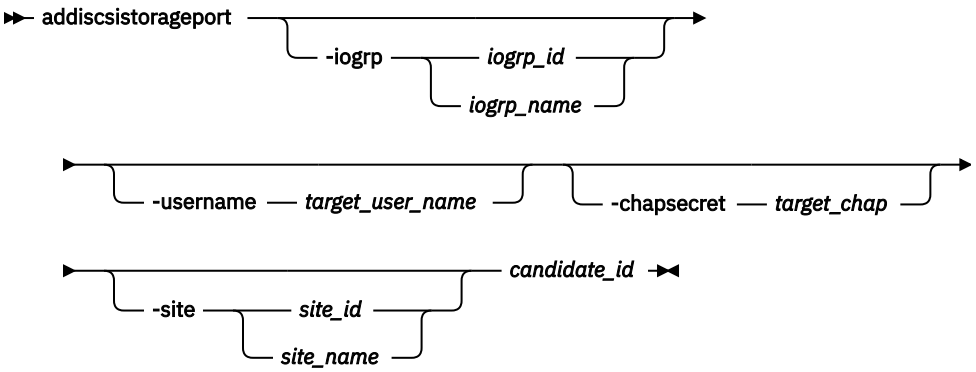
The resulting output:

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

addiscsistorageport

Use the **addiscsistorageport** command to establish Internet Small Computer Systems Interface (iSCSI) login sessions from any (or all) nodes in a specified I/O group (or the entire clustered system) to a discovered backend target iSCSI controller.

Syntax



Parameters

-iogrp *iogrp_id | iogrp_name*
(Optional) Specifies I/O group ID or name that is added. The *iogrp_id* value must be 0, 1, 2, or 3. The *iogrp_name* value must be an alphanumeric string.

Specifying this parameter triggers discovery through both nodes for the specified I/O group. The port number on each node (used to establish a session) is displayed in the selected row of discovery results from specifying **detectiscsistorageportcandidate**.

-username *target_user_name*
(Optional) Specifies the target controller user name that is added. The value must be an alphanumeric string up to 256 characters.

If the target controller requires a *target_user_name* and *target_chap* for discovery, the user name for the target controller must be specified.

Some controllers might require that you use the iSCSI qualified name (IQN) user name for discovery. Each nodes IQN is picked up automatically and used if required.

-chapsecret *target_chap*
(Optional) Specifies the Challenge-Handshake Authentication Protocol (CHAP) secret, *target_chap*, that is required for discovery of the target iSCSI controller that is being added. The value must be an alphanumeric string (case-sensitive) up to 79 characters. This keyword is required if **-username** is specified.

-site *site_id | site_name*
(Optional) Specifies the site ID or site name of the host that is being detected. The site ID must be 1 (the default) or 2. The site name must be an alphanumeric value.

candidate_id
(Required) Indicates the row ID indicating the selected row in the **lsiscsistorageportcandidate** output.

Description

This command establishes iSCSI login sessions from a specified I/O group (or, if I/O group is not specified, the entire clustered system) to a discovered backend iSCSI target controller.

To use this command, you must first:

1. Specify **detectiscsistorageportcandidate** to detect or discover the backend controller target ports.
2. Specify **lsiscsistorageportcandidate** to display the discovery output and locate a unique IQN and Internet Protocol (IP) combination on the discovered iSCSI storage controller.
3. Specify **addiscsistorageport**. You can establish sessions (from all nodes in a single I/O group) to the iSCSI controller port by specifying the I/O group number of the source port that the iSCSI sessions are initiated from.

If you do not specify an I/O group, sessions are established from all nodes in the system. The source port identifier is in the discovery results. Because some iSCSI controllers represent logical unit numbers (LUNs) as IQNs and might require a different *target_user_name* and *target_chap* for every IQN, these values can also specified for authentication when establishing the session.

Note: You can use the **chiscsistorageport** command to add more initiator node authentication credentials.

A detailed invocation example

This example shows target discovery that uses an IPv4 IP address for a target iSCSI controller through source port ID 0. For example, first you might specify:

```
detectiscsistorageportcandidate -targetip 192.168.81.91  -srcportid 3 -chapsecret Vardy -site 1
```

You would then specify **lsiscsistorageportcandidate** to list iSCSI port information:

```
id src_port_id target_ipv4    target_ipv6 target_iscsiname   iogroup_list configured status
0  4           192.168.213.33             IQN1                1:1:1:1      yes        full
```

Then establish a session using **addiscsistorageport** for discovery output row 0.

Note: If you specify **-username** or **-chapsecret** with **detectiscsistorageportcandidate** during discovery, you must specify **-username** or**-chapsecret** for session establishment.

```
addiscsistorageport 0
```

The detailed resulting output:

```
No feedback
```

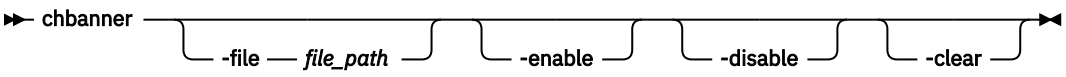
cfgportip (Deprecated)

The **cfgportip** command has been deprecated. Instead of using **cfgportip** command use **mkip** command to create a new IP address and associate it with a portset. You can use **chportethernet** command to configure the permitted uses of ethernet ports and the maximum transmission unit (MTU) setting for the ethernet ports.

chbanner

Use the **chbanner** command to configure the login message that is displayed during CLI Secure Shell (SSH) login.

Syntax



Parameters

-file *file_path*
(Optional) Specifies the path to the file on the configuration node that contains the new login message.

-enable
(Optional) Enables the login message.

-disable
(Optional) Disables the login message.

-clear
(Optional) Clears the login message.

Description

This command configures the login message that is displayed during CLI SSH login. Use this command for warnings or disclaimers or anything else you need to display in your login screen before logging in.

The file that contains the login message must be copied to the configuration node before you specify `chbanner -file`. If a configuration node failover occurs between copying the file that contains the login message and running the command, the temporary file must be copied to the new configuration node.

To set a login message that uses a SAN administrator's workstation:

1. Use a suitable text editor to create the message and save the file with a recognizable name.
2. Use a secure copy client to copy the file to the configuration node of the system to be configured.
3. Specify the management IP address of the system that is to be configured.
4. Log into the system to be configured.
5. Use the **chbanner** command to set the login message.

An invocation example

```
chbanner -file /tmp/loginmessage
```

The detailed resulting output:

```
No feedback
```

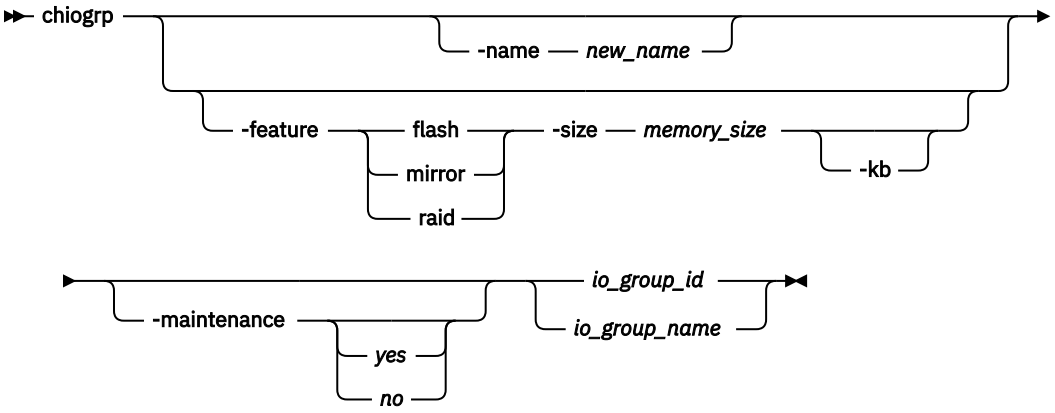
chcluster (Discontinued)

Attention: The **chcluster** command has been discontinued. Use the **chsystem** command instead.

chiogrp

Use the **chiogrp** command to modify I/O group properties or enable maintenance mode.

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 | mirror | raid
(Optional) Specifies the feature to modify the amount of memory for RAID arraysor volume mirroring. You must specify this parameter with the `-size` parameter. You cannot specify this parameter with the **-name** parameter.

- `flash` specifies the amount of memory that is used for FlashCopy®.
- `mirror` specifies the amount of memory that is used for volume mirroring operations.
- `raid` specifies the amount of memory that is used for RAID arrays.

-size *memory_size*
(Optional) Specifies the amount of memory that is available for the specified RAID arrays, Copy Services, or volume mirroring function. Valid input is `0` or any integer. The default unit of measurement for this parameter is MiB. Use the **-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 MiB to KiB. If you specify this parameter, the **-size** *memory_size* value must be any number divisible by 4. You must specify this parameter with the **-feature** and **-size** parameters. You cannot specify this parameter with the **-name** parameter.

-maintenance *yes* **/** *no*
(Optional) Enable maintenance mode for the I/O group. This is used for some service procedures. After you enter maintenance mode, it continues until either:

- It is explicitly cleared.
- Thirty minutes elapse.

io_group_id | io_group_name
(Required) Specifies the I/O group to modify.

Description

The **chiogrp** command modifies the I/O group properties, or enables maintenance mode.

Use the **-feature** and **-size** parameters (together) to change the amount of memory available in the I/O group to one of the following types:

- FlashCopy
- Volume mirroring
- RAID

For example:

```
chiogrp -feature flash -size 40 0
```

The amount of memory can be decreased or increased. Consider the following memory sizes when you use this command:

- The default amount of memory for FlashCopy is 20 MiB.
- The default memory size for mirrored volumes is 20 MiB.
- The default memory size for RAID arrays is 40 MiB.
-
-

This table demonstrates the amount of memory that is required for copy services and volume mirroring. For the listed grain size, each 1 MiB of memory provides the indicated volume capacities.

Table 23. Memory that is required for copy services and volume mirroring

Feature	**Grain size**	**1 MiB of memory provides the following volume capacity for the specified I/O group**
	256 KiB	2 TiB of total target volume capacity
	64 KiB	512 GiB of total target volume capacity
Incremental	256 KiB	1 TiB of total Incremental target volume capacity
Incremental	64 KiB	256 GiB of total Incremental target volume capacity
Volume mirroring	256 KiB	2 TiB of mirrored volumes

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

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

Note: For FlashCore Modules, the **effective capacity** should be used and not the **physical capacity**.

An invocation example to change the name of an existing I/O group to `testiogrpone`

```
chiogrp -name testiogrpone io_grp0
```

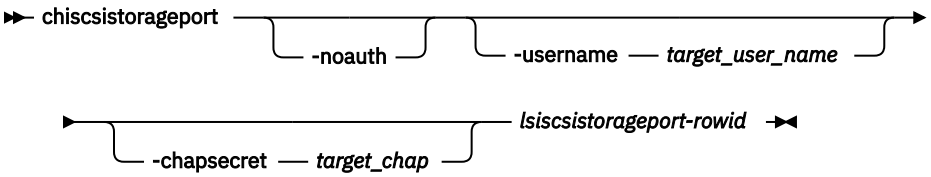
The resulting output:

```
No feedback
```

chiscsistorageport

Use the **chiscsistorageport** command to change authentication parameters, such as setting authentication credentials, removing authentication parameters, or updating credentials.

Syntax



Parameters

-noauth
(Optional) Clears all authentication parameters for a session.

Note: The **-noauth** parameter cannot be used with other parameters.

-username *target_user_name*
(Optional) Specifies the target controller user name. The value must be an alphanumeric string up to 256 characters.

If the target controller requires a *target_user_name* and *target_chap* for discovery, the user name for the target controller must be specified.

Warning: Changing the **-username** for a target controller can be a disruptive operation, so exercise caution when you change the authentication details for the session. Be sure to make the controller side authentication credentials changes before you change the authentication credentials for the session.

Some controllers might require that you use the iSCSI qualified name (IQN) user name for discovery. The IQN of each node is picked up automatically and used if required.

-chapsecret *target_chap*
(Optional) Specifies the Challenge-Handshake Authentication Protocol (CHAP) secret *target_chap* required for discovery of the target iSCSI controller. The value must be an alphanumeric string up to 80 characters. This keyword is required when not using the **-noauth** parameter.

Warning: Changing the **-chapsecret** for a target controller can be a disruptive operation, so exercise caution when you change the authentication details for the session. Be sure to make the controller side authentication credentials changes before you change the authentication credentials for the session.

lsiscsistorageport-rowid
(Required) Specifies the row ID of an existing **lsiscsistorageport** output row.

Description

The **chiscsistorageport** command operates on a row ID specified by the output of **lsiscsistorageport** command. Because some storage controllers do not drop the existing active

session after you change the authentication credentials, this command forces the session to drop and reconnect to confirm that the changed authentication credentials work.

Note: You cannot change the session mode between target-specific authentication and initiator node-specific authentication by using this command.

An invocation example to clear authentication for an iSCSI session

This example shows how to clear the authentication details of a session. For example, first specify **lsiscsistorageport** to list iSCSI port information:

```
lsiscsistorageport
id   port_id target_ipv4 target_ipv6 target_iscsiname controller_id controller_name
iogroup_list status site_id site_name
5     2      10.10.10.1              IQN1             1             ctlr1
1:1:1:1      full
```

The following example shows how to clear authentication where the target has a single user name or CHAP secret.

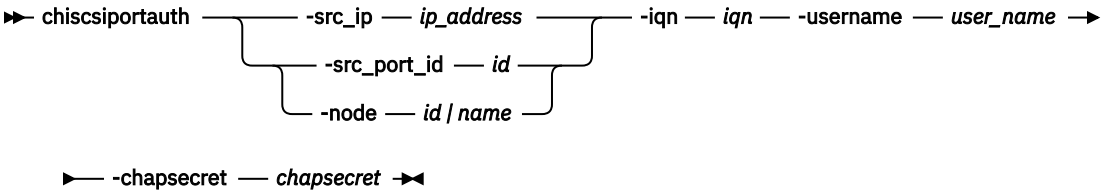
```
chiscsistorageport -noauth 5
```

The result is clearing the authentication of iSCSI sessions from all initiator nodes.

chiscsiportauth

Use the **chiscsiportauth** command to set or configure the iSCSI Initiator authentication or authorization information that is used to connect to the backend IBM Cloud storage.

Syntax



Parameters

-src_ip *ip_address*
(Required if you do not specify **-src_port_id** and **-node**) The IPv4 or IPv6 address of the system initiator port for which the credentials are configured. You cannot specify this value with **-src_port_id** or **-node**.

-src_port_id *id*
(Required if you do not specify **-src_ip**) The source port identifier of the initiator port for which the credentials are configured. This value is the numeric ID of the Ethernet port starting with 1 up to the maximum ports in the system. You cannot specify this value with **-src_ip**, and you must specify this value with **-node**.

-node *id | name*
(Required if you do not specify **-src_ip**) Specify the node ID or name for which the credentials are configured. You cannot specify this value with **-src_ip**, and you must specify this value with **-src_port_id**.

-iqn *iqn*
(Required) The iSCSI Qualified name (IQN) used to connect to the target storage. This value is an alphanumeric with a maximum of 256 characters.

-username *user_name*
(Required) The username used to authenticate to the target storage. This value is an alphanumeric with a maximum of 32 characters.

-chapsecret *chap_secret*
(Required) The chapsecret used to authenticate to the target storage. This value is an alphanumeric with a maximum of 32 characters.

Description

This command is used to set or configure the iSCSI Initiator authentication or authorization information that is used to connect to the backend IBM Cloud storage. Before you run this command, the system authentication must be set to *ip* with the `svctask chsystem -force ip` command.

An invocation example

```
svctask chiscsiportauth -src_ip 192.168.4.21 -iqn iqn.1986.ibm.com:cluster1.node1 -username
marvel -chapsecret phantom
```

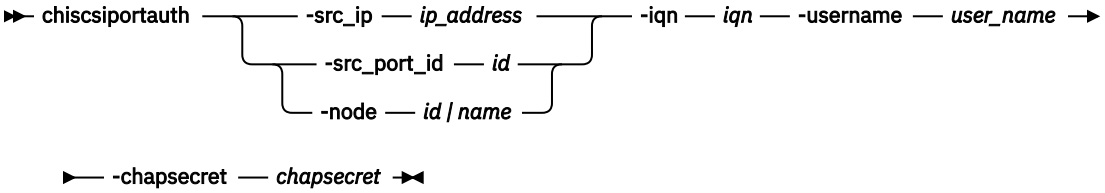
The resulting output:

```
No feedback
```

chiscsiportauth

Use the **chiscsiportauth** command to set or configure the iSCSI Initiator authentication or authorization information that is used to connect to the backend IBM Cloud storage.

Syntax



Parameters

-src_ip *ip_address*
(Required if you do not specify **-src_port_id** and **-node**) The IPv4 or IPv6 address of the system initiator port for which the credentials are configured. You cannot specify this value with **-src_port_id** or **-node**.

-src_port_id *id*
(Required if you do not specify **-src_ip**) The source port identifier of the initiator port for which the credentials are configured. This value is the numeric ID of the Ethernet port starting with 1 up to the maximum ports in the system. You cannot specify this value with **-src_ip**, and you must specify this value with **-node**.

-node *id | name*
(Required if you do not specify **-src_ip**) Specify the node ID or name for which the credentials are configured. You cannot specify this value with **-src_ip**, and you must specify this value with **-src_port_id**.

-iqn *iqn*
(Required) The iSCSI Qualified name (IQN) used to connect to the target storage. This value is an alphanumeric with a maximum of 256 characters.

-username *user_name*
(Required) The username used to authenticate to the target storage. This value is an alphanumeric with a maximum of 32 characters.

-chapsecret *chap_secret*
(Required) The chapsecret used to authenticate to the target storage. This value is an alphanumeric with a maximum of 32 characters.

Description

This command is used to set or configure the iSCSI Initiator authentication or authorization information that is used to connect to the backend IBM Cloud storage. Before you run this command, the system authentication must be set to *ip* with the `svctask chsystem -force ip` command.

An invocation example

```
svctask chiscsiportauth -src_ip 192.168.4.21 -iqn iqn.1986.ibm.com:cluster1.node1 -username
marvel -chapsecret phantom
```

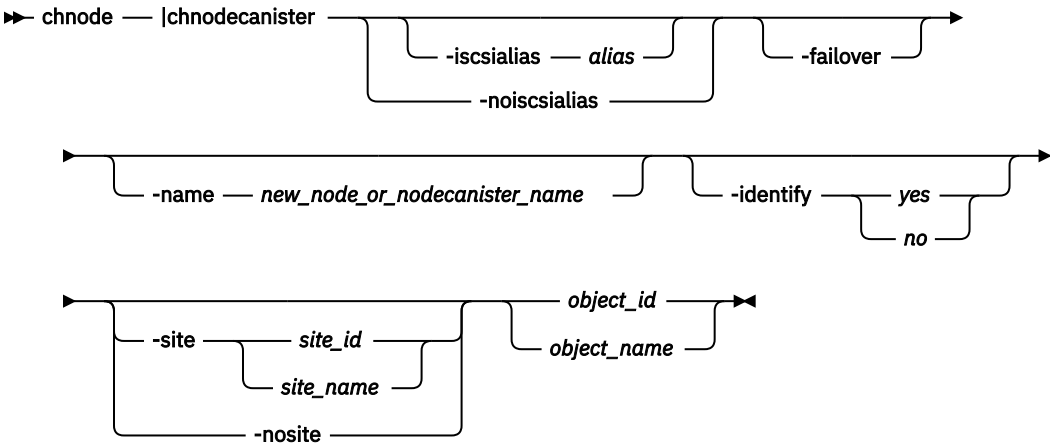
The resulting output:

```
No feedback
```

chnode (SVC) / chnodecanister (Storwize family products)

Use the **chnode** / **chnodecanister** command to change the name that is assigned to a node or node canister as well as other options. You can then use the new name when running subsequent commands. All parameters that are associated with this command are optional. However, you must specify one or more parameters.

Syntax



Parameters

-iscsialias *alias*
(Optional) Specifies the iSCSI name of the node or node canister. The maximum length is 79 characters. Do not use spaces for the iSCSI alias name.

Important: You can specify this parameter for online spares nodes.

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

Important: You can specify this parameter for online spares nodes.

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

Important: You can specify this parameter for online spares nodes.

-name ***new_node_or_nodecanister_name***
(Optional) Specifies the name to assign to the node or node canister.

Note: Node or node canister names supplied with -name on **chnode** / **chnodecanister** commands must not be in use already as node or node canister names or as node or node canister failover names.

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

-identify ***yes | no***
(Optional) Allows you to control the light-emitting diode (LED) used on the node. The values are *yes* or *no*.

Important: You can specify this parameter for online spares nodes.

-site ***site_id | site_name***
(Optional) Specifies the numeric site value or site name for the existing node. The value is 1 or 2.

-nosite
(Optional) Resets the site value.

object_id | object_name
(Required) Specifies the object name or ID that you want to modify. The variable that follows the parameter is either:

- The object name that you assigned when you added the node to the clustered system
- The object ID that is assigned to the node (not the worldwide node name)

Description

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

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

To change the name of the node (with I/O running) :

1. Make sure the host system has active sessions with both node canisters in the I/O group (hosting the volume on which the I/O occurs).
2. Change name of one node canister using **chnode** command.
3. From the host system, log out of the node canister whose name changes.

4. Rediscover the target iSCSI qualified name (IQN) from the host using the host operating system's discovery mechanism.
5. Login with the new target IQN discovered on the host system, and make sure the login succeeds.
6. Repeat steps 2-5 with the other node canister.

Note: When using VMware ESX, delete the static paths (in the iSCSI initiator properties) that contain the old target IQN.

This ensures that the node canister name change does not impact iSCSI I/O during events such as a target failover.

An invocation example

```
chnode -name newname -identify yes node8
```

The resulting output:

```
No feedback
```

An invocation example

```
chnodecanister -name testnodeone nodeone
```

The resulting output:

```
No feedback
```

An invocation example

```
chnodecanister -site 1 node2
```

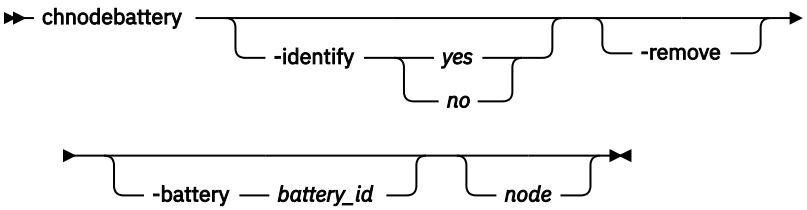
The resulting output:

```
No feedback
```

chnodebattery

Use the **chnodebattery** command to set or clear the light-emitting diode (LED) on a hot-swappable battery (in a node).

Syntax



Parameters

-identify
(Optional) Allows you to control the light-emitting diode (LED).

-battery *battery_id*
(Optional) Specifies the battery that is in the node.

-remove
(Optional) Specifies battery removal and terminates any calibration that runs on another battery.

node
(Optional) Specifies the node that the battery is in.

Description
This command notifies the battery back-up (BBU) driver that a user wants to remove a battery.

An invocation example to make the fault LED flash on battery 1 in node 3

```
chnodebattery -identify yes -battery 1 3
```

The resulting output:

```
No feedback
```

An invocation example to remove battery 1 in node 3

```
chnodebattery -remove -battery 1 3
```

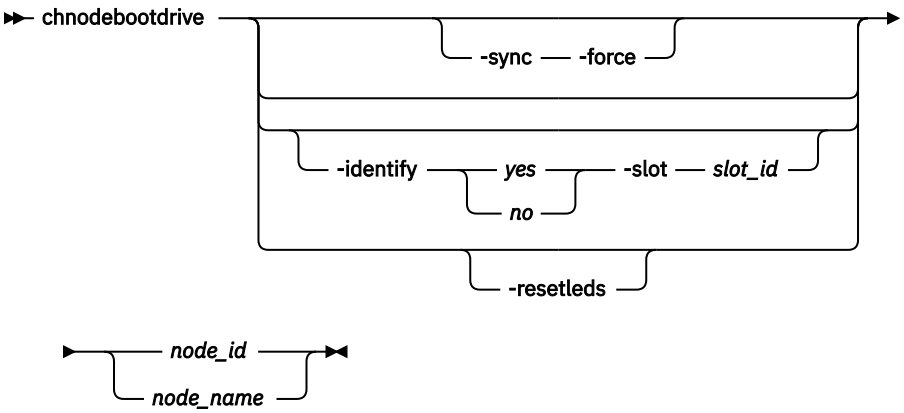
The resulting output:

```
No feedback
```

chnodebootdrive

Use the **chnodebootdrive** command to change a drive or synchronize the drives on a system if a drive or field-replaceable unit (FRU) replacement drive breaks.

Syntax



Parameters

-sync
(Optional) Specifies synchronization of drives marked `can_sync`.

-force
(Optional) Forces synchronization (though taking the node offline might cause a volume to go offline).

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

-identify *yes* | *no*
(Optional) Controls the operation of the light-emitting diode (LED) of the drive in the specified slot.

-slot *slot_id*
(Optional) Specifies the boot drive slot. It must be used with the **-identify** parameter.

resetleds
(Optional) Clears the identify LEDs of all drives in the specified node and indicates `-identify no` is specified.

node_id | node_name
(Optional) Specifies the ID or name of the node.

Description

The command identifies and synchronizes drive information for system drives.

Specifying **-sync** causes a node restart on the specified node. This restart is not successful if any volume depends on that node.

Important: If **-force** is also specified, the system does not check for dependent volumes.

An invocation example

```
chnodebootdrive
```

The following output is displayed:

```
No feedback
```

An invocation example

```
chnodebootdrive -identify yes -slot 1 1
```

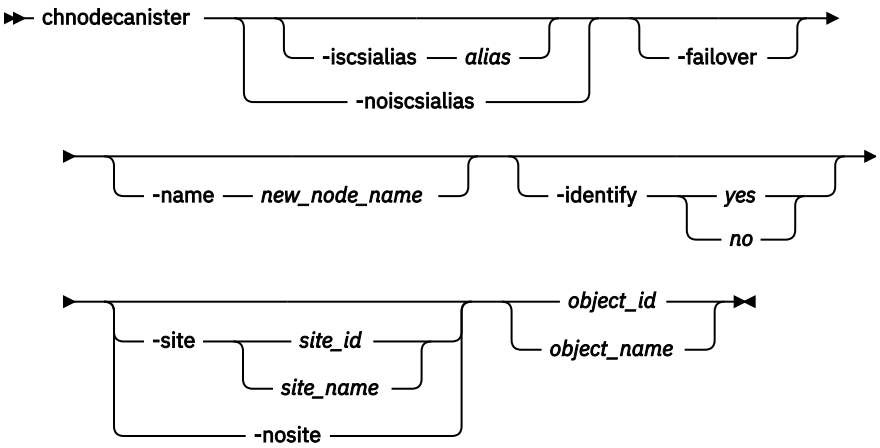
The following output is displayed:

```
No feedback
```

chnodecanister

You can use the **chnodecanister** 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



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

-name *new_nodecanister_name*
(Optional) Specifies the name to assign to the node canister.

Note: Node canister names supplied with **-name** on **chnodecanister** commands must not be in use already as node canister names or as node canister failover names.

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

-identify *yes | no*
(Optional) Allows you to control the light-emitting diode (LED) used on the node. The values are *yes* or *no*.

-site *site_id | site_name*
(Optional) Specifies the numeric site value or site name for the existing node. The value for *site_id* is 1 or 2. The value for *site_name* is an alphanumeric string.

-nosite
(Optional) Resets the site value.

object_id | object_name
(Required) Specifies the object name or ID that you want to modify. The variable that follows the parameter is either:

- The object name that you assigned when you added the node canister to the system
- The object ID that is assigned to the node canister (not the worldwide node name)

Description

If the **failover** parameter is not specified, this command changes the name of the node canister. You can then use the name to identify the node canister in subsequent commands.

If the **failover** parameter is not specified, this command changes the name or iSCSI alias of the node canister. You can then use the name to identify the node canister 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 canister 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 canister they are set on handles iSCSI I/O requests to the iSCSI alias specified, or the IQN that is created using the node canister name. If the partner node canister in the I/O group is online when these parameters are set, the partner node canister handles iSCSI requests to the iSCSI alias specified, and its node canister name and IQN change.

To change the name of the node (with I/O running) :

1. Make sure the host system has active sessions with both node canisters in the I/O group (hosting the volume on which the I/O occurs).
2. Change name of one node canister using **chnode** command.
3. From the host system, log out of the node canister whose name changes.

4. Rediscover the target iSCSI qualified name (IQN) from the host using the host operating system's discovery mechanism.
5. Login with the new target IQN discovered on the host system, and make sure the login succeeds.
6. Repeat steps 2-5 with the other node canister.

Note: When using VMware ESX, delete the static paths (in the iSCSI initiator properties) that contain the old target IQN.

This ensures that the node canister name change does not impact iSCSI I/O during events such as a target failover.

An invocation example

```
chnode -name newname -identify yes node8
```

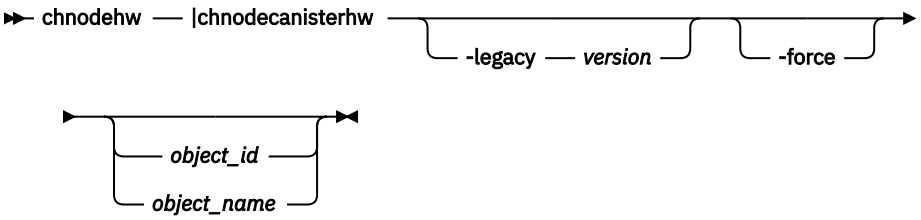
The resulting output:

```
No feedback
```

chnodehw

Use the **chnodehw** / **chnodecanisterhw** command to update the hardware configuration for a node or node canister.

Syntax



Parameters

-legacy *version*
(Optional) Sets the hardware configuration to make it compatible with the 6.3.0.0 code level. The format is four decimal numbers that are separated by periods, and there can be up to 16 characters.

-force
(Optional) Allow the node to restart and change its hardware configuration even if it causes volumes to go offline.

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

object_id | object_name
(Optional) Specifies the object name or ID.

Description

This command automatically reboots the node or node canister if the node or node canister hardware is different than its configured hardware. After rebooting, the node or node canister uses its hardware, and does not use the previous configuration.

Attention: When you run the **chnodehw** command to change the configured hardware for a node:

- Small Computer System Interface-3 (SCSI-3) reservations (through that node) are removed.
- Small Computer System Interface-3 (SCSI-3) registrations (through that node) are removed.

Note:

- This command fails if you remove the last compression card of a node and try to commit that change while compressed volumes still exist in that I/O group.
- This command cannot be used for memory upgrades. If any new hardware configuration contains extra memory, this command fails.

Use the **-legacy** parameter if you want to establish a partnership with another clustered system that is running an earlier level of code than the local system. The value that is supplied for the **-legacy** parameter must be the code level of the other clustered system.

An invocation example of how to update the node hardware configuration of node ID 7

```
chnodehw 7
```

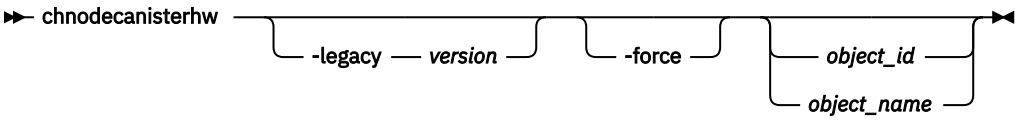
The resulting output:

```
No feedback
```

chnodecanisterhw

Use the **chnodecanisterhw** command to update the hardware configuration for a node canister.

Syntax



Parameters

-legacy *version*
(Optional) Sets the hardware configuration to make it compatible with the 6.3.0.0 code level. The format is four decimal numbers that are separated by periods, and there can be up to 16 characters.

-force
(Optional) Allows the node to change its hardware configuration, even if it might cause the volumes to go offline.

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

object_id | object_name
(Optional) Specifies the object name or ID.

Description

This command automatically reboots the node canister if the canister hardware is different than its configured hardware. After rebooting, the node canister uses its hardware, and does not use the previous configuration.

Attention: When you run the **chnodecanisterhw** command to change the configured hardware for a node canister, these items are removed:

- Small Computer System Interface-3 (SCSI-3) reservations (through that node canister) are removed.
- Small Computer System Interface-3 (SCSI-3) registrations (through that node canister) are removed.

Note: This command fails if you remove the last compression card of a node and try to commit that change while compressed volumes still exist in that I/O group.

Use the **-legacy** parameter if you want to establish a partnership with another system that is running an earlier level of code than the local system. The value that is supplied for with **-legacy** must be the code level of the other system.

An invocation example of how to update the node canister hardware configuration of node canister ID 7

```
chnodecanisterhw 7
```

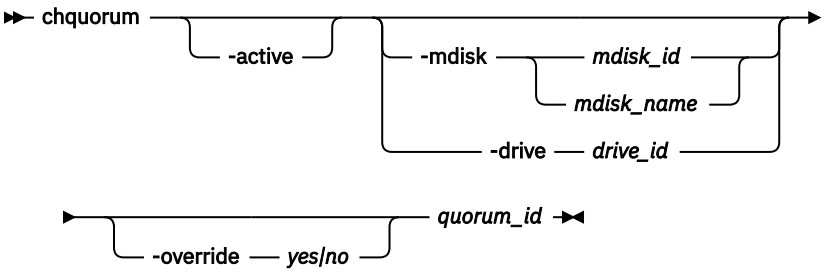
The resulting output:

```
No feedback
```

chquorum

Use the **chquorum** command to change the quorum association.

Syntax



Parameters

-active
(Optional) Makes the specified quorum ID the active one. The **active** parameter must be used if neither the **mdisk** nor the **drive** parameters are specified.

-mdisk *mdisk_id* | *mdisk_name* | **-drive** *drive_id*
(Optional) Specifies the MDisk or drive to be this quorum ID.

-override *yes|no*
Enables the automatic quorum selection to be overridden. In this state, the quorum disk is only moved if the resources are offline. Do not use this parameter unless a specific quorum disk is required for the configuration.

quorum_id
(Required) Specifies which quorum ID to change. Permitted values are values are 0, 1, and 2.

Description

Use the **chquorum** command to change the quorum association. To identify the drive or MDisk that is the current active quorum disk, use the **lsquorum** command.

Remember: You cannot use this command to change the active quorum device when you use an IP quorum application. To change the active IP quorum application, the quorum application must be restarted. The quorum application that connects first is chosen and is active (if valid).

The **chquorum** command is not synchronous, but usually takes only a few seconds to complete. In some situations it can take several minutes.

The system uses the quorum disk or drive as a tie breaker when exactly half of the nodes that were previously a member of the system are present.

Attention: Only assign quorum disks to drives in the control enclosure or to external MDisks. Some maintenance procedures require that quorum is moved temporarily to expansion enclosures. Once that procedure is complete, return the quorum drives to the control enclosure.

The use of a quorum disk or drive allows the system to manage a SAN fault that splits the system exactly in half. One half of the system continues to operate and the other half stops until SAN connectivity is restored.

There is only one quorum disk or drive; however, the system uses three as quorum candidates. The system selects the actual quorum disk or drive from the pool of quorum candidates. The quorum candidates also hold a copy of important system metadata. Just over 256 MB is reserved for this purpose on each quorum candidate disk. When using an MDisk as quorum disk, this space is allocated from the storage pool.

The number of extents required depends on the extent size for the storage pool containing the MDisk. This table provides the number of extents reserved for quorum use by extent size.

Table 24. Number of extents reserved by extent size

Extent size (MB)	Number of extents reserved for quorum use
16	17
32	9
64	5
128	3
256	2
512	1
1024	1
2048	1
4096	1
8192	1

When you issue this command, the MDisk or drive that currently is assigned the quorum index number is set to a nonquorum disk. The system automatically assigns quorum indexes.

You can set the active quorum disk or drive with the **-active** parameter. This can be useful in a system configuration to ensure that the most highly-available quorum disk or drive is used.

An invocation example

```
chquorum -mdisk 45 2
```

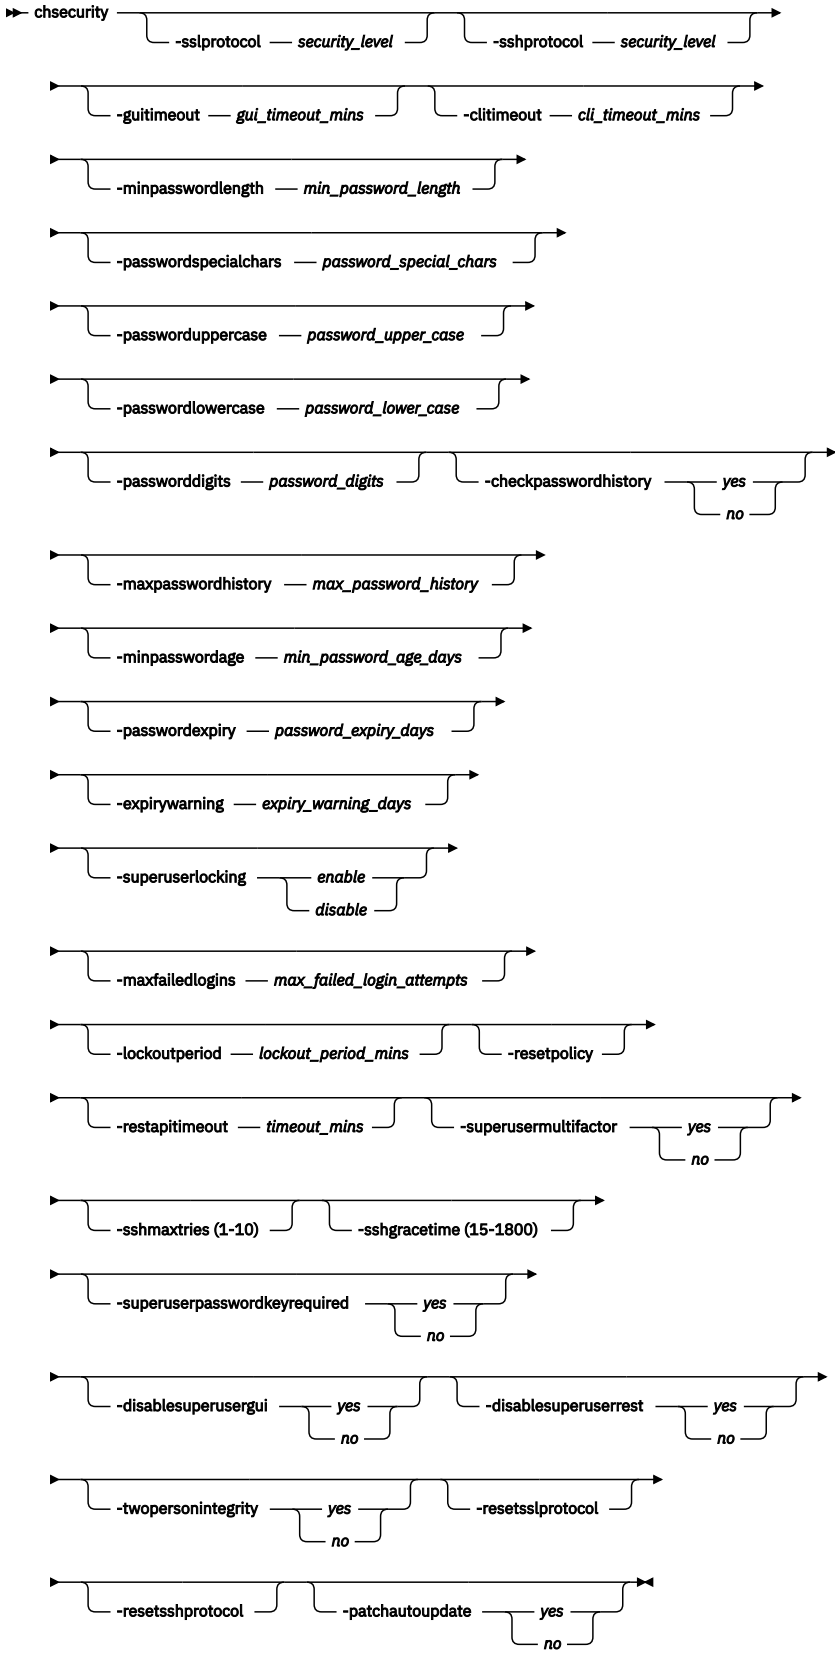
The resulting output:

```
No feedback
```

chsecurity

Use the **chsecurity** command to change the security settings for a system.

Syntax


chsecurity
-sslprotocol security_level
-sshprotocol security_level
-guitimeout gui_timeout_mins
-clitimeout cli_timeout_mins
-minpasswordlength min_password_length
-passwordspecialchars password_special_chars
-passworduppercase password_upper_case
-passwordlowercase password_lower_case
-passworddigits password_digits
-checkpasswordhistory
yes
no
-maxpasswordhistory max_password_history
-minpasswordage min_password_age_days
-passwordexpiry password_expiry_days
-expirywarning expiry_warning_days
-superuserlocking
enable
disable
-maxfailedlogins max_failed_login_attempts
-lockoutperiod lockout_period_mins
-resetpolicy
-restapitimeout timeout_mins
-superusermultifactor
yes
no
-sshmaxtries (1-10)
-sshgracetime (15-1800)
-superuserpasswordkeyrequired
yes
no
-disablesuperusergui
yes
no
-disablesuperuserrest
yes
no
-twopersonintegrity
yes
no
-resetsslprotocol
-resetsshprotocol
-patchautoupdate
yes
no

Parameters

-sslprotocol *security_level*
(Optional) Specifies the numeric value for the SSL security level setting, which can take any value from 2 to 7. A setting of 5 is the default value.

Use these **sslprotocol** security level settings:

- 2 - Allows TLS 1.2, but disallows TLS 1.0 and TLS 1.1.
- 3 - Additionally disallows TLS 1.2 cipher suites that are not exclusive to 1.2.
- 4 - Additionally disallows RSA key exchange ciphers and static key exchange ciphers.
- 5 (Compatibility mode) - Initially allows TLS 1.3, which is the preferred method of connection. If TLS 1.3 fails, TLS 1.2 is used for connections.
- 6 - Allows TLS 1.3 and the five ciphers of TLS 1.3.
- 7 - Allows TLS 1.3 and a single FIPS cipher.

Changing the SSL security level might disable the GUI connection on older web browsers. If connection is lost, use the CLI prompt to change the security level back to a known good level.

-sshprotocol *security_level*
(Optional) Specifies the numeric value for the SSH security level setting, which can take a value of 1 to 4. A setting of 3 is the default value.

Use these **sshprotocol** security level settings.

- 1 Allows the following key exchange methods:
  - curve25519-sha256
  - curve25519-sha256@libssh.org
  - ecdh-sha2-nistp256
  - ecdh-sha2-nistp384
  - ecdh-sha2-nistp521
  - diffie-hellman-group-exchange-sha256
  - diffie-hellman-group16-sha512
  - diffie-hellman-group18-sha512
  - diffie-hellman-group14-sha256
  - diffie-hellman-group14-sha1
  - diffie-hellman-group1-sha1
  - diffie-hellman-group-exchange-sha1
- 2 Allows the following key exchange methods:
  - curve25519-sha256
  - curve25519-sha256@libssh.org
  - ecdh-sha2-nistp256
  - ecdh-sha2-nistp384
  - ecdh-sha2-nistp521
  - diffie-hellman-group-exchange-sha256
  - diffie-hellman-group16-sha512
  - diffie-hellman-group18-sha512
  - diffie-hellman-group14-sha256
  - diffie-hellman-group14-sha1
- 3 Allows the following key exchange methods:

  - curve25519-sha256
  - curve25519-sha256@libssh.org
  - ecdh-sha2-nistp256
  - ecdh-sha2-nistp384
  - ecdh-sha2-nistp521
  - diffie-hellman-group-exchange-sha256
  - diffie-hellman-group16-sha512
  - diffie-hellman-group18-sha512
  - diffie-hellman-group14-sha256
- 4 Allows the following key exchange methods:
  - curve25519-sha256
  - curve25519-sha256@libssh.org
  - ecdh-sha2-nistp256
  - ecdh-sha2-nistp384
  - ecdh-sha2-nistp521

  Restriction: The 3-site-orchestrator does not support SSH protocol level 4.

-guitimeout *gui_timeout_mins*
(Optional) Specifies the amount of time (in minutes) before a session expires and the user is logged out of the GUI for inactivity. The value must be an integer in the range 5 - 240.

-clitimeout *cli_timeout_mins*
(Optional) Specifies the amount of time (in minutes) before a session expires and the user is logged out of the CLI for inactivity. The value must be an integer in the range 5 - 240.

-minpasswordlength *min_password_length*
(Optional) Specifies the minimum length requirement for user account passwords on the system. The value must be an integer in the range 6 - 64.

-passwordspecialchars *password_special_chars*
(Optional) Specifies how many special characters are required in passwords for local users. A value of 0 means that no special characters are required. The value must be an integer in the range 0 - 3.

-passworduppercase *password_upper_case*
(Optional) Specifies how many uppercase characters are required in passwords for local users. A value of 0 means that no uppercase characters are required. The value must be an integer in the range 0 - 3.

-passwordlowercase *password_lower_case*
(Optional) Specifies how many lowercase characters are required in passwords for local users. A value of 0 means that no lowercase characters are required. The value must be an integer in the range 0 - 3.

-passworddigits *password_digits*
(Optional) Specifies how many digits are required in passwords for local users. A value of 0 means that no numbers are required. The value must be an integer in the range 0 - 3.

-checkpasswordhistory *yes | no*
(Optional) Specifies whether the system prevents the user from reusing a previous password. The value is either **yes** or **no**.

This parameter is not supported on FlashSystem 5015, FlashSystem 5035 and FlashSystem 5045.

-maxpasswordhistory *max_password_history*
(Optional) Specifies the number of previous passwords to compare with if **checkpasswordhistory** is enabled. A value of 0 means that the new password is compared with the current password only. The value must be an integer in the range 0 - 10.

-minpasswordage *min_password_age_days*
(Optional) Specifies the minimum number of days between password changes. This setting is enforced if **checkpasswordhistory** is enabled. This restriction is ignored if the password is expired. The setting does nothing if the value is greater than the **passwordexpiry** value. The value must be an integer in the range 0 - 365.

-passwordexpiry *password_expiry_days*
(Optional) Specifies the number of days before a password expires. A value of 0 means that the feature is disabled and passwords do not expire. The value must be an integer in the range 0 - 365.

-expirywarning *expiry_warning_days*
(Optional) Specifies how many days before a password expires to raise a warning. The warning is displayed on every CLI login until the password is changed. A value of 0 means that the feature is disabled and warnings are not displayed. The value must be an integer in the range 0 - 30.

-superuserlocking *enable | disable*
(Optional) Specifies whether the locking policy that is configured on the system also applies to the superuser. The value is either **enable** or **disable**. This parameter is only supported on systems with a dedicated technician port.

-maxfailedlogins *max_failed_login_attempts*
(Optional) Specifies the number of failed login attempts before the user account is locked for the amount of time that is specified in **lockoutperiod**. A value of 0 means that the feature is disabled and accounts are not locked out after failed login attempts. The value must be an integer in the range 0 - 10.

-lockoutperiod *lockout_period_mins*
(Optional) Specifies the number of minutes that a user is locked out for if the max failed logins value is reached. A value of 0 implies the user is indefinitely locked out when the max failed login attempts are reached. The value must be an integer in the range 0 - 10080.

-resetpolicy
(Optional) Resets all security settings to their default values. A **yes / no** warning prompt is displayed to confirm the action. This parameter cannot be run with any other parameters.

-restapitimeout
(Optional) Specifies the amount of time (in minutes) before a token expires. The value is in the range 10 - 120.

-superusermultifactor *yes | no*
(Optional) Specifies whether the superuser should be prompted for multifactor authentication. The value is either **yes** or **no**.

-sshmaxtries
(Optional) Specifies the amount of allowed login attempts per a single SSH connection. The value is in the range 1 - 10.

-sshgracetime
(Optional) Specifies the duration of time a user must enter login factors per SSH connection before the connection is terminated. The value is number of seconds in the range 15 - 1800.

-superuserpasswordkeyrequired
(Optional) Specifies whether the superuser must be expected to provide both a password and SSH key for authentication. The value is either **yes** or **no**.

-disablesuperusergui
(Optional) Specifies whether GUI access must be disabled for the superuser. The value is either **yes** or **no**.

-disablesuperuserrest
(Optional) Specifies whether REST API access must be disabled for the superuser. The value is either **yes** or **no**.

-twopersonintegrity
(Optional) Specifies whether two person integrity (TPI) is enabled or disabled. The value is either **yes** or **no**.

-resetsslprotocol

(Optional) Resets the SSL protocol security level to the default value 5 and configures the system to automatically follow the suggested level.

-resetsshprotocol

(Optional) Resets the SSH protocol security level to default value 3 and configures the system to automatically follow the suggested level.

-patchautoupdate
(Optional) Specifies whether the Patch Auto Update process is enabled or disabled. The value is either **yes** or **no**.

Description

This command changes the security settings on a system.

Important: If you use SSL or TLS, changing the security level might disrupt these services.

Use this procedure if disruption occurs.

1. Wait 5 minutes and try again. (Wait for any services to restart.)
2. Confirm that the SSL or TLS implementation is up to date and supports the specified level of security.
3. If necessary, revert to an earlier version of SSL or TLS security.

Note:

- When you enable two person integrity, **superuser_locking** is also enabled. You can't use the **chsecurity -superuserlocking** command to disable it.
- You must have the Security Administrator role to configure security settings with the **chsecurity** command. On a system where TPI is enabled, any changes that can be made with the **chsecurity** command require an approved role elevation.

An invocation example

```
lsservicestatus | grep superuser
```

The resulting output

```
superuser_locked no
superuser_multi_factor yes
superuser_password_sshkey_required no
superuser_gui_disabled no
superuser_rest_disabled no
```

An invocation example

```
lssecurity | grep grace
```

The resulting output

```
ssh_grace_time_seconds
```

An invocation example

```
chsecurity -sshmaxtries 3
```

The resulting output

```
No feedback.
```

An invocation example

```
lssecurity | grep max_tries
```

The resulting output

```
ssh_max_tries 3
```

An invocation example

```
lsservicestatus | grep superuser
```

The resulting output

```
superuser_locked no
superuser_multi_factor no
superuser_password_sshkey_required yes
superuser_gui_disabled no
superuser_rest_disabled no
```

An invocation example

```
lsservicestatus | grep superuser
```

The resulting output

```
superuser_locked no
superuser_multi_factor no
superuser_password_sshkey_required no
superuser_gui_disabled no
superuser_rest_disabled yes
```

Invocation examples

```
chsecurity -minpasswordlength 12
No feedback.


chsecurity -guitimeout 60
No feedback.


chsecurity -clitimeout 60
No feedback.


chsecurity -superuserlocking enable
Changing the system security settings could result in a loss of access to the system via SSH
or the management GUI. Refer to the Command Line Interface help for more information about the
risks associated with each parameter. Are you sure you wish to continue? (y/yes to confirm)


chsecurity -maxfailedlogins 4
No feedback.


chsecurity -lockoutperiod 60
No feedback.

chsecurity -lockoutperiod 0
No feedback.

chsecurity -passwordexpiry 90
No feedback.

chsecurity -passwordexpiry 60 -expirywarning 14
No feedback.

chsecurity -checkpasswordhistory yes
No feedback.
```

```
chsecurity -maxpasswordhistory 3
No feedback.

chsecurity -minpasswordage 1
No feedback.



svctask chsecurity -passwordspecialchars 0
No feedback.

svctask chsecurity -passworddigits 1
No feedback.

svctask chsecurity -passwordlowercase 2
No feedback.

svctask chsecurity -passworduppercase 3
No feedback.



chsecurity -resetpolicy
Changing the system security settings could result in a loss of access to the system via SSH
or the management GUI. Refer to the Command Line Interface help for more information about the
risks associated with each parameter. Are you sure you wish to continue? (y/yes to confirm)
```

An invocation example

Enabling TPI on the system:

```
svctask chsecurity -twopersonintegrity yes
```

The resulting output

```
No feedback.
```

An invocation example

```
svctask chsecurity -resetsslprotocol
```

The resulting output

```
No feedback.
```

An invocation example

```
svctask chsecurity -resetsshprotocol
```

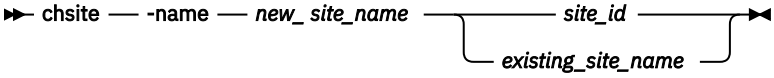
The resulting output

```
No feedback.
```

chsite

Use the **chsite** command to change the site name. This command applies only to models that support multiple I/O groups.

Syntax



Parameters

-name ***new_site_name***
(Required) Specifies the new name for the site.

site_id | existing_site_name
(Required) Specifies the existing site ID or site name that is being changed.

Description

This command changes the site name.

An invocation example

```
chsite -name Quorum 3
```

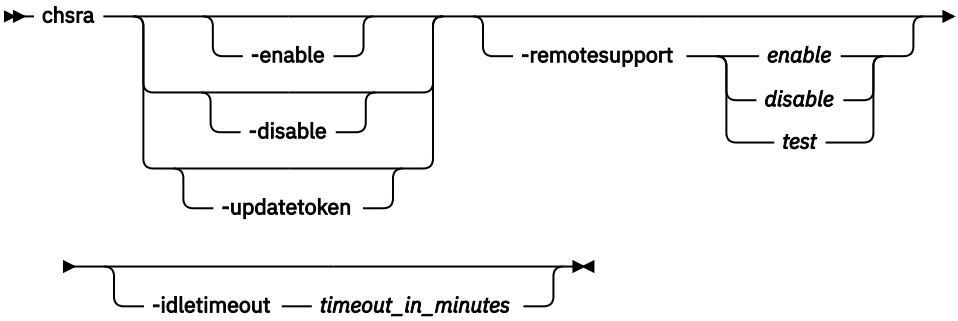
The resulting output:

```
No feedback
```

chsra

Use the **chsra** command to configure support assistance.

Syntax



Parameters

-enable
(Optional) Creates remote access accounts and enables local support assistance.

-disable
(Optional) Deletes all remote access accounts and disables local and remote support assistance.

Note: Remote Support Assistance cannot be disabled on IBM Storage Virtualize for Public Cloud (SVPC) if multifactor authentication has been enabled for the superuser.

-updatetoken
(Optional) Updates the shared security token that is used for support assistance.

-remotesupport ***enable | disable | test***
(Optional) Configures remote support assistance directly over the internet or by a configured Remote Support Proxy server. The values are:

- `enable`
- `disable`
- `test`

No default value exists.

-idletimeout *timeout_in_minutes*
(Optional) Enables remote support for a limited amount of time (specified in minutes). The value must be a positive number (integer) denoting how many minutes remote support assistance is idle (and timed out). This parameter does not time out when a support session is in progress on any of your system nodes. It is renewed as many times as needed and only times out after all active sessions are terminated.

Remember: If the idle timeout expires on all participating nodes in a system, remote system support is disabled. If remote system support is disabled, remote system support is not automatically started on events that include:

- When new nodes join the system.
- T3 recovery procedures.
- Node warm or cold starts.

If you specify **-idletimeout**, you must also specify **-remotesupport**.

Description

This command configures local or remote support assistance.

Note: Turn on both local and remote support assistance to more efficiently resolve any problems that are encountered.

Remote support assistance is available either directly over the internet or by using a Remote Support Proxy server. Remote system support is routed by using the Remote Support Proxy server if any Remote Support Proxy servers are configured.

Note:

- Remote Support Proxy servers are deprecated. You can configure HTTP proxy servers instead by using **mkproxy** or **chproxy** commands. You can configure your firewall to allow traffic to pass directly from the system or you can route traffic through an HTTP proxy server within your environment. Do not configure an HTTP proxy server by using the **mksystemsupportcenter** command. For more information, see HTTP proxy server.
- With the addition of the HTTP proxy support, Remote Support Proxy servers are no longer necessary, but they are still supported for existing configurations. Alternately, a Remote Support Proxy can be configured to consolidate firewall traffic from a number of storage systems. Remote upgrades cannot be completed through the Remote Support Proxy server.

You must do the following to turn on remote support assistance:

1. Configure service IP addresses on all nodes.
2. Configure call home on the system.
3. Configure local support assistance on your system, which creates support and sets up authentication. For storage systems that have direct access to the internet, the firewall must allow outbound connections to Internet Protocol (IP) addresses `170.225.126.11`, `170.225.126.12`, `170.225.127.11`, and `170.225.127.12` on port 22 or 443. If you must use a Remote Support Proxy server, configure it by using the **mksystemsupportcenter** command.

   For remote access, port 443 is suggested as it provides more security. If the connection from port 443 remote access is unsuccessful, then port 22 is used as a backup. When you configure the firewall rules with port 443, the fallback of remote access from port 443 to port 22 is not applicable.

An invocation example for creating support assistance accounts and enabling local support assistance

```
chsra -enable
```

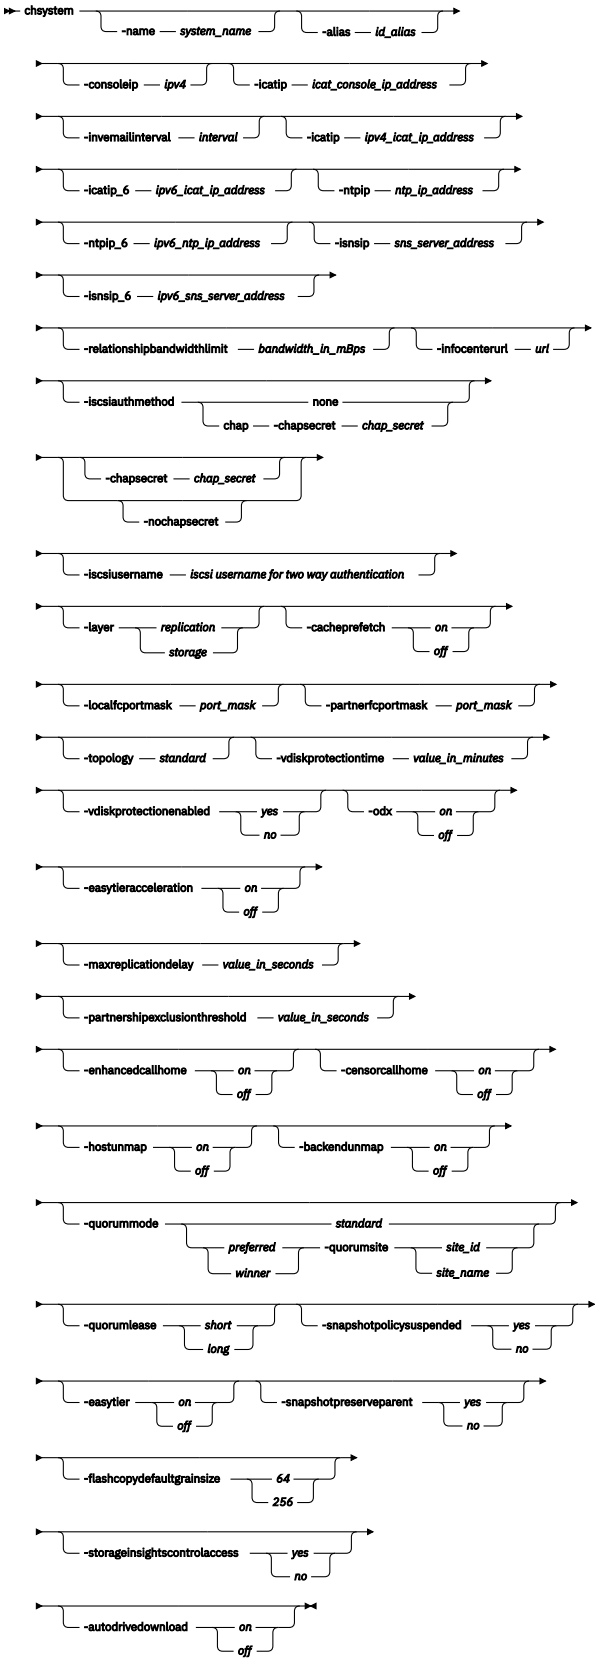
The detailed resulting output:

```
No feedback
```

chsystem

Use the **chsystem** command to modify the attributes of an existing system. Enter this command any time after a system is created. All the parameters that are associated with this command are optional. However, you must specify one or more parameters with this command.

Syntax


chsystem
-name — system_name
-alias — id_alias
-consoleip — ipv4
-icatip — icat_console_ip_address
-invemailinterval — interval
-icatip — ipv4_icat_ip_address
-icatip_6 — ipv6_icat_ip_address
-ntpip — ntp_ip_address
-ntpip_6 — ipv6_ntp_ip_address
-isnsip — sns_server_address
-isnsip_6 — ipv6_sns_server_address
-relationshipbandwidthlimit — bandwidth_in_mBps
-infocenterurl — url
-iscsiauthmethod
none
chap — -chapsecret — chap_secret
-chapsecret — chap_secret
-nochapsecret
-iscsiusername — iscsi username for two way authentication
-layer
replication
storage
-cacheprefetch
on
off
-localfcportmask — port_mask
-partnerfcportmask — port_mask
-topology — standard
-vdiskprotectiontime — value_in_minutes
-vdiskprotectionenabled
yes
no
-odx
on
off
-easytieracceleration
on
off
-maxreplicationdelay — value_in_seconds
-partnershipexclusionthreshold — value_in_seconds
-enhancedcallhome
on
off
-censorcallhome
on
off
-hostunmap
on
off
-backendunmap
on
off
-quorummode
standard
preferred
winner
-quorumsite
site_id
site_name
-quorumlease
short
long
-snapshotpolicysuspended
yes
no
-easytier
on
off
-snapshotpreserveparent
yes
no
-flashcopydefaultgrainsize
64
256
-storageinsightscontrolaccess
yes
no
-autodrivedownload
on
off

Parameters

-name *system_name*
(Optional) Specifies a new name for the system.

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

-alias *id_alias*
(Optional) Specifies an alternative name that does not change the basic ID for the system, but does not influence the `VDisk_UID` of newly created volumes. Existing `VDisk_UIDs` are not changed.

-consoleip *ipv4*

(Optional) Specifies the system's new management IPv4 address that is used by the system. The format of this IP address must be a dotted decimal notation with the port; for example, 255.255.255.255.

-invemailinterval *interval*
(Optional) Specifies the interval at which inventory emails are sent to the designated email recipients. The interval range is 0 to 15. The interval is measured in days. Setting the value to 0 turns off the inventory email notification function.

-icatip *icat_console_ip_address*
(Optional) Specifies the system's new IPv4 address that is used by the system. The format of this IP address must be a dotted decimal notation with the port; for example, `255.255.255.255:8080`. If you specify this parameter, it overwrites any existing `-icatip_6` address.

-icatip_6 *icat_console_ipv6_address*
(Optional) Specifies the system's new IPv6 address. If you specify this parameter, it overwrites any existing **-icatip** address. The format of the IPv6 address must be:

- Eight colon-separated groups of four hexadecimal digits; for example:

```
[1234:1234:abcd:0123:0000:0000:7689:6576]:23
```

- Eight colon-separated groups of hexadecimal digits with leading zeros omitted; for example:

```
[1234:1234:abcd:123:0:0:7689:6576]:23
```

- Suppression of one or more consecutive all 0 groups; for example:

```
[1234:1234:abcd:123::7689:6576]:23
```

-ntpip *ntp_ip_address*
(Optional) Specifies the IPv4/IPv6 address or fully qualified domain name (FQDN) for the Network Time Protocol (NTP) server. The maximum number of characters is 256. Configuring an NTP server address causes the system to use that NTP server as its time source. Specify the **-ntpip** parameter with a zero address to use another time source:

```
chsystem -ntpip 0.0.0.0
```

-ntpip_6 *ipv6_ntp_ip_address*

Note: Before you specify **-ntpip_6**, an IPv6 prefix and gateway must be set for the system.

(Optional) Specifies the IPv6 address for the NTP server. Configuring an NTP server address causes the system to immediately start using that NTP server as its time source. To choose another time source, specify the **-ntpip_6** parameter with a zero address, as follows:

```
chsystem -ntpip_6 0::0
```

-isnsip *sns_server_address*
(Optional) Specifies the IPv4 address for the iSCSI storage name service (SNS). Specify the **-isnsip** parameter with a zero address to select another IPv4 iSCSI SNS server:

```
chsystem -isnsip 0.0.0.0
```

-isnsip_6 *ipv6_sns_server_address*
(Optional) Specifies the IPv6 address for the iSCSI SNS. Specify the **-isnsip_6** parameter with a zero address to select another configured IPv6 iSCSI SNS server:

```
chsystem -isnsip_6 0::0
```

-relationshipbandwidthlimit *bandwidth_in_mBps*
(Optional) Deprecated. This setting is not used by the system.

-infocenterurl *url*
Specifies the preferred online documentation URL to override the one used by the GUI. Because this information is interpreted by the Internet browser, the specified information might contain a hostname or an IP address.

Remember: View the currently configured URL in the GUI preferences window. This window can also help reset this value to the default setting.

-iscsiauthmethod none | chap -chapsecret *chap_secret*
(Optional) Sets the authentication method for the iSCSI communications of the system:

- chap indicates Internet Small Computer System Interface (iSCSI) authentication is turned on.

  Remember: This turns on iSCSI partnership authentication when a Challenge Handshake Authentication Protocol (CHAP) secret key is set for the system.
- none indicates that iSCSI partnership authentication is turned off.

-chapsecret*chap_secret*
(Optional) Sets the CHAP secret to be used to authenticate the system that uses iSCSI. This parameter is required if the **iscsiauthmethod chap** parameter is specified. The specified CHAP secret cannot begin or end with a space.

-nochapsecret
(Optional) Clears any previously set CHAP secret for iSCSI authentication. The **nochapsecret** parameter cannot be specified when **chapsecret** is specified.

-iscsiusername
(Optional) Specifies the user name for the entire SVC system that is used for two-way authentication for iSCSI host attach login. If this parameter is specified, this value is taken as "username" for two-way authentication in iSCSI host attach login. If user name is not specified, user name for two-way chap authentication is NULL.

-layer *replication | storage*
(Optional) Specifies which layer a system is in. The system can create partnerships with systems in the same layer.

Note: If you specify **-layer** you must specify either *replication* or *storage*. This option can be used if no other systems are visible on the fabric, and no system partnerships are defined.

-cacheprefetch *on | off*
(Optional) Indicates whether cache prefetching is enabled or disabled across the system. Adjust this only when following direction from your product support information.

-localfcportmask *port_mask*
(Optional) Indicates the Fibre Channel (FC) input/output (I/O) port mask for local system node-to-node communications only. Fibre Channel port mask does not affect host or storage system traffic. The port mask is 64 binary bits and is made up of a combination of 0's and 1's, where 0 indicates that the corresponding FC I/O port cannot be used and 1 indicates that it can be used. The mask is applied to all nodes in the local system. At least two ports must be selected for local system node-to-node communications. The mask must result in at least

2 FC connections between each node in the local system, using only the selected ports and FC zones visible to those ports. Valid mask values might range from `0011` (only ports 1 and 2 enabled) to `1111111111111111111111111111111111111111111111111111111111111111` (all ports enabled). For example, a mask of `111111101101` enables ports 1, 3, 4, 6, 7, 8, 9, 10, 11, and 12.

Remember: A partial mask (fewer than 64 characters) is zero-extended, meaning that any ports that are not specified are not enabled.

Specify the **lsportfc** command to display FC I/O port IDs.

-partnerfcportmask *port_mask*
(Optional) Indicates the FC I/O port mask for partnered system-to-system communications only. Fibre Channel port mask does not affect host or storage system traffic. The port mask is 64 binary bits and is made up of a combination of 0's and 1's, where `0` indicates that the corresponding FC I/O port cannot be used and 1 indicates that it can be used. The mask is applied to all nodes in the local system. Valid mask values might range from `0000` (no ports enabled) to `1111111111111111111111111111111111111111111111111111111111111111` (all ports enabled). For example, a mask of `111111101101` enables ports 1, 3, 4, 6, 7, 8, 9, 10, 11, and 12.

Remember: A partial mask (fewer than 64 characters) is zero-extended, meaning that any ports not specified are not enabled.

Specify the **lsportfc** command to display FC I/O port IDs.

-hightempmode *on | off*
(Optional) Sets (or removes) High Temperature Mode (HTM). The values are on and `off`.

-topology *standard*
(Optional) Indicates the intended *standard* system topology.

-vdiskprotectiontime *value_in_minutes*
(Optional) Sets volume protection time (in minutes). The range is 15 - 1440.

-vdiskprotectionenabled *yes | no*
(Optional) Enables or disables volume protection. The values are `yes` and no.

-odx *on | off*
(Optional) Enables or disables offloaded data transfers (ODX). The values are on and `off`.

-easytieracceleration *on | off*
(Optional) Enables Easy Tier® and pool balancing acceleration. The values are on and `off`.

-maxreplicationdelay *value_in_seconds*
(Optional) Deprecated. This setting is not used by the system.

-partnershipexclusionthreshold *value_in_seconds*
(Optional) Deprecated. This setting is not used by the system.

-enhancedcallhome *on | off*
(Optional) Specifies that the call home function is to send enhanced reports to the support center. Valid values are `yes` or no.

The enhanced reports include operational and event-related data and specific configuration information that is included in the inventory report. This function alerts the support center about hardware failures and potentially serious configuration or environmental issues. The support center can use the configuration information to automatically generate best practices or recommendations that are based on your actual configuration.

-censorcallhome *on | off*
(Optional) Specifies that sensitive data is deleted from the enhanced call home data. The values are `yes` or no.

-hostunmap *on | off*
(Optional) Specifies whether the administrator wants the volumes to advertise support for unmap to hosts. The values are on (default) or `off`.

-backendunmap ***on* | *off***
(Optional) Specifies whether the administrator wants to enable unmap support for backend storage. Enabling this parameter allows the system to issue Unmap or write same commands to compatible backend storage controllers and drives to reclaim capacity when capacity is freed within the system. The values are on (default) or off.

Note: If **-backendunmap** off is specified, the system may still issue unmap commands to drives if they are in a storage pool where the format_required property is yes.

-snapshotpolicysuspended ***yes* | *no***
(Optional) Specifies the suspend or resume of the scheduler function. The valid values are *yes* or *no*.

Note: If **-snapshotpolicysuspended** parameter *yes* is specified, snapshots will no longer be automatically deleted or created by the snapshot scheduler.

Specify **-snapshotpolicysuspended** parameter *no* to resume the normal operations.

-easytier ***on* | *off***
(Optional) Enables or disables the Easy Tier function on a system level. The values are *on* and *off*.

Note: After the completion of T3 recovery, the system disables the Easy Tier function by default. You can specify the **-easytier** *on* to enable this function on the system level.

-snapshotpreserveparent ***yes* | *no***
(Optional) Specifies the snapshot and parent preservation behavior. The valid values are *yes* and *no*.

Note: If the value is set to *no*, the parent volume for snapshots goes offline when the system cannot maintain the snapshot. If the value is set to *yes*, the parent volume remains online and the snapshot is deleted. You can specify a single pool for the mirrored parent volumes if the **-snapshotpreserveparent** parameter is set to yes and the snapshot that is created is mirrored with a single copy in the specified pool.

-flashcopydefaultgrainsize ***64* | *256***
(Optional) Use this parameter to set the default grain size of all FlashCopy mappings created after this value has been set. If the volume of a new FlashCopy mapping is already associated with mappings of a different grain size then the existing grain size that volume is using will be continued to be used.

-storageinsightscontrolaccess ***yes* | *no***
(Optional) Indicates whether the storage insights control access for the system is enabled or disabled.

-autodrivedownload ***on* | *off***
(Optional) Specifies whether the auto drive download feature for the system is enabled or disabled. The values are on (default) or off.

Description

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

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

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

Use the **chsystemip** command to modify the system IP address and service IP address.

Remember: Setting a CHAP secret key for the system does not turn on authentication for iSCSI hosts partnerships. Turn off authentication by issuing **-iscsiauthmethod**.

Note: Only Security Admin role users can execute the parameters **-ntpip** and **-ntpip_6**.

An invocation example to set up an external NTP server

```
chsystem -ntpip 123.234.123.234
```

The resulting output:

```
No feedback
```

An invocation example to change the preferred `infocenterurl` value

```
chsystem -infocenterurl http://miscserver.company.com/ibm/infocenter
```

The resulting output:

```
No feedback
```

An invocation example to change the local port mask value

To indicate that FC I/O ports with IDs 3 and 4 of an eight port system are capable of local node communication:

```
chsystem -localfcportmask 00001100
```

The resulting output:

```
No feedback
```

An invocation example to change the partner port mask value

To indicate that FC ports with IDs 7 and 8 of an eight port system are capable of remote node communication:

```
chsystem -partnerfcportmask 11000000
```

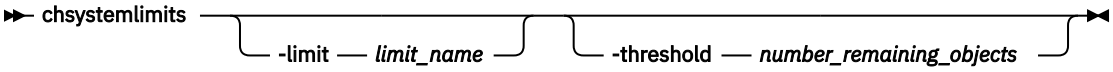
The resulting output:

```
No feedback
```

chsystemlimits

Use the **chsystemlimits** command to modify the warning threshold for a particular object.

Syntax



Parameters

-limit *limit_name*
(Optional) Specifies the name of the limit to modify.

-threshold *number_remaining_objects*
(Optional) Specifies the value at which warning is raised if the available object goes below.

An invocation example

```
lssystemlimits
```

The resulting output:

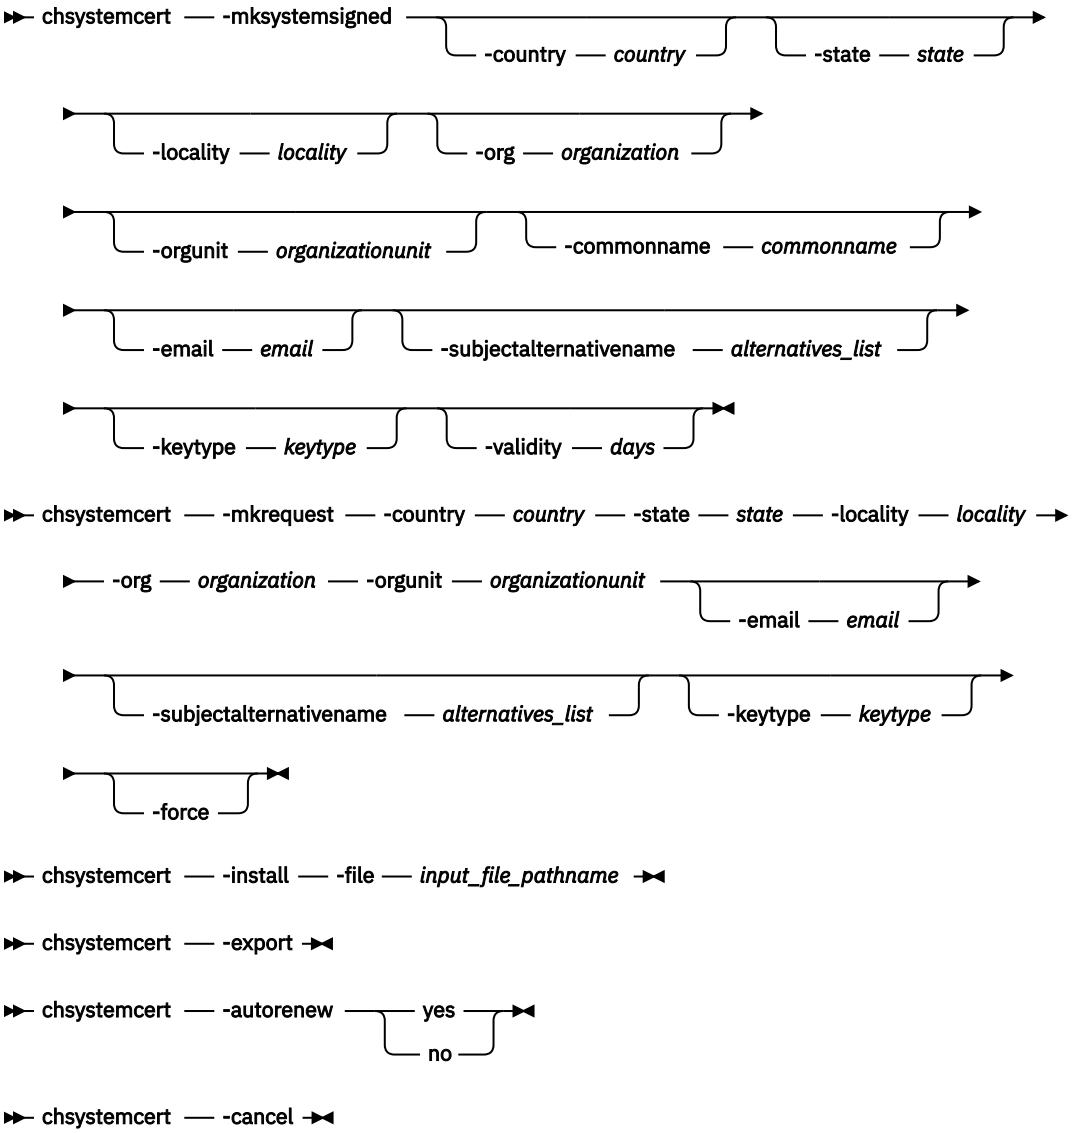
```
type               configured max   available warning_threshold
volume_copies      0          32500 32500     0
volume_snapshots   1          31999 31997     0
flashcopy_mappings 0          15864 15864     0
flashcopy_groups   0          500   500       0
volumes            1          15864 15863     0
```

Note: A **warning_threshold** of 0 indicates that no warning is logged for this type of object.

chsystemcert

Use the **chsystemcert** command to manage the Secure Sockets Layer (SSL) certificate that is installed on a system.

Syntax



Parameters

-mksystemsigned
(Optional) Generates a new certificate that is signed by the system's root certificate authority. The certificate is activated immediately.

-mkrequest
(Optional) Generates a new certificate signing request in `/dumps/certificate.csr`. The file can be copied from the system and provided to a trusted third-party certificate authority. The certificate authority signs the request and provides a signed certificate that can be installed using the **-install** and **-file** parameters.

-country *country*
(Optional) Specifies the two digit country code.

-state *state*
(Optional) Specifies the state information for the certificate request. The value can be an ASCII string in the range `0` - 128 characters.

-locality *locality*
(Optional) Specifies the locality information for the certificate request. The value can be an ASCII string in the range 0 - 128 characters.

-org *organization*
(Optional) Specifies the organization information for the SSL certificate. The value can be an ASCII string in the range 0 - 64 characters.

-orgunit *organizationunit*
(Optional) Specifies the organization unit information for the SSL certificate. The value can be an ASCII string in the range 0 - 64 characters.

-commonname *commonname*
(Optional) Specifies the common name for the SSL certificate. The value can be an ASCII string in the range 0 - 64 characters.

-email *email*
(Optional) Specifies the email address that is used in the SSL certificate. The value can be an ASCII string in the range 0 - 64 characters.

-subjectalternativename *subject_alternative_name*
(Optional) This parameter allows value to be specified for the Subject Alternative Name certificate extension field permitted in X.509 version 3 certificates. You can specify this parameter only with, **-mksystemsigned** or **-mkrequest** options. The parameter can specify a value up to 512 characters in length. To include some characters such as whitespace, newline or other special characters, apply the appropriate bash command line modifications to ensure the value is specified correctly. This is particularly important if non-character delimiters are used to specify multiple alternative names.

-keytype *keytype*
(Optional) Specifies the SSL certificate key type.

- rsa2048
- ecdsa384
- ecdsa521
- rsa4096

-validity *days*
(Optional) Specifies the number of days (`1-9000`) that the internally-signed certificates is valid.

-install
(Optional) Specifies the installing of certificate.

-file *input_file_pathname*
(Optional) Specifies the absolute path name of the certificate to install.

-export
(Optional) Exports the full chain of installed SSL certificates. The certificate is exported to the `/dumps/certificate.pem` directory on the configuration node.

-autorenew yes | no
(Optional) Turns automatic renewal on or off. When turned on, the system certificate will be renewed automatically 30 days before it expires. The **-autorenew** option is only supported when the certificate is signed by the system's root certificate authority. The default value when using an internally signed certificate is yes.

-cancel
(Optional) Cancel the certificate signing request to the external certificate authority that is currently in progress.

Description

Use this command to manage the SSL certificate that is installed on a system. The command can be used for the following items.

- Generate an internally-signed certificates that is signed by the system's root certificate authority (CA). The root certificate has a long validity period and can be installed on browsers, devices and applications that support chain of trust checking. Internally-signed certificates can be renewed automatically.
- Create a certificate signing request which is copied from the system and sent to an external certificate authority to sign.
- Install an externally signed certificate on to the system.
- Export the full chain of installed certificates.

Important: You must specify one of the following parameters:

- **-mkrequest**
- **-mksystemsigned**
- **-install**
- **-export**
- **-cancel**
- **-autorenew**

An invocation example

The following example shows creating a new system certificate that is signed by the system’s internal root CA with automatic renewal enabled:

```
svctask chsystemcert -mksystemsigned -commonname virtualize -country GB -locality Manchester
-org IBM -orgunit Systems -email support@ibm.com -keytype rsa2048
-validity 355 -subjectalternativename "DNS:test.ibm.com"
```

The resulting output

```
No feedback.
```

An invocation example

The following example shows creating a new certificate signing request that can be provided to an external CA. The resulting CSR file is located at /dumps/certificate.csr on the configuration node:

```
chsystemcert -mkrequest -commonname virtualize -country GB -locality Manchester -org IBM
-orgunit Systems -state England -email support@ibm.com -keytype rsa4096 -subjectalternativename
"DNS:test.ibm.com"
```

The resulting output

```
No feedback.
```

An invocation example

The following example shows installing a certificate that has been signed by an external CA and copied to the /dumps directory on the configuration node:

```
svctask chsystemcert -install -file /dumps/mycertificate.pem
```

The resulting output

```
No feedback.
```

An invocation example

The following example shows exporting the currently installed certificate chain to the /dumps directory on the configuration node:

```
svctask chsystemcert -export
```

The resulting output

```
No feedback.
```

An invocation example

The following example shows cancelling an outstanding certificate request:

```
svctask chsystemcert -cancel
```

The resulting output

```
No feedback.
```

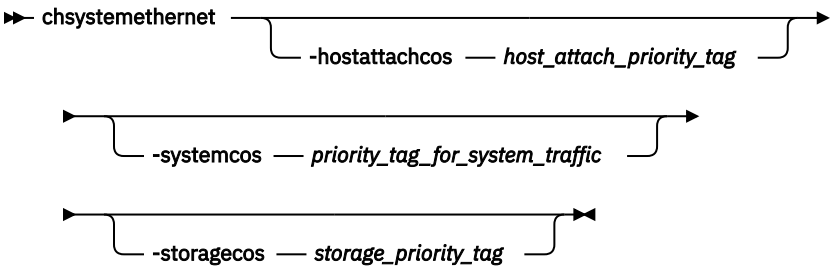
chsystemip (Deprecated)

The **chsystemip** command is deprecated and can not be used to change or remove IP addresses configured on a port other than 1 and 2 or if it is configured using VLAN. To change the port use **chip** command and to remove the port use **rmip** command.

chsystemethernet

The **chsystemethernet** command is used to set the priority tag for host attachment, system, and storage traffic.

Syntax



Parameters

-hostattachcos ***host_attach_priority_tag***
(Optional) Specify the system-wide host attach priority tag that you want to be used by all host attach traffic over iSCSI, or NVMeF. The value must be in the range 0 (default) - 7.

-systemcos ***priority_tag_for_system_traffic***
(Optional) Specify the system-wide clustering priority tag that you want to be used by system traffic. The value must be in the range 0 (default) - 7.

Note: After changing `-systemcos`, you must run **satask stopnode -warmstart** orderly on all the nodes in the system.

-storagecos ***storage_priority_tag***
(Optional) Specify the system-wide storage priority tag that you want to be used by storage traffic. The value must be in the range 0 (default) - 7.

Description

Use the **chsystemethernet** command to set system-wide Ethernet related properties. This command affects all existing and new iSCSI sessions. The priority of iSCSI packets changes immediately after the priority tag is changed by using the **chsystemethernet** command.

Note: PFC is supported on Lancer, Mellanox, and Chelsio. PFC is not supported on Intel.

An invocation example to set the host attach priority tag (class of service) to 4, system priority tag to 5, and storage priority tag to 6.

```
chsystemethernet -hostattachcos 4 -systemcos 5 -storagecos 6
```

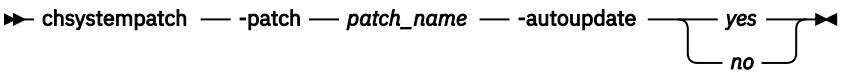
The resulting output:

```
No feedback
```

chsystempatch

Use the **chsystempatch** the command to enable or disable the patch auto-update process for a specific installed patch. If the patch does not support this process, a corresponding error message is returned.

Syntax



Parameters

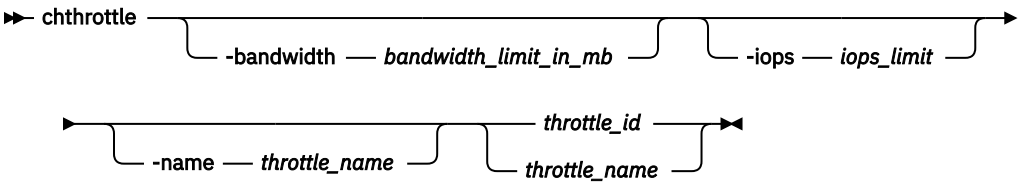
-patch ***patch_name***
Indicates the name of the installed patch for which the auto-update process is enabled or disabled.

-autoupdate ***yes no***
Indicates whether the auto-update process is enabled or disabled for a patch.

chthrottle

Use the **chthrottle** command to change attributes associated with a specified throttle object.

Syntax



Parameters

-bandwidth *bandwidth_limit_in_mb*
(Optional) Specifies the bandwidth in MBps. This must be a numeric value 0 - 268435456.

Note: No bandwidth limit is set unless you specify this keyword.

-iops *iops_limit*
(Optional) Specifies the I/O operations limit. This must be a numeric value 0 - 33554432.

Note: No I/O operations limit is set unless you specify this keyword.

-name *throttle_name*
(Optional) Specifies the throttling object's name. This value must be an alphanumeric string up to 63 characters long.

throttle_id | throttle_name
(Required) Specifies the volume ID or name of the volume to throttle. The value must be a numeric or alphanumeric string up to 15 characters long.

Description

This command changes attributes associated with a specified throttle object.

Note: The throttle limits are per node.

An invocation example for changing the bandwidth limit to 100 for an offloaded throttle

```
chthrottle -bandwidth 100 offloadThrottle
```

The detailed resulting output:

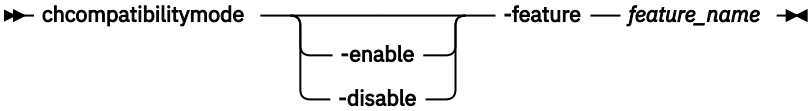
```
No feedback
```

chcompatibilitymode

Use the **chcompatibilitymode** command to change whether a feature is enabled in compatibility mode.

Note: Multiple I/O groups have been replaced by FlashSystem grid as the method to scale out the storage system.

Syntax



Parameters

-enable
(Optional) Specifies that the feature is enabled.

-disable
(Optional) Specifies that the feature is disabled.

-feature ***feature_name***
(Required) Specifies the name of the feature to enable or disable.

cpdumps

Use the **cpdumps** command to copy dump files from a nonconfiguration node onto the configuration node.

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

Syntax



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 memory dump or log files in that directory are retrieved. You can use the following directory arguments (filters):

- `/dumps` (retrieves all files in all subdirectories).
- `/dumps/audit`
- `/dumps/cimom`
- `/dumps/cloud`
- `/dumps/configs`
- `/dumps/elogs`
- `/dumps/easytier`
- `/dumps/feature`
- `/dumps/iostats`
- `/dumps/iotrace`
- `/dumps/mdisk`
- `/home/admin/update`
- `/home/tejas`

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 CLI:

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

```
>cleardumps -prefix "/dumps/elogs/*.txt"
```

node_id | node_name
(Required) Specifies the node from which to retrieve the memory 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 system.
- The node ID that is assigned to the node (not the worldwide node name).

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

Description

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

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

You can view the contents of the directories by using the **lsdumps** command. You can track the status of a copy using the **lscopystatus** command.

An invocation example

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

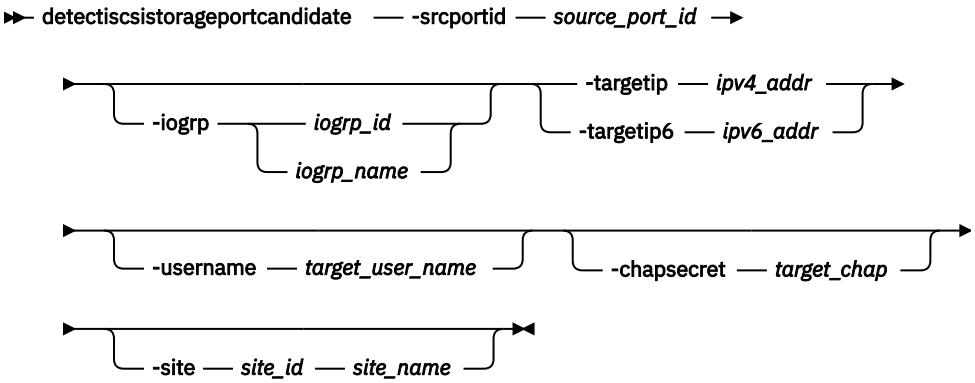
The resulting output:

```
No feedback
```

detectiscsistorageportcandidate

Use the **detectiscsistorageportcandidate** command to establish Internet Small Computer Systems Interface (iSCSI) login sessions from any nodes in a specified I/O group to a discovered backend iSCSI target controller.

Syntax



Parameters

-srcportid ***source_port_id***
(Required) Specifies the source Ethernet port ID (indicated in the **lsportethernet** output) used to complete target controller discovery. The value must be a number 1 - 8.

- If you also specify **-iogrp**, you trigger discovery through the Ethernet port by using the *source_port_id* specified on all nodes in the I/O group.
- If you do not specify **-iogrp**, you trigger discovery through the Ethernet port by using the *source_port_id* specified on all nodes in the system.

-iogrp *iogrp_id | iogrp_name*
(Optional) Specifies I/O group ID or name being detected. The *iogrp_id* value must be 0, 1, 2, or 3. The *iogrp_name* value must be an alphanumeric string.

If you specify this parameter, you trigger discovery through the *source_port_id* of both nodes for the specified I/O group. If no value is specified, discovery is triggered by using the *source_port_id* of all nodes in the system.

-targetip *ipv4_addr*
(Required if you do not specify **-targetip6**) Specifies the target iSCSI controller IPv4 address being detected that receives target discovery requests by using the *source_port_id* for all nodes in the specified I/O group.

-targetip6 *ipv6_addr*
(Required if you do not specify **-targetip**) Specifies the target iSCSI controller IPv6 address being detected that receives target discovery requests by using the *source_port_id* for all nodes in the specified I/O group.

-username *target_user_name*
(Optional) Specifies the target controller user name being detected. The value must be an alphanumeric string up to 256 characters.

If the target controller requires a *target_user_name* and *target_chap* for discovery, use the target user name and Challenge-Handshake Authentication Protocol (CHAP) secret to discover the iSCSI target controller.

Some controllers might require that you use the iSCSI qualified name (IQN) user name for discovery. Each nodes IQN is picked up automatically and used if required.

-chapsecret *target_chap*
(Optional) Specifies the *target_chap* required for discovery of the target iSCSI controller that is being detected. The value must be an alphanumeric string (case-sensitive) up to 79 characters.

-site *site_id | site_name*
(Optional) Specifies the site ID or site name of the target iSCSI controller that is being detected or discovered. If no I/O group is specified, discovery is done from all nodes present in the specified site and if an I/O group is specified discovery is done only from a node that is part of the specified site and I/O group. The site ID must be 1 (the default) or 2. The site name must be an alphanumeric value.

Description

This command detects iSCSI backend storage controllers for migration and virtualization. This command helps with target iSCSI controller discovery. Use the **lsiscsistorageportcandidate** command to list information about the discovered candidate iSCSI target controller(s).

The target data is available until either another discovery is completed or the system undergoes a recovery procedure, which clears the previous discovery data. The command completes when either discovery from all source nodes completes or the command fails.

Use the **addiscsistorageport** command to establish sessions from any nodes in a specified I/O group to a discovered backend iSCSI controller. Use the **lsiscsistorageportcandidate** command to list information about the candidate iSCSI target controller.

A target discovery invocation example with IP address

This example shows target discovery using an IPv4 IP address with a target ISCSI controller and source port ID 1.

```
detectiscsistorageportcandidate -srcportid 1 -targetip IP1
```

The detailed resulting output:

```
No feedback
```

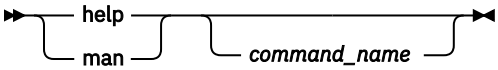
dumpconfig (Discontinued)

Attention: The **dumpconfig** command is discontinued.

help

Use the **help** (or **man**) command to display help information for system commands.

Syntax



Parameters

command_name
(Optional) Indicates the command name.

Description

Use this command to display help information for system commands. If you specify a command name using *command_name*, the complete help file text for the command is displayed. If you do not specify a command name, a comprehensive list of all commands is displayed (with one brief descriptive line).

An invocation example

```
help
```

The resulting output:

```
addhostiogrp - Maps I/O groups to a host object.
addhostport - Adds worldwide port names (WWPNs) or iSCSI names to a host object.
addmdisk - Adds managed disks to a storage pool.
addnode - Adds a new (candidate) node canister to an existing system.
...
```

lsavailablepatch

Use the **lsavailablepatch** command to list the patches that are compatible with the software version on the system and those of which that can be auto-updated. The command output also indicates whether the specific patches listed are already installed on the system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

Description

This table provides the attribute values that can be displayed as output view data.

Table 25. ***lsavailablepatch*** *output*

Attribute	Possible Values
id	Indicates the index of the listed patch.
patch_name	Indicates the name of the patch.
installed	Indicates whether the patch is currently installed on the system.
auto_updateable	Indicates whether the patch can be auto-updated.
patch_description	Indicates the description of the patch.

lsclustercandidate (Discontinued)

Attention: The **lsclustercandidate** command has been discontinued. Use the **lspartnershipcandidate** command instead.

lscluster (Discontinued)

Attention: The **lscluster** command is discontinued. Use a combination of the **lspartnership**, **lspartnershipcandidate**, and **lssystem** commands instead.

lsclusterip (Discontinued)

Attention: The **lsclusterip** command has been discontinued. Use the **lssystemip** command instead.

lsclusterstats (Discontinued)

Attention: The **lsclusterstats** command is discontinued. Use the **lssystemstats** command instead.

lsdiscoverystatus

Use the **lsdiscoverystatus** command to determine whether a discovery operation is in progress.

Syntax

Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

Description

This command displays the state of all discoveries in the cluster. During discovery, the system updates the drive and MDisk records. You must wait until the discovery finishes and is inactive before you attempt to use the system. This command displays one of the following results:

active
There is a discovery operation in progress at the time that the command is issued.

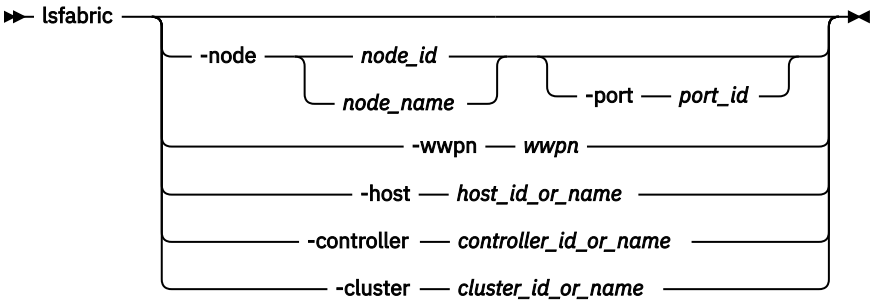
inactive
There are no discovery operations in progress at the time that the command is issued.

If the Fibre Channel functions are used only to enable the nodes to cluster, then the Fibre Channel line are not displayed in the **lsdiscoverystatus** command. The *fc_fabric* line appears if there is at least one Fibre Channel controller.

lsfabric

Use the **lsfabric** command to generate a report that displays the Fibre Channel (FC) connectivity between nodes, controllers, and hosts.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-node *node_name | node_id*
(Optional) Displays the output for all ports for the specified node. The only parameter that you can specify with the **-node** parameter is the **-port** parameter.

-port *port_id*
(Optional) Displays a concise view of all worldwide port names (WWPNs) that are logged into the specified port ID and node. The **-port** parameter must be specified with only the **-node** parameter. A valid *port_id* value is from a minimum of one through a maximum equal to the number of node Fibre Channel (FC) IO ports. It specifies the port number in the vital product data (VPD) or the hexadecimal WWPN of the local port.

-wwpn *wwpn*
(Optional) Displays a list of all ports that have a login to the specified WWPN. You cannot use the **-wwpn** parameter with any other parameter.

-host *host_id_or_name*
(Optional) Specifies a host name or ID. Issuing the **lsfabric** command with the **-host** parameter is equivalent to issuing the **lsfabric wwpn** *wwpn* command for every configured WWPN of the specified host. For example, a host with two ports that are zoned to one port of every node in an eight-node system produces 16 lines of output. You cannot use the **-host** parameter with any other parameter.

-controller *controller_id_or_name*
(Optional) Specifies a controller ID or name. You cannot use the **-controller** parameter with any other parameter in this command. Issuing the **lsfabric** command with the **-controller** parameter is equivalent to issuing the **lsfabric wwpn** *wwpn* command for every configured WWPN of the specified controller. For example, a controller with four ports that are connected to an eight node system with two counterpart SANs produces 64 lines of output.

-cluster *cluster_id_or_name*
(Optional) Specifies a system ID or name. You cannot use the **-cluster** parameter with any other parameter. Issuing the **lsfabric** command with the **-cluster** parameter is equivalent to issuing the **lsfabric wwpn** *wwpn* command for every known WWPN in the specified system. Output is sorted by remote WWPNs and then system WWPNs. This parameter can be used to check the state of connections within the local system or between the local and remote system. When the local system ID or name is specified, each node-to-node connection is listed twice: once from each end. For example, an eight-node system with two counterpart SANs produces eight nodes, which are multiplied by seven other nodes, which are 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.

Remember: The value of the `local_port` field is the number of the node's Fibre Channel (FC) IO port.

This table provides the attribute values that can be displayed as output view data.

Table 26. lsfabric output

Attribute	Description
remote_wwpn	Indicates the ports that have a login to the specified WWPN
remote_nportid	Indicates the port identifier on the remote.
id	Indicates the node ID on which the login is established
node_name	Indicates the node name on which the login is established
local_wwpn	Indicates the ports that have a login to the specified WWPN
local_port	Indicates a concise view of all NQNs that are logged into the specified port ID and node
local_nportid	Indicates the port identifier on the local.

Table 26. lsfabric output (continued)	
Attribute	**Description**
state	Indicates whether a login is `active`, `inactive` or `blocked`. **state active** The meaning of this value depends on the object that it applies to, as follows: **host or controller** Small Computer System Interface (SCSI) commands were issued within the last 5 minutes. **node** Node ports can see other ports. **state inactive** No transactions completed within the last 5 minutes. **Note:** It can take up to 10 seconds after a command for a controller port to change from inactive to active. It can take up to 5 minutes after a command for a host port to change from inactive to active. **state blocked** This value shows connections that are `blocked` due to the system's port mask settings. The meaning of this value depends on either of the following: 1. Node ports that are blocked due to the system's port mask settings. 2. Logins are blocked due to violating portset membership between the host and the Fibre Channel IO port where host has logged in. The Fibre Channel IO port is not added to the portset which the host is a part of. Host IO will not be allowed from this login path for the host.
name	Indicates the name of the FC connectivity between nodes, controllers, and hosts.

Remember: The **lsfabric** command is limited to displaying 16,384 entries. If you have a large system configuration that exceeds these limits you must filter the output (for example, by node or node port) to view all fabric login records.

lsfabricport

Use the **lsfabricport** command to display the FDMI information that is discovered by the system.

Syntax



Parameters

-callhome
(Required) Displays all the FDMI fields for the port.

Description

To view the FDMI information, run the **lsfabricport** command. FDMI discovery runs automatically every 6 hours and collects information about the fabric.

This table shows the possible outputs:

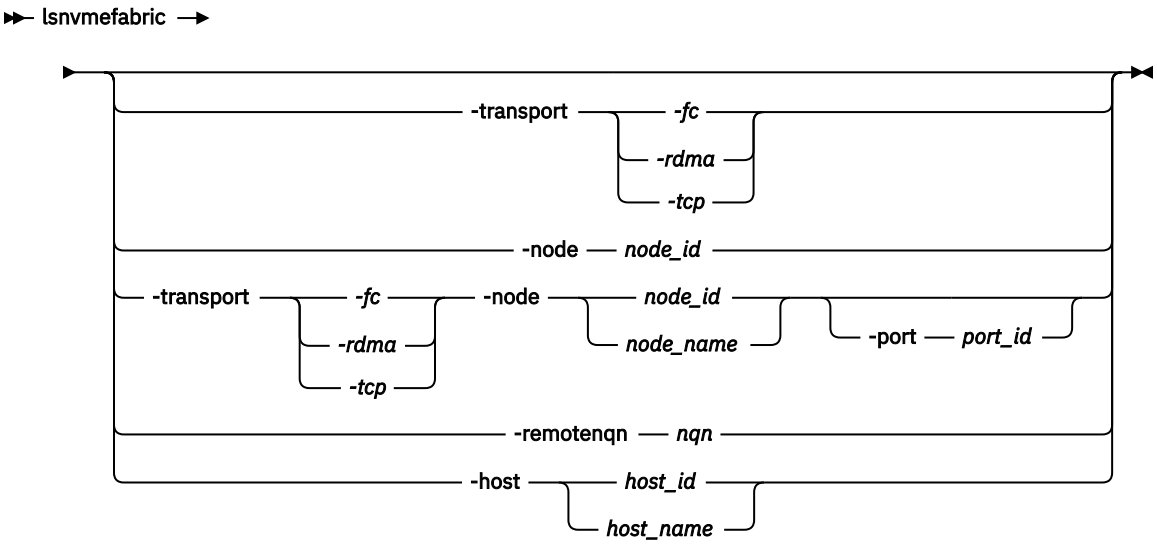
*Table 27. **lsfabricport** output*	
Attribute	**Description**
fc_WWPN	Indicates the port name.
fdmi_host_name	Indicates the host name that is registered by the port.
nportid	Indicates the port identifier on the switch.
fabric_name	Indicates the name of the port fabric.
speed	Indicates the current port speed in Gb.
supported_speed	Indicates the multiple supported speed by the port in Gb.
manufacturer	Indicates the manufacturer of the port or HBA.
serial_number	Indicates the serial number of HBA for the port.
model	Indicates the model of HBA for the port.
model_desc	Indicates the model description of HBA for the port.
firmware_version	Indicates the firmware version of HBA for the port.
driver_version	Indicates the driver version of HBA for the port.
os	Indicates the name and version of the operating software.
symbolic_name	Indicates the symbolic name of the node.
host_name	Indicates the name of the host object if WWPN is configured in a host object.
logged_in	Indicates whether the WWPN is logged in to any node in IBM Storage Virtualize system.

Note: If an attribute listed in Table 27 on page 136 is not registered with the switch, its value is displayed as blank.

lsnvmefabric

Use the **lsnvmefabric** command to display information about NVMe logins from configured hosts.

Syntax



Parameters

-transport *fc | rdma | tcp*
(Optional) Displays the output of all logins with the specified NVMe transport.

-node *node_name | node_id*
(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 NQNs that are logged into the specified port ID and node. The **-port** parameter must be specified with only the **-node** parameter. A valid *port_id* value is from a minimum of 1 through a maximum equal to the number of node Fibre Channel (FC) ports.

-remotenqn *nqn*
(Optional) Displays a list of all logins from or to the specified remote NQN. You cannot use the **-remotenqn** parameter with any other parameter.

-host *host_id | host_name*
(Optional) Specifies a host name or ID. Issuing the **lsnvmefabric** command with the **-host** parameter is equivalent to issuing the **lsnvmefabric** command for every configured NQN of the specified host. You cannot use the **-host** parameter with any other parameter.

Description

The **lsnvmefabric** command displays NVMe logins from configured hosts.

This table provides the attribute values that can be displayed as output view data.

*Table 28. **lsnvmefabric** output*

Attribute	Description
remote_nqn	Indicates the remote NQN of the login.
node_id	Indicates the node ID on which the login is established.
node_name	Indicates the node name on which the login is established.

Table 28. **lsnvmefabric** output (continued)	
Attribute	**Description**
port_id	The node port ID that indicates the NVMe port that is displayed by the **lstargetportfc** view with the same **port_id** field.
state	Indicates whether a login is **active** or **inactive**. • active - NVMe commands were issued within the last 5 minutes. • inactive - No transactions completed within the last 5 minutes. **Note:** It might take up to 5 minutes after a command is issued for a host port to change from **inactive** to **active**.
object_type	The type of object to which the login belongs to.
object_name	The name of the object to which the login belongs to.
transport	Indicates the output of all logins with the specified NVMe transport.
port	Indicates a concise view of all NQNs that are logged into the specified port ID and node. The **-port** parameter must be specified either **-node** or **-transport**. A valid **port_id** value is in the range of 1 to the number of node transport ports. **Note:** • If the **-node** and **-transport** parameters are specified then it will show all the logins specific to the node and transport. • If **-node** and **-transport** parameters are not specified then it will show all the logins.

lsfcportcandidate

Use the **lsfcportcandidate** command to list the Fibre Channel (FC) ports. This information is used to find open FC ports.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

Description

This command returns a list of unconfigured, logged in FC ports.

Note: The **lsfcportcandidate** command presents a list of host FC ports that are logged in to nodes. However, there are situations when the information that is presented might include host FC ports that are no longer logged in or even part of the SAN fabric. For example, if a host FC port is unplugged from a switch but **lsfcportcandidate** shows the worldwide port name (WWPN) that is logged in to all nodes, the incorrect entry is removed when another device is plugged in to the same switch port that previously contained the removed host FC port.

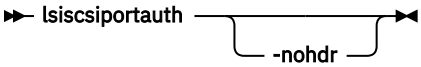
Table 29 on page 139 shows the possible output:

Table 29. lsfcportcandidate output	
Attribute	**Description**
fc_WWPN	Indicates that the FC WWPN is logged in but unconfigured (not assigned to a host). This value must be 16 hexadecimal characters.

lsiscsiportauth

Use the **lsiscsiportauth** command to display the per initiator port authentication and authorization information that is configured.

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.

Description

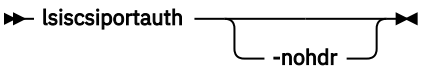
This table shows the possible output:

Table 30. ***lsiscsiportauth*** *output*	
Attribute	**Description**
io_grp	Displays the I/O group of the node. This value is a number in the range of 0 to 4.
position	Displays the position of the node in the I/O group. This value is either 0 or 1.
src_port_id	Displays the local Ethernet port ID of the node. This value is a number in the range of 1 to 16.
iqn	Displays the iSCSI Qualified Name (IQN) to be mandated by IBM Cloud for connecting to the storage that uses this initiator port. This value is an alphanumeric with a maximum of 256 characters.
username	Displays the username to be mandated by IBM Cloud for connecting to the storage that uses this initiator port. This value is an alphanumeric with a maximum of 32 characters.

lsiscsiportauth

Use the **lsiscsiportauth** command to display the per initiator port authentication and authorization information that is configured.

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.

Description

This table shows the possible output:

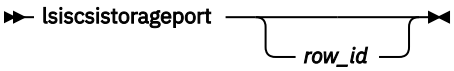
*Table 31. **lsiscsiportauth** output*

Attribute	Description
io_grp	Displays the I/O group of the node. This value is a number in the range of 0 to 4.
position	Displays the position of the node in the I/O group. This value is either 0 or 1.
src_port_id	Displays the local Ethernet port ID of the node. This value is a number in the range of 1 to 16.
iqn	Displays the iSCSI Qualified Name (IQN) to be mandated by IBM Cloud for connecting to the storage that uses this initiator port. This value is an alphanumeric with a maximum of 256 characters.
username	Displays the username to be mandated by IBM Cloud for connecting to the storage that uses this initiator port. This value is an alphanumeric with a maximum of 32 characters.

lsiscsistorageport

Use the **lsiscsistorageport** command to display the details of the iSCSI login sessions that are established from the initiator's iSCSI source ports to the iSCSI backend target controller ports.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

row_id
(Optional) Specifies the row ID view. The ID denotes sessions that are established from the specified initiator node ports to a backend controller target iSCSI qualified name (IQN) through a target Internet Protocol (IP) address. The value must be a number 0 - 1024.

Description

This command displays the details of sessions that are established from the initiator's iSCSI source ports to iSCSI backend target controller ports.

This table provides the attribute values that can be displayed as output view data.

Table 32. lsiscsistorageport output

Attribute	Description
id	Indicates the object identifier for any sessions from any system nodes to the iSCSI backend controller iSCSI qualified name (IQN) through an iSCSI backend controller target IP. The value must be a number 0 - 1023.
src_port_id	Indicates the source port identifier for the node Ethernet port number that is displayed in the **lsportethernet** output. The value is a number 1 - 8.
target_ipv4	Indicates the IPv4 address of the iSCSI backend controller target port that establishes a session from the initiator source port that is identified by the source port ID. The default value is blank.
target_ipv6	Indicates the IPv6 address of the iSCSI backend controller target port that establishes a session from the initiator source port that is identified by the source port ID. The default value is blank.
target_iscsiname	Indicates the IQN of the iSCSI backend controller target that establishes a session. The value must be an alphanumeric string of no more than 256 characters. The default value is blank.
controller_id	Indicates the controller ID that is displayed in the **lscontroller** output. The value must be a number 0 - 1023. The default value is 1024.
iogroup_list	Indicates a colon-separated list of discovery result codes: The value must be an alphanumeric string of up to 32 characters. This field cannot be blank. The values are 0 and 1. • 0 indicates that the I/O group is available in the system, but discovery is either not triggered through the I/O group or discovery through the I/O group fails. • 1 indicates that the I/O group is present and discovery is successful through the I/O group. **Note:** The value - (dash) indicates that the I/O group is not valid or is not present in the system.

Table 32. lsiscsistorageport output (continued)	
Attribute	**Description**
status	Indicates the connectivity status from all nodes in the system to the target port. • `full` If you specify a single I/O group by using the **addiscsistorageport** command and you establish the session from all nodes in the specified I/O group, the status is `full`. • `partial` If you specify a single I/O group by using the **addiscsistorageport** command and you establish the session from a single node in the specified I/O group, the status is `partial`. • `none` If you specify a single I/O group by using the **addiscsistorageport** command and you do not establish the session from any node in the specified I/O group, the status is `none`. No default value exists. This field cannot be blank.
connected	Indicates whether the established connection is from a specified Ethernet port of a target IQN and IP address. The values are `yes` and `no`.
node_id	Indicates the node ID of the initiator node that establishes the session. The value must be a numeric value.
node_name	Indicates the node name of the initiator node that establishes the session. The value must be an alphanumeric string of no more than 16 characters. The default value is blank.
src_ipv4	Indicates the IPv4 IP address of the source port ID on a specified node. The default value is blank.
src_ipv6	Indicates the IPv6 IP address of the source port ID on a specified node. The default value is blank.
src_iscsiname	Indicates the IQN of the source node for which connectivity is being displayed for the target port. The value must be an alphanumeric string of no more than 256 characters. The default value is blank.

Before you specify **lsiscsistorageport** such as in the following examples, you must use this procedure.

1. Complete target discovery by using an IPv4 IP address of a target ISCSI controller through source port ID `0`.

   ```
   detectiscsistorageportcandidate -targetip IP1 –srcportid 2
   ```

2. You would then specify session establishment by using **addiscsistorageport** for discovery output row `0` through I/O group `1`.

   ```
   addiscsistorageport -iogrp 1 0
   ```

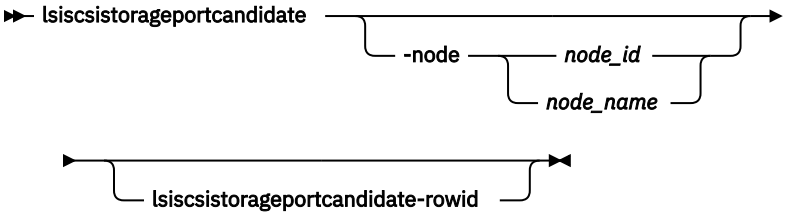
3. Specify **lsiscsistorageport** to view the output (no *tgt_user_name* or *target_chap* is required for discovery or session establishment).

Specify **rmiscsistorageport** to remove a session.

lsiscsistorageportcandidate

Use the **lsiscsistorageportcandidate** command to display a concise or detailed list of information about the candidate Internet Small Computer Systems Interface (iSCSI) target controller iSCSI Qualified Name (IQN) that is specified with the target IP from the specified initiator source ports.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-node *node_name | node_id*
(Optional) Specifies the ID or name of a node in the system. The value must be an alphanumeric string.

lsiscsistorageportcandidate-rowid
(Optional) Specifies a row ID shown in the concise view to provide a detailed view of information about discovered portal IPs.

Description

This command lists information about the last invocation of the **detectiscsistorageportcandidate** command. This command can also display two types of lists:

- A concise list of information about the candidate iSCSI target controller IQNs that are visible at the specified target IP from the specified initiator ports along with indication of which initiator ports can see each discovered iSCSI target IQN.
- The detailed information about all the Target Controller Portal IPs discovered during the last invocation of the **detectiscsistorageportcandidate** command.

This table provides the attribute values that can be displayed as output view data.

Table 33. ***lsiscsistorageportcandidate*** *output*

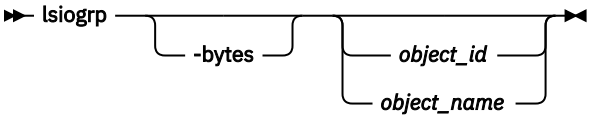
Attribute	Description
id	Indicates the row ID for the discovery output. Enter the **detectiscsistorageportcandidate** command before you use the **lsiscsistorageportcandidate** command. Display the concise view first to show one row per IQN. Use a row ID from the concise view to specify the detailed view of the **lsiscsistorageportcandidate** *lsiscsistorageportcandidate-rowid* command. The detailed view displays the list of discovered target IP addresses for the IQN.
src_port_id	Indicates the source port ID. The value is a number 1 - 8.
target_ipv4	Indicates the target IPv4 address.
target_ipv6	Indicates the target IPv6 address.

*Table 33. **lsiscsistorageportcandidate** output (continued)*	
Attribute	**Description**
target_iscsiname	Indicates the (discovered) IQN that uses the target controller. The value is an alphanumeric string that is 256 characters long.
iogroup_list	Indicates a colon-separated list of discovery result codes: The value must be an alphanumeric string of up to 32 characters. This field cannot be blank. The values are 0 and 1: • 0 indicates that the I/O group is available in the system, but discovery is either not triggered through the I/O group or discovery through the I/O group fails. • 1 indicates that the I/O group is present and discovery is successful through the I/O group. **Note:** The value - (dash) indicates that the I/O group is not valid or is not present in the system.
status	Indicates whether discovery was successful. The status is one of the following values: • full • partial • none
configured	Indicates whether the discovered target IQN has any established sessions with source ports or target ports. The values are yes and no (default).
SendTargets Portal IPs	Provides list of all the (Maximum 32) discovered Target Portal IPs by using the **detectiscsistorageportcandidate** command.

lsiogrp

Use the **lsiogrp** command to display a concise list or a detailed view of input/output (I/O) groups visible to the system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-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 that match the filtering requirements that are specified by the **-filtervalue** parameter are displayed.

Description

This command returns a concise list or a detailed view of I/O groups visible to the system.

The following attributes are displayed:

id
Indicates the ID of the I/O group.

name
Indicates the name of the I/O group.

node_count
Indicates the number of nodes in the I/O group.

vdisk_count
Indicates the number of volumes in the I/O group.

host_count
Indicates the number of hosts in the I/O group.

flash_copy_total_memory
Indicates the total amount of memory that is allocated to FlashCopy.

flash_copy_free_memory
Indicates the total amount of memory that is allocated to FlashCopy, but unused.

remote_copy_total_memory (Deprecated)

remote_copy_free_memory (Deprecated)

mirroring_total_memory
Indicates the total amount of memory that is allocated to mirroring.

raid_total_memory
Indicates the total amount of memory that is allocated to RAID.

raid_free_memory
Indicates the total amount of memory that is allocated to RAID, but is unused.

maintenance
Indicates whether the I/O group is in maintenance mode.

compression_active
Indicates whether Real-time compression is used in the selected I/O group.

accessible_vdisk_count
The number of accessible volumes in this I/O group.

compression_supported
Indicates whether the I/O group supports Real-time compression or data reduction compression.

max_enclosures
Indicates the maximum number of enclosures that are supported by this I/O group.

encryption_supported
Indicates whether the I/O group supports encryption for attached drives. The values are:

- yes
- no

flash_copy_maximum_memory
Identifies the maximum memory that can be set for use by FlashCopy in the specified I/O group.

site_id
Identifies the site ID for the I/O group. The possible values are:

- 1
- 2

site_name
Identifies the site name for the I/O group. The value must be an alphanumeric string or blank.

fctargetportmode
Indicates the current N_Port ID Virtualization (NPIV) status in the specified I/O group. The values are:

- disabled
- transitional
- enabled
- changing_disabled_to_transitional
- changing_transitional_to_disabled
- changing_enabled_to_transitional
- changing_transitional_to_enabled

compression_total_memory
Indicates the total amount of memory that is allocated for Real-time compression per node in the specified I/O group.

deduplication_supported
Indicates whether this I/O group supports data deduplication. The value that is displayed is either yes or no.

deduplication_active
Indicates whether data deduplication is in use in the I/O group. The value that is displayed is either yes or no.

nqn
The NVMe Qualified Name (NQN) presented by nodes in this I/O group. This field is blank if NVMe host attachment is not supported.

flash_copy_max_legacy_memory
The maximum amount of bitmap memory available that can be used by FlashCopy legacy mappings.

flash_copy_used_legacy_memory
The amount of bitmap memory currently used by FlashCopy legacy mapping.

remote_copy_maximum_memory (Deprecated)

mirroring_maximum_memory
The maximum amount of bitmap memory usable by Mirroring for this I/O group.

raid_maximum_memory
The maximum amount of bitmap memory usable by RAID for this I/O group.

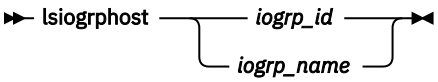
lshbaportcandidate (Deprecated)

The **lshbaportcandidate** command is deprecated. Use either the **lsfcportcandidate** or **lssasportcandidate** command instead.

lsiogrphost

Use the **lsiogrphost** command to display a list of the hosts mapped to a specified I/O group.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

iogrp_id | iogrp_name
(Required) The ID or name of the I/O group for which a list of all mapped hosts is required.

Description

The **lsiogrphost** command displays a list of hosts that are mapped to a specified I/O group.

lsiogrpcandidate

Use the **lsiogrpcandidate** command to list the I/O groups that can have nodes added to them.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

Description

The **lsiogroupcandidate** command displays a list of I/O groups to which nodes can be added. Only the I/O group IDs are displayed.

lsiostatsdumps (Deprecated)

Attention: The **lsiostatsdumps** command is deprecated. Use the **lsdumps** command to display a list of files in a particular dumps directory.

lsiotracedumps (Deprecated)

Attention: The **lsiotracedumps** command is deprecated. Use the **lsdumps** command to display a list of files in a particular dumps directory.

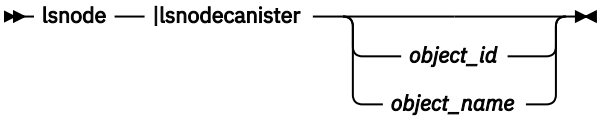
lsnode

Use the **lsnode**/ **lsnodecanister** command to return a concise list or a detailed view of nodes or node canisters that are part of the system .

The list report style can be used to obtain two styles of report:

- A list containing concise information about all the nodes or node canister on a system. Each entry in the list corresponds to a single node or node canister.
- The detailed information about a single node or node canister.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

object_id | object_name
(Optional) Specifies the object ID or name. When you use this parameter, the detailed view of the specific object is returned and any value that is specified by the **-filtervalue** parameter is ignored. If you do not specify the *object_id* | *object_name* parameter, the concise view of all objects that match the filtering requirements that are specified by the **-filtervalue** parameter are displayed.

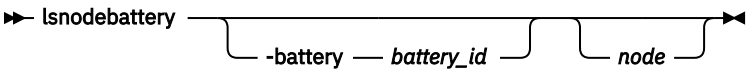
Description

The first four Fibre Channel (FC) input/output (I/O) ports display the worldwide port name (WWPN), state, and speed. If there are less than four FC I/O ports, the fields display with a WWPN of `0000000000000000`, port_status of `inactive`, and port_speed of `N/A`. To examine the FC ports, use the **lsportfc** command.

lsnodebattery

Use the **lsnodebattery** command to display information about the batteries in a node.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-battery ***battery_id***
(Optional) Specifies the battery ID. If you specify this parameter, you must also specify *node*.

node
(Optional) Specifies the node ID or name.

Description

The command displays information about the batteries in a node. The concise view displays a line for each battery slot in all nodes.

This table provides the possible values that are applicable to the attributes that are displayed as data in the output views.

Table 34. ***lsnodebattery*** *attribute values*

Attribute	Value
node_id	Identifies the node that contains the battery.
node_name	Identifies the node that contains the battery.
battery_id	Identifies the battery in the node.
status	Identifies the status of the battery: • `online` if the battery is present and working as usual (which includes a battery calibration). • `degraded` indicates that the battery is present but not working as usual (it has an error logged against it). • `offline` indicates that the battery cannot be detected or is failed (a node error indicating it is missing or failed is logged against the battery). **Remember:** If the status is `offline`, all other fields display the most recent battery data. If no data was shown, all fields remain blank.

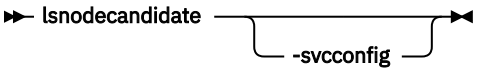
*Table 34. **lsnodebattery** attribute values (continued)*	
Attribute	**Value**
charging_status	Identifies the charging state of the battery: • `charged` indicates that the battery is fully charged. • `charging` indicates that the battery is charging. • `discharging` indicates that the battery is losing voltage (life) and it is recalibrating the gas gauge after the battery recharges. • `idle` indicates that the battery is not charging and it is not discharging, but it is not fully charged. • `reconditioning` indicates that the battery is reconditioning itself by discharging and after recharging. **Important:** Reconditioning occurs approximately every three months (on redundant systems) and can take from 12 - 48 hours.
recondition_needed	Identifies that the battery needs to be reconditioned or it must start reconditioning soon. A recalibration is required because the reported values from the gas gauge are not trusted. The values are `yes` and `no`.
node_percentage_charge	Identifies the battery's percentage charged to determine if it can support the node.
end_of_life_warning	Identifies the battery's end of life (with a warning noise). The values are `yes` and `no`. **Important:** Replace the battery.
present	Identifies a battery is present. The values are `yes` and `no`.
redundant	Identifies if it is safe to remove the battery. The values are `yes` and `no`.
remaining_charge_capacity_mAh	Identifies the remaining capacity of the battery (when the battery recharges) in milliamps hours (mAh).
full_charge_capacity_mAh	Identifies the fully charged capacity of the battery in mAh.
FRU_part_number	Identifies the FRU part number of the battery. The value contains seven alphanumeric characters.
FRU_identity	Identifies the 11S number (combining the manufacturing part number and the serial number). The value contains 22 alphanumeric characters.
compatibility_level	Identifies the battery driver software must support the current software level to operate with this battery (this comes from the battery vital product data or VPD).
last_recondition_timestamp	Indicates a system timestamp for the last successful recalibration of the gas gauge. The format is `YYMMDDHHMMSS`, where: • `YY` indicates year. • The first `MM` indicates month. • `DD` indicates day. • `HH` indicates hour. • The second `MM` indicates minute. • `SS` indicates second.
powered_on_hours	Indicates the number of hours the battery is in a powered node.

Table 34. **lsnodebattery** attribute values (continued)	
Attribute	**Value**
cycle_count	Identifies the number of charge or discharge cycles the battery performs.
error_sequence_number	Identifies the error log number of the highest priority error. This field is generally blank, but if the status is `degraded` or `offline` an error sequence number is displayed.

lsnodecandidate (SAN Volume Controller)

Use the **lsnodecandidate** command to list all of the nodes that are available to add to the clustered system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-svcconfig
(Optional) Lists all nodes in the enclosure that are in a candidate state.

Description

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

This command displays a list of nodes that are available to add to the clustered system. This list includes nodes that are not already part of a clustered system, but are compatible with the clustered system code level. Nodes with hardware types that are incompatible with the installed code are not listed.

This table describes the possible outputs:

Table 35. lsnodecandidate outputs	
Attribute	**Description**
panel_name	Unique identifier for the node.
UPS_serial_number	The serial number of the UPS.
UPS_unique_id	The unique ID of UPS.
hardware	Describes the type of nodes.

lsnodecanister

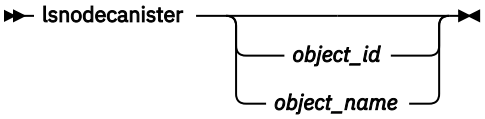
The **lsnodecanister** command returns a concise list or a detailed view of node canisters that are part of the system.

The list report style can be used to obtain two styles of report:

- A list that contains concise information about all the node canisters on a system. Each entry in the list corresponds to a single node canister.

- The detailed information about a single node canister.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

object_id | object_name
(Optional) Specifies the object ID or name. When you use this parameter, the detailed view of the specific object is returned and any value that is specified by the **-filtervalue** parameter is ignored. If you do not specify the *object_id* | *object_name* parameter, the concise view of all objects that match the filtering requirements that are specified by the **-filtervalue** parameter are displayed.

Description

This command returns a concise list or a detailed view of nodes that are part of the system. This table provides the possible values that are applicable to the attributes that are displayed as data in the output views.

Table 36. ***lsnodecanister*** *attribute values*

Attribute	Value
status	Indicates the status. The values are: • offline • service • flushing • pending • online • adding • deleting
config_node	Indicates if the node is a configuration node. The values are: • yes • no
IO_group_id	Indicates the I/O group of the node.
IO_group_name	Indicates the name of the I/O group in the node.

*Table 36. **lsnodecanister** attribute values (continued)*	
Attribute	**Value**
port_status	Indicates the port status. The values are: • active • inactive • not_installed
hardware	Indicates the node canister hardware type, for example 300.
UPS_serial_number	Blank
UPS_unique_id	Blank
panel_name	Indicates a compound of the enclosure ID and the canister ID, separated with a hyphen. On a system this value is the same as the serial number on the front panel.
enclosure_id	Indicates the enclosure that the node is in. This field is blank for a node-based system.
canister_id	Indicates the canister location within the enclosure. This field is blank for a node-based system.
enclosure_serial_number	Indicates the serial number of the enclosure. This field is blank for a node-based system.
service_IP_mode	Indicates the current mode of the service IPv4 • Empty if IPv4 is not active • The values are: – static (if the service IP is set by the user) – dhcp (if the service IP is set successfully by using DHCP server) – dhcpfallback (if the service IP is set to a default value after a DHCP server request failed)
service_IP_mode_6	Indicates the current mode of the service IPv6 • Empty if IPv6 is not active • Either static (if the service IP is set by the user) or dhcp (if the service IP set successfully by using DHCP server).
identify_LED	Indicates the node or node canister state - on, off, or blank.
code_level	Indicates the current level of machine code that is running on the node canister. on, off, or blank.
product_mtm	Indicates the product machine type.
update_complete	Indicates whether the node canister completes its update. The values are yes or no.
site_id	Indicates the site node canister value.
site_name	Indicates the site name.
nvme_name	Indicates the system NVMe Qualified Name (NQN).

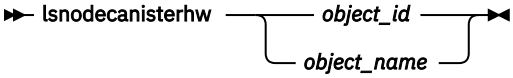
Note: Use the **lsenclosure** command to display information about the serial_number or machine_signature.

The first four Fibre Channel (FC) input/output (I/O) ports display the worldwide port name (WWPN), state, and speed. If there are less than four FC I/O ports, the fields display with a WWPN of 0000000000000000, port_status of `inactive`, and port_speed of N/A. To examine the FC ports, use the **lsportfc** command.

lsnodecanisterhw

Use the **lsnodecanisterhw** command to display the configured and actual hardware configuration of node canisters in the clustered system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

object_id | object_name
(Required) Specifies the object name or ID.

Description

This table provides the possible values that are applicable to the attributes that are displayed as data in the output views.

*Table 37. **lsnodecanisterhw** attribute values*

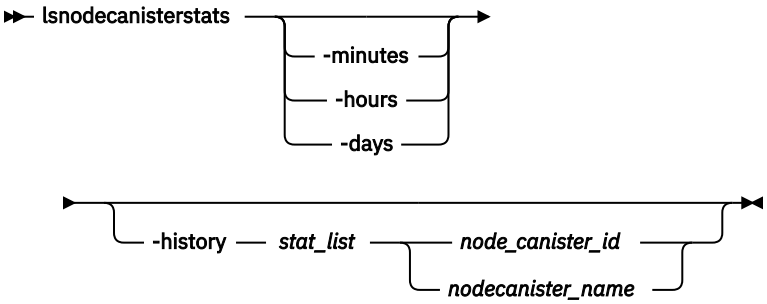
Attribute	Value
id	Indicates the node canister unique ID.
name	Indicates the node canister name.
status	Indicates the status. The values are: • offline • service • flushing • pending • online • adding • deleting
IO_group_id	The I/O group ID.
IO_group_name	Indicates the I/O group name.
hardware	Indicates the node canister hardware type, such as 300.
actual_different	Indicates whether the node canister hardware is different from the configured hardware.
actual_valid	Indicates whether the node canister hardware is valid.

Table 37. **lsnodecanisterhw** attribute values (continued)	
Attribute	**Value**
memory_configured	Indicates the configured amount of memory (in GB).
memory_actual	Indicates the currently installed amount of memory (in GB).
memory_valid	Indicates whether the actual memory is a valid configuration.
cpu_count	Indicates the maximum number of CPUs for the node canister.
cpu_socket	Indicates the ID of socket the CPU fields refer to.
cpu_configured	Indicates the configured CPU for this socket.
cpu_actual	Indicates the currently installed CPU in this socket.
cpu_valid	Indicates whether the currently installed CPU is a valid configuration.
adapter_count	Indicates the maximum number of adapters for the node canister (differs by node canister type).
adapter_location	Indicates the location of this adapter.
adapter_configured	Indicates the configured adapter for this location.
adapter_actual	Indicates the currently installed adapter for this location.
adapter_valid	Indicates whether the adapter in this location is valid.
ports_different	Indicates whether the current hardware is able to provide more I/O ports. The values are yes and no.

lsnodecanisterstats

The **lsnodecanisterstats** command displays the most recent values of all node canister statistics, and displays all node canister statistics for a particular canister. This command also can be used to display a history of values for a subset of available statistics.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-minutes
(Optional) Displays stats at minutes granularity.

-hours
(Optional) Displays stats at hours granularity.

-days
(Optional) Show stats at days granularity.

Note:

- Flags are mutually exclusive. In absence of any the live (5 sec) samples are displayed and used with or without **-history** parameter.
- Each granularity will show up to 60 samples per stat when used with **-history** parameter.

-history stat_list
(Optional) Provides the most recent node canister statistical values, specific canister statistical values, or historical data for any canister.

nodecanister_id | nodecanister_name
(Optional) Identifies the node canister for which you want to request statistics.

Description

This command returns a concise list or a detailed view of node canisters that are part of the clustered system. This table provides the possible values that are applicable to the attributes that are displayed as data in the output views.

Table 38. Node canister attribute values

Attribute	Value
nodecanister_id	The ID of the node canister.
nodecanister_name	The name of the node canister.
stat_current	The current value of the statistic field. The value is accurate up to 3 decimal places (microseconds).
stat_list	The system history of the reported statistics.
stat_name	The name of the statistic field.
stat_peak	The peak value of the statistic field in the last five minutes. The value is accurate up to 3 decimal places (microseconds).
stat_peak_time	The time that the peak occurred.
sample_time	The time of the sample occurrence.
stat_value	The statistical value at the epoch interval.

Note: Filtering is supported on `nodecanister_id`, `nodecanister_name`, and `stat_name` fields by using the concise view.

The following table provides the possible values that are applicable to the values that are displayed for the **stat_name** attribute.

*Table 39. **stat_name** field values*

Value	Description
cpu_pc	The total percentage of CPU utilization for the system.

*Table 39. **Stat_name** field values (continued)*	
Value	**Description**
fc_mb	Displays the total number of megabytes transferred per second (MBps) for Fibre Channel traffic on the system. This value includes host I/O and any bandwidth that is used for communication within the system.
fc_io	Displays the total input/output (I/O) operations that are transferred per seconds for Fibre Channel traffic on the system. This value includes host I/O and any bandwidth that is used for communication within the system.
sas_mb	Displays the total number of megabytes transferred per second (MBps) for serial-attached SCSI (SAS) traffic on the system. This value includes host I/O and bandwidth that is used for background RAID activity.
sas_io	Displays the total I/O operations that are transferred per second for SAS traffic on the system. This value includes host I/O and bandwidth that is used for background RAID activity.
iscsi_mb	Displays the total number of megabytes transferred per second (MBps) for iSCSI traffic on the system.
iscsi_io	Displays the total I/O operations that are transferred per second for iSCSI traffic on the system.
write_cache_pc	Displays the percentage of the write cache usage for the node.
total_cache_pc	Displays the total percentage for both the write and read cache usage for the node.
vdisk_mb	Displays the average number of megabytes transferred per second (MBps) for read and write operations to volumes during the sample period.
vdisk_io	Displays the average amount of I/O operations that are transferred per second for read and write operations to volumes during the sample period.
vdisk_ms	Displays the average amount of time in microseconds that the system takes to respond to read and write requests to volumes over the sample period.
mdisk_mb	Displays the average number of megabytes transferred per second (MBps) for read and write operations to MDisks during the sample period.
mdisk_io	Displays the average amount of I/O operations that are transferred per second for read and write operations to MDisks during the sample period.
mdisk_ms	Displays the average amount of time in microseconds that the system takes to respond to read and write requests to MDisks over the sample period.
drive_mb	Displays the average number of megabytes transferred per second (MBps) for read and write operations to drives during the sample period.
drive_io	Displays the average amount of I/O operations that are transferred per second for read and write operations to drives during the sample period.
drive_ms	Displays the average amount of time in microseconds that the system takes to respond to read and write requests to drives over the sample period.
vdisk_w_mb	Displays the average number of megabytes transferred per second (MBps) for read and write operations to volumes during the sample period.
vdisk_w_io	Displays the average amount of I/O operations that are transferred per second for write operations to volumes during the sample period.
vdisk_w_ms	Displays the average amount of time in microseconds that the system takes to respond to write requests to volumes over the sample period.

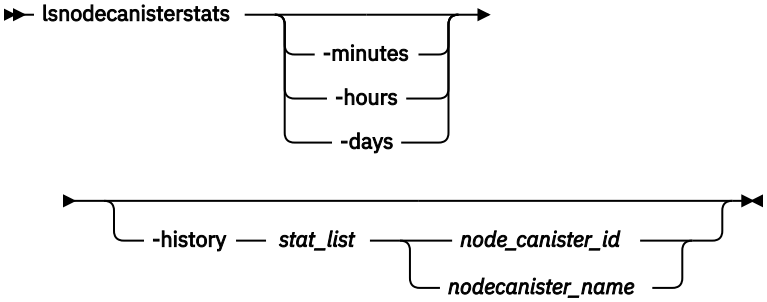
*Table 39. **Stat_name** field values (continued)*	
Value	**Description**
mdisk_w_mb	Displays the average number of megabytes transferred per second (MBps) for write operations to MDisks during the sample period.
mdisk_w_io	Displays the average amount of I/O operations that are transferred per second for write operations to MDisks during the sample period.
mdisk_w_ms	Displays the average amount of time in microseconds that the system takes to respond to write requests to MDisks over the sample period.
drive_w_mb	Displays the average number of megabytes transferred per second (MBps) for write operations to drives during the sample period.
drive_w_io	Displays the average amount of I/O operations that are transferred per second for write operations to drives during the sample period.
drive_w_ms	Displays the average amount of time in microseconds that the system takes to respond write requests to drives over the sample period.
vdisk_r_mb	Displays the average number of megabytes transferred per second (MBps) for read operations to volumes during the sample period.
vdisk_r_io	Displays the average amount of I/O operations that are transferred per second for read operations to volumes during the sample period.
vdisk_r_ms	Displays the average amount of time in microseconds that the system takes to respond to read requests to volumes over the sample period.
mdisk_r_mb	Displays the average number of megabytes transferred per second (MBps) for read operations to MDisks during the sample period.
mdisk_r_io	Displays the average amount of I/O operations that are transferred per second for read operations to MDisks during the sample period.
mdisk_r_ms	Displays the average amount of time in microseconds that the system takes to respond to read requests to MDisks over the sample period.
drive_r_mb	Displays the average number of megabytes transferred per second (MBps) for read operations to drives during the sample period
drive_r_io	Displays the average amount of I/O operations that are transferred per second for read operations to drives during the sample period.
drive_r_ms	Displays the average amount of time in microseconds that the system takes to respond to read requests to drives over the sample period.
iplink_mb	The total number of megabytes transferred per second (MBps) for Internet Protocol (IP) replication traffic on the system. This value does not include iSCSI host input/output (I/O) operations.
iplink_comp_mb	Displays the average number of compressed megabytes transferred per second (MBps) over the IP replication link during the sample period.
iplink_io	The total input/output (I/O) operations that are transferred per second for IP partnership traffic on the system. This value does not include Internet Small Computer System Interface (iSCSI) host I/O operations.
nvme_rdma_mb	Displays the total number of megabytes transferred per second (MBps) for NVMe over RDMA traffic on the system.
nvme_rdma_io	Displays the total I/O operations that are transferred per second for NVMe over RDMA traffic on the system.

Table 39. ***Stat_name*** *field values (continued)*	
Value	**Description**
nvme_tcp_mb	Displays the total number of megabytes transferred per second (MBps) for NVMe over TCP traffic on the system.
nvme_tcp_io	Displays the total I/O operations that are transferred per second for NVMe over TCP traffic on the system.

lsnodecanisterstats

The **lsnodecanisterstats** command displays the most recent values of all node canister statistics, and displays all node canister statistics for a particular canister. This command also can be used to display a history of values for a subset of available statistics.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-minutes
(Optional) Displays stats at minutes granularity.

-hours
(Optional) Displays stats at hours granularity.

-days
(Optional) Show stats at days granularity.

Note:

- Flags are mutually exclusive. In absence of any the live (5 sec) samples are displayed and used with or without **-history** parameter.
- Each granularity will show up to 60 samples per stat when used with **-history** parameter.

-history stat_list
(Optional) Provides the most recent node canister statistical values, specific canister statistical values, or historical data for any canister.

nodecanister_id | nodecanister_name
(Optional) Identifies the node canister for which you want to request statistics.

Description

This command returns a concise list or a detailed view of node canisters that are part of the clustered system. This table provides the possible values that are applicable to the attributes that are displayed as data in the output views.

Table 40. Node canister attribute values

Attribute	Value
nodecanister_id	The ID of the node canister.
nodecanister_name	The name of the node canister.
stat_current	The current value of the statistic field. The value is accurate up to 3 decimal places (microseconds).
stat_list	The system history of the reported statistics.
stat_name	The name of the statistic field.
stat_peak	The peak value of the statistic field in the last five minutes. The value is accurate up to 3 decimal places (microseconds).
stat_peak_time	The time that the peak occurred.
sample_time	The time of the sample occurrence.
stat_value	The statistical value at the epoch interval.

Note: Filtering is supported on `nodecanister_id`, `nodecanister_name`, and `stat_name` fields by using the concise view.

The following table provides the possible values that are applicable to the values that are displayed for the **`stat_name`** attribute.

*Table 41. **`Stat_name`** field values*

Value	Description
cpu_pc	The total percentage of CPU utilization for the system.
fc_mb	Displays the total number of megabytes transferred per second (MBps) for Fibre Channel traffic on the system. This value includes host I/O and any bandwidth that is used for communication within the system.
fc_io	Displays the total input/output (I/O) operations that are transferred per seconds for Fibre Channel traffic on the system. This value includes host I/O and any bandwidth that is used for communication within the system.
sas_mb	Displays the total number of megabytes transferred per second (MBps) for serial-attached SCSI (SAS) traffic on the system. This value includes host I/O and bandwidth that is used for background RAID activity.
sas_io	Displays the total I/O operations that are transferred per second for SAS traffic on the system. This value includes host I/O and bandwidth that is used for background RAID activity.
iscsi_mb	Displays the total number of megabytes transferred per second (MBps) for iSCSI traffic on the system.
iscsi_io	Displays the total I/O operations that are transferred per second for iSCSI traffic on the system.

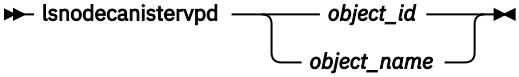
*Table 41. **Stat_name** field values (continued)*	
Value	**Description**
write_cache_pc	Displays the percentage of the write cache usage for the node.
total_cache_pc	Displays the total percentage for both the write and read cache usage for the node.
vdisk_mb	Displays the average number of megabytes transferred per second (MBps) for read and write operations to volumes during the sample period.
vdisk_io	Displays the average amount of I/O operations that are transferred per second for read and write operations to volumes during the sample period.
vdisk_ms	Displays the average amount of time in microseconds that the system takes to respond to read and write requests to volumes over the sample period.
mdisk_mb	Displays the average number of megabytes transferred per second (MBps) for read and write operations to MDisks during the sample period.
mdisk_io	Displays the average amount of I/O operations that are transferred per second for read and write operations to MDisks during the sample period.
mdisk_ms	Displays the average amount of time in microseconds that the system takes to respond to read and write requests to MDisks over the sample period.
drive_mb	Displays the average number of megabytes transferred per second (MBps) for read and write operations to drives during the sample period.
drive_io	Displays the average amount of I/O operations that are transferred per second for read and write operations to drives during the sample period.
drive_ms	Displays the average amount of time in microseconds that the system takes to respond to read and write requests to drives over the sample period.
vdisk_w_mb	Displays the average number of megabytes transferred per second (MBps) for read and write operations to volumes during the sample period.
vdisk_w_io	Displays the average amount of I/O operations that are transferred per second for write operations to volumes during the sample period.
vdisk_w_ms	Displays the average amount of time in microseconds that the system takes to respond to write requests to volumes over the sample period.
mdisk_w_mb	Displays the average number of megabytes transferred per second (MBps) for write operations to MDisks during the sample period.
mdisk_w_io	Displays the average amount of I/O operations that are transferred per second for write operations to MDisks during the sample period.
mdisk_w_ms	Displays the average amount of time in microseconds that the system takes to respond to write requests to MDisks over the sample period.
drive_w_mb	Displays the average number of megabytes transferred per second (MBps) for write operations to drives during the sample period.
drive_w_io	Displays the average amount of I/O operations that are transferred per second for write operations to drives during the sample period.
drive_w_ms	Displays the average amount of time in microseconds that the system takes to respond write requests to drives over the sample period.
vdisk_r_mb	Displays the average number of megabytes transferred per second (MBps) for read operations to volumes during the sample period.

Table 41. ***Stat_name*** *field values (continued)*	
Value	**Description**
vdisk_r_io	Displays the average amount of I/O operations that are transferred per second for read operations to volumes during the sample period.
vdisk_r_ms	Displays the average amount of time in microseconds that the system takes to respond to read requests to volumes over the sample period.
mdisk_r_mb	Displays the average number of megabytes transferred per second (MBps) for read operations to MDisks during the sample period.
mdisk_r_io	Displays the average amount of I/O operations that are transferred per second for read operations to MDisks during the sample period.
mdisk_r_ms	Displays the average amount of time in microseconds that the system takes to respond to read requests to MDisks over the sample period.
drive_r_mb	Displays the average number of megabytes transferred per second (MBps) for read operations to drives during the sample period
drive_r_io	Displays the average amount of I/O operations that are transferred per second for read operations to drives during the sample period.
drive_r_ms	Displays the average amount of time in microseconds that the system takes to respond to read requests to drives over the sample period.
iplink_mb	The total number of megabytes transferred per second (MBps) for Internet Protocol (IP) replication traffic on the system. This value does not include iSCSI host input/output (I/O) operations.
iplink_comp_mb	Displays the average number of compressed megabytes transferred per second (MBps) over the IP replication link during the sample period.
iplink_io	The total input/output (I/O) operations that are transferred per second for IP partnership traffic on the system. This value does not include Internet Small Computer System Interface (iSCSI) host I/O operations.
nvme_rdma_mb	Displays the total number of megabytes transferred per second (MBps) for NVMe over RDMA traffic on the system.
nvme_rdma_io	Displays the total I/O operations that are transferred per second for NVMe over RDMA traffic on the system.
nvme_tcp_mb	Displays the total number of megabytes transferred per second (MBps) for NVMe over TCP traffic on the system.
nvme_tcp_io	Displays the total I/O operations that are transferred per second for NVMe over TCP traffic on the system.

lsnodecanistervpd

The **lsnodecanistervpd** command displays the vital product data (VPD) for each node canister.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

object_id | object_name
(Required) Specifies the object name or ID.

Description

This command displays the VPD for the specified node canister. Each field is reported on a new line. All fields are strings. The VPD is split into sections. Each section has a section heading. The number of fields in that section follows each heading. Each section is separated by an empty line.

For example:

```
section name:3 fields
field1:value
field2:value
field3:value

new section:x fields
(more)
```

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
(more)
```

lsnodedependentvdisks (Deprecated)

Attention: The **lsnodedependentvdisks** command is deprecated. Use the **lsdependentvdisks** command instead.

lsnodehw

Use the **lsnodehw** / **lsnodecanisterhw** command to display the configured and actual hardware configuration of nodes in the clustered system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

object_id | object_name
(Required) Specifies the object name or ID.

Description

This table provides the possible values that are applicable to the attributes that are displayed as data in the output views.

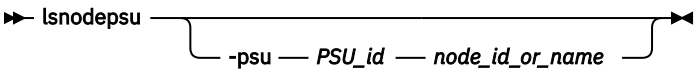
Table 42. Attribute values for ***lsnodehw*** *and* ***lsnodecanisterhw***

Attribute	Value
id	Indicates the node or node canister unique ID.
name	Indicates the node or node canister name.
status	Indicates the node or node canister status.
IO_group_id	Indicates the input/output (I/O) group ID.
IO_group_name	Indicates the I/O group name.
hardware	Indicates the hardware model, such as DH8.
actual_different	Indicates whether the node or node canister hardware is different from the configured hardware.
actual_valid	Indicates whether the node or node canister hardware is valid.
memory_configured	Indicates the configured amount of memory (in GB).
memory_actual	Indicates the currently installed amount of memory (in GB).
memory_valid	Indicates whether the actual memory is a valid configuration.
cpu_count	Indicates the maximum number of CPUs for the node.
cpu_socket	Indicates the ID of socket the CPU fields refer to.
cpu_configured	Indicates the configured CPU for this socket.
cpu_actual	Indicates the currently installed CPU in this socket.
cpu_valid	Indicates whether the currently installed CPU is a valid configuration.
adapter_count	Indicates the maximum number of adapters for the node (differs by node type).
adapter_location	Indicates the location of this adapter.
adapter_configured	Indicates the configured adapter for this location.
adapter_actual	Indicates the currently installed adapter for this location.
adapter_valid	Indicates whether the adapter in this location is valid.
ports_different	Indicates whether the current hardware is able to provide more I/O ports? The values are yes and no.

lsnodepsu

Use the **lsnodepsu** command to display the power supply unit information of nodes in a system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-psu *PSU_id*
(Optional) Specifies the ID of the PSU. You must specify a **node_id** or **node_name** if you use the **-psu** parameter.

node_id | node_name
(Optional) Specifies the node ID or node name (as listed in the **lsnode** command).

Description

This table provides the possible values that are applicable to the attributes that are displayed as data in the output views.

Table 43. Attribute values for ***lsnodepsu***

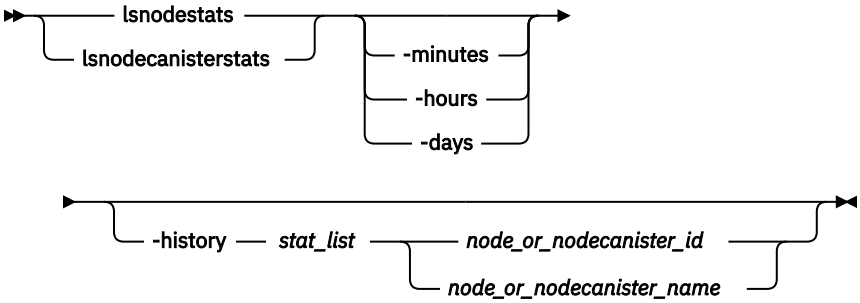
Attribute	Value
node_id	Indicates the ID, as shown in **lsnode**, of the node in which the PSU is located.
node_name	Indicates the name, as shown in **lsnode**, of the node in which the PSU is located.
PSU_id	Indicates the ID of the PSU.
status	Indicates the overall status of the PSU. The possible values are: • offline • online • degraded • unknown
input_failed	Indicates the state of the input failed LED on the PSU. The value is either **on** or **off**.
output_failed	Indicates the state of the output failed LED on the PSU. The value is either **on** or **off**.
fan_failed	Indicates the state of the fan failed LED on the PSU. The value is either **on** or **off**.

Table 43. Attribute values for **lsnodepsu** (continued)	
Attribute	**Value**
swapped	Indicates whether the PSU has been swapped since power-on or the last clear swap. The value is either **yes** or **no**.
redundant	Indicates whether the PSU is redundant. The value is either **yes** or **no**.
FRU_part_number	Indicates FRU part number of the PSU. The value is an ASCII string up to 22 characters.
FRU_identity	Indicates FRU identity of the PSU. The value is an ASCII string up to 22 characters.
firmware_level_1	Indicates the firmware level of the PSU. The value is an ASCII string up to 22 characters.

lsnodestats

Use the **lsnodestats** / **lsnodecanisterstats** command to display the most recent values of statistics for all nodes or node canisters, and display all statistics for a particular node or node canister. Additionally, You can use this command to display a history of values for a given subset of available statistics. This command is only applicable for IBM SAN Volume Controller.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-minutes
(Optional) Displays stats at minutes granularity.

-hours
(Optional) Displays stats at hours granularity.

-days
(Optional) Show stats at days granularity.

Note:

- Flags (**-minutes**, **-hours**, and **-days**) are mutually exclusive. When no options are selected, the real-time (5-second) samples are shown and can be utilized, with or without the inclusion of the **-history** parameter.
- Each granularity will show up to 60 samples per stat when used with **-history** parameter.

-history ***stat_list***
(Optional) Provides a table of statistical values for the specified node. The *stat_list* is a colon-delimited list of one or more statistical values. A table is generated for each entry in the *stat_list*.

Remember: If **-history** is specified, a node ID or name must be specified as well.

node_or_nodecanister_id | node_or_nodecanister_name
(Optional) Identifies the node or node canister for which you want to request statistics.

Description

This command returns a concise list or a detailed view of nodes or node canisters that are part of the clustered system. This table provides the possible values that are applicable to the attributes that are displayed as data in the output views.

Table 44. Attribute values for ***lsnodestats*** *or* ***lsnodecanister***

Attribute	Value
node_id	The ID of the node or node canister.
node_name	The name of the node or node canister.
stat_current	The current value of the statistic field.
stat_list	The system history of the reported statistics. The list of statistics can contain multiple items that are separated by colons.
stat_name	The name of the statistic field. See Table 45 on page 167 for descriptions of available statistics.
stat_peak	The peak value of the statistic field in the last 5 minutes.
stat_peak_time	The time that the peak occurred.
sample_time	The time of the sample occurrence.
stat_value	The statistical value at the epoch interval.

Remember: Filtering is supported on the `stat_name` field by using the concise view.

The following table provides the possible values that are applicable to the values that are displayed for the **stat_name** attribute.

Table 45. ***Stat_name*** *field values*

Value	Description
compression_cpu_pc (Deprecated)	Displays the percentage of allocated CPU capacity that is utilized for compression.
cpu_pc	Displays the percentage of allocated CPU capacity that is utilized for the system.
fc_mb	Displays the total number of megabytes transferred per second for Fibre Channel traffic on the system. This value includes host I/O and any bandwidth that is used for communication within the system.
fc_io	Displays the total input/output (I/O) operations that are transferred per seconds for Fibre Channel traffic on the system. This value includes host I/O and any bandwidth that is used for communication within the system.

*Table 45. **Stat_name** field values (continued)*	
Value	**Description**
sas_mb	Displays the total number of megabytes transferred per second for serial-attached SCSI (SAS) traffic on the system. This value includes host I/O and bandwidth that is used for background RAID activity.
sas_io	Displays the total I/O operations that are transferred per second for SAS traffic on the system. This value includes host I/O and bandwidth that is used for background RAID activity.
iscsi_mb	Displays the total number of megabytes transferred per second for iSCSI traffic on the system.
iscsi_io	Displays the total I/O operations that are transferred per second for iSCSI traffic on the system.
write_cache_pc	Displays the percentage of the write cache usage for the node.
total_cache_pc	Displays the total percentage for both the write and read cache usage for the node.
vdisk_mb	Displays the average number of megabytes transferred per second for read and write operations to volumes during the sample period.
vdisk_io	Displays the average number of I/O operations that are transferred per second for read and write operations to volumes during the sample period.
vdisk_ms	Displays the average amount of time in milliseconds that the system takes to respond to read and write requests to volumes over the sample period.
mdisk_mb	Displays the average number of megabytes transferred per second for read and write operations to MDisks during the sample period.
mdisk_io	Displays the average number of I/O operations that are transferred per second for read and write operations to MDisks during the sample period.
mdisk_ms	Displays the average amount of time in milliseconds that the system takes to respond to read and write requests to MDisks over the sample period.
drive_mb	Displays the average number of megabytes transferred per second for read and write operations to drives during the sample period.
drive_io	Displays the average number of I/O operations that are transferred per second for read and write operations to drives during the sample period.
drive_ms	Displays the average amount of time in milliseconds that the system takes to respond to read and write requests to drives over the sample period.
vdisk_w_mb	Displays the average number of megabytes transferred per second for read and write operations to volumes during the sample period.
vdisk_w_io	Displays the average number of I/O operations that are transferred per second for write operations to volumes during the sample period.
vdisk_w_ms	Displays the average amount of time in milliseconds that the system takes to respond to write requests to volumes over the sample period.
mdisk_w_mb	Displays the average number of megabytes transferred per second for write operations to MDisks during the sample period.
mdisk_w_io	Displays the average number of I/O operations that are transferred per second for write operations to MDisks during the sample period.

*Table 45. **Stat_name** field values (continued)*	
Value	**Description**
mdisk_w_ms	Displays the average amount of time in milliseconds that the system takes to respond to write requests to MDisks over the sample period.
drive_w_mb	Displays the average number of megabytes transferred per second for write operations to drives during the sample period.
drive_w_io	Displays the average number of I/O operations that are transferred per second for write operations to drives during the sample period.
drive_w_ms	Displays the average amount of time in milliseconds that the system takes to respond write requests to drives over the sample period.
power_w	Displays the power that is consumed in watts.
temp_c	Displays the ambient temperature in Celsius.
temp_f	Displays the ambient temperature in Fahrenheit.
vdisk_r_mb	Displays the average number of megabytes transferred per second for read operations to volumes during the sample period.
vdisk_r_io	Displays the average number of I/O operations that are transferred per second for read operations to volumes during the sample period.
vdisk_r_ms	Displays the average amount of time in milliseconds that the system takes to respond to read requests to volumes over the sample period.
mdisk_r_mb	Displays the average number of megabytes transferred per second for read operations to MDisks during the sample period.
mdisk_r_io	Displays the average number of I/O operations that are transferred per second for read operations to MDisks during the sample period.
mdisk_r_ms	Displays the average amount of time in milliseconds that the system takes to respond to read requests to MDisks over the sample period.
drive_r_mb	Displays the average number of megabytes transferred per second for read operations to drives during the sample period
drive_r_io	Displays the average number of I/O operations that are transferred per second for read operations to drives during the sample period.
drive_r_ms	Displays the average amount of time in milliseconds that the system takes to respond to read requests to drives over the sample period.
iplink_mb	The total number of megabytes transferred per second for Internet Protocol (IP) replication traffic on the system. This value does not include iSCSI host input/output (I/O) operations.
iplink_comp_mb	Displays the average number of compressed megabytes transferred per second (MBps) over the IP replication link during the sample period.
iplink_io	The total input/output (I/O) operations that are transferred per second for IP partnership traffic on the system. This value does not include Internet Small Computer System Interface (iSCSI) host I/O operations.
cloud_up_mb	Displays the average number of megabits transferred per second (Mbps) for upload operations to a cloud account during the sample period.
cloud_up_ms	Displays the average amount of time (in milliseconds) it takes for the system to respond to upload requests to a cloud account during the sample period.

Table 45. ***Stat_name*** *field values (continued)*	
Value	**Description**
cloud_down_mb	Displays the average number of Mbps for download operations to a cloud account during the sample period.
cloud_down_ms	Displays the average amount of time (in milliseconds) it takes for the system to respond to download requests to a cloud account during the sample period.

lsnodevpd

Use the **lsnodevpd** / **lsnodecanistervpd** command to display the vital product data (VPD) for each node.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

object_id | object_name
(Required) Specifies the object name or ID.

Description

This command displays the VPD for the specified node or node canister. Each field is reported on a new line. All fields are strings. The VPD is split into sections. Each section has a section heading. The number of fields in that section follows each heading. Each section is separated by an empty line.

For example:

```
section name:3 fields
field1:value
field2:value
field3:value

new section:x fields
...
```

Some sections contain information about multiple objects of that type. Each object within the section is separated by an empty line.

For example:

```
section name:4 fields
object1 field1:value
object1 field2:value

object2 field1:value
object2 field2:value
```

```
new section: x fields
...
```

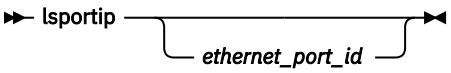
Note: For SAN Volume Controller 2145-8G4 nodes, the VPD displays the device serial number of the Fibre Channel adapter as N/A.

Table 46. Attribute values for ***lsnodevpd*** *and* ***lsnodecanistervpd***	
Value	**Description**
system board	Displays the system board information.
part_number	Displays the total number of system part numbers.
system_serial_number	Displays the total number of system serial numbers.
number_of_processors	Displays the total number of system processors.
number_of_memory_modules	Displays the total number of memory modules.
number_of_fans	Displays the total number of system fans.
number_of_FC_cards	Displays the total number of Fibre Channel (FC) cards.
number_of_Ethernet_cards	Displays the total number of Ethernet cards.
iscsi_initiator_name	Displays the iSCSI IQN that is stored in node vital product data (VPD).

lsportip (Deprecated)

The **lsportip** command has been deprecated and maintained only as stub view for backward compatibility purpose. Instead of **lsportip** command use **lsip** command to view information about configured IP addresses for all type (host, storage or replication) of connectivity and **lsportethernet** command to provide information about physical ethernet ports.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

ethernet_port_id
(Optional) Specifies the ID of an Ethernet port (1, 2, 3, or 4). If omitted, a concise view is displayed for all ports. When you use this parameter, the detailed view of the specified port is returned and any value that is specified by the **-filtervalue** parameter is ignored. If you do not use the *ethernet_port_id* parameter, the concise view displays all ports that match the filtering requirements that are specified by the **-filtervalue** parameter.

Description

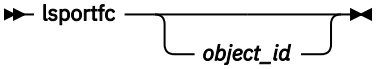
This command lists the configuration of the Ethernet ports for each node in the IBM Storage Virtualize system. This command lists all port IP addresses for each node in the clustered system. The concise view displays two rows of output for each Ethernet port. Each node has two Ethernet ports.

The **lsportip** command has been deprecated and maintained only as stub view for backward compatibility purpose. Instead of **lsportip** command use **lsip** command to view information about configured IP addresses for all type (host, storage or replication) of connectivity and **lsportethernet** command to provide information about physical ethernet ports.

lsportfc

Use the **lsportfc** command to view the status and properties of the Fibre Channel (FC) input/output (I/O) ports for the clustered system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-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.

object_id
(Optional) Specifies the ID of an object that is displayed in the view.

Description

This command enables you to view information about clustered system port status.

This table provides the attribute values that can be displayed as output view data.

The following table shows the possible outputs:

Table 47. ***lsportfc*** *output*

Attribute	Description
id	Indicates a unique value for the object. The value must be a numeric 0 or greater.
fc_io_port_id	Indicates the FC I/O port ID. The value must be a positive integer.
port_id	Indicates the platform port ID. The value must be a positive integer.

*Table 47. **lsportfc** output (continued)*	
Attribute	**Description**
type	Indicates the type of platform port. The value can be either `fc` or `ethernet`.
port_speed	Indicates the I/O port speed. The value is XGb. The value is N/A if the port has never been active. If the port is inactive, it shows the last-known port speed.
node_id	Indicates the ID of the node containing the port. The value must be a positive integer.
node_name	Indicates the name of the node containing the port.
WWPN	Indicates the I/O port worldwide port name (WWPN). The value must be in 16-character hexadecimal format.
nportid	Indicates the most recent NPort ID used by the port. The value must be in 6-character hexadecimal format, and all zeroes if never active.
status	Indicates whether the port is configured to a device of Fibre Channel (FC) port. The values are: • `active` • `inactive_configured` • `inactive_unconfigured`. • `disabled`
switch_WWPN	Indicates the WWPN of the device that was most recently attached to the port. The value must be in 16-character hexadecimal format, or all zeroes if the port has never been active.
fpma	Indicates the Fabric Provided MAC Address (FPMA) assigned to the Fiber Channel VN_PORT port. The value is a formatted 48-bit MAC address. If the switch is connected to an FCF, the value is N/A for ports that are never active. If the port is inactive the last-known FPMA is used.
vlan_id	Indicates the VLAN ID on which a specific VN port is communicating. The value is up to a 4-character decimal string. The value is `N/A` for ports that are never active. If the port is inactive, the last-known VLAN ID is used.
fcf_MAC	Indicates the MAC address for the switch attached to the VN port. The value is N/A for ports that are never active. The value is a formatted 48-bit MAC address. If the port is inactive, the last known `fcf_MAC` value is used.
attachment	Indicates if the port is attached to an FC switch or directly to an FC host. (Or, if the port is offline, it specifies which one it was attached to when last online.)
cluster_use	Indicates the node's current capability for local or partner cluster communications: • `none` indicates the port cannot be used for any node communication • `local` indicates the port can be used for local system node communication • `partner` indicates the port can be used for partner system node communication • `local_partner` indicates the port can be used for both local and partner system node communication
topology	Indicates the current topology: • pp refers to point-to-point • `al` refers to arbitrated loop

Table 47. **lsportfc** output (continued)	
Attribute	**Description**
topology_auto	Indicates whether the topology is currently automatic - yes or no.
adapter_id	Indicates the adapter in a canister the port is in: 1 is left, 2 is right
adapter_location	Indicates the location of the adapter containing the Ethernet port (any number from 0 to 6).
adapter_port_id	Indicates the location of the Ethernet port that is in the adapter (any number from 1 to 4).
fabric_WWN	Indicates the World Wide Name (WWN) value for the connected fabric string. The value must be in 16-character hexadecimal format, or all zeroes if the port has never been active or the port is not connected to a fabric.

lsportsas

Use the **lsportsas** command to display the status of all SAS ports in the clustered system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

Description

This command displays information about system port status.

This command output shows all available paths, which are defined by zoning, independent of their usage. This means that the command output includes paths that are not used because of port masking.

This table provides the attribute values that can be displayed as output view data.

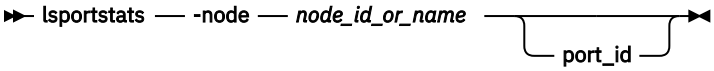
Table 48. **lsportsas** output	
Attribute	**Description**
id	Indicates the line number within the displayed information (numeric string).
port_id	Indicates the ID of the port.
port_speed	Indicates the speed of the I/O port (in XGb). This speed is the fastest local link speed for the SAS port. The value is the last-known port speed if the port is inactive, and N/A if port has is unused and was never active.
node_id	Indicates the ID of the node that contains the port (numeric string).

*Table 48. **lsportsas** output (continued)*	
Attribute	**Description**
node_name	Indicates the name of the node that contains the port (alphanumeric string).
WWPN	Indicates the worldwide port name (WWPN) for the I/O port (16-character hexadecimal string).
status	Indicates the status of the port (numeric string). The values can be: • online if the port is functioning • offline if the port is not functioning • offline_unconfigured if not functioning but not configured by the user • degraded if one or more ports are not functioning or have a lower speed than others • excluded if excluded by the user or system
switch_WWPN	Indicates the WWPN of the switch port if attached to switch (16-character hexadecimal string), or is blank. If the port is offline, the last-known value is displayed.
attachment	Indicates what the port is attached to. The possible values are: • drive • controller • switch • host • enclosure • none If the port is offline, the field shows what was attached when the port was last online.
type	Indicates how the port is configured. This field also shows the devices that can be attached to the SAS port. The possible values are: • drive • enclosure • enclosure_controller • host_controller • none
adapter_location	Displays the location of the adapter that contains the SAS port (any number in the range 0 - 6).
adapter_port_id	Displays the location of the SAS port that is in the adapter (any number in the range 1 - 4).

lsportstats

Use the **lsportstats** command to view the port transfer and failure counts and Small Form-factor Pluggable (SFP) diagnostics data that is recorded in the statistics file for a node.

Syntax



Parameters

-node *node_id_or_name*
(Required) Specifies the node ID or node name of the node that has the port that you are interested in. This value is the **node_id** or **node_name** as seen in the output of **lsnode** or **lsnodecanister**.

port_id
(Optional) Specifies the specific port ID as present in the statistics file for the node. This value is the port ID as seen in output of `lsportstats -node node`.

Description

This command displays the statistics of a port from the statistics file for a node.

The command displays the file name of the node statistic file from which the data displayed is taken.

- Pay attention to **node_dump_name** between "Nn_stats_" and the date stamp in the file name to be sure that it corresponds to the node that has the ports that you are interested in.
- Pay attention to date and timestamp suffix of the file name to understand how recently the file was created. (YYMMDD_HHMMSS)

If a port ID is not provided, attributes for every port in the node is displayed so that you can find which port ID has the worldwide port name (WWPN) or iSCSI qualified name (IQN) that you are interested in. If a port ID is provided, a meaningful description is displayed, instead of the attribute, together with a value for the port ID given. Statistics files are only created for nodes that are online in the system and when statistics is on in the output of **lssystem**. See the **starstats** command. If the software version on your system does not support **lsportstats**, then you can get a utility package from https://www.ibm.com/support/home/, which installs **utilportstats** command such as **utilporstats -node 1**.

This table provides the attribute values that can be displayed as output.

*Table 49. **lsportstats** output*

Attribute	Values
port id	Indicates the port ID. The value is a positive number (integer) in the range 1 - 24.
type	Indicated the connection type. The possible values are FC, iSCSI, SAS, PCIe, or NVMe.
type_id	Indicates the type ID. The value is a positive integer in the range of 1 - 5.
wwpn	Indicates the WWPN or the value is blank.
fc_wwpn	Indicates the Fibre Channel WWPN or the value is blank.
fcoe_wwpn	Indicates the Fibre Channel over Ethernet WWPN or the value is blank.

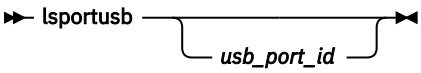
*Table 49. **lsportstats** output (continued)*	
Attribute	**Values**
sas_wwn	Indicates the Serial Attached SCSI WWPN or the value is blank.
iqn	Indicates the IQN for an Ethernet port or the value is blank.
hbt	Indicates the bytes sent to hosts.
hbr	Indicates the bytes received from hosts.
het	Indicates the exchanges that are initiated to hosts.
her	Indicates the exchanges that are received from hosts.
cbt	Indicates the bytes sent to controllers.
cbr	Indicates the bytes received from controllers.
cet	Indicates the exchanges that are initiated to controllers.
cer	Indicates the exchanges that are received from controllers.
lnbt	Indicates the bytes sent to local nodes.
lnbr	Indicates the bytes received from local nodes.
lnet	Indicates the exchanges that are initiated to local nodes.
lenr	Indicates the exchanges that are received from local nodes.
rmbt	Indicates the bytes sent to remote nodes.
rmbr	Indicates the bytes received from remote nodes.
rmet	Indicates the exchanges that are initiated to remote nodes.
rmer	Indicates the exchanges that are received from remote nodes.
dtdt	Indicates the total time in microseconds for which data transmission was excessively delayed.
dtdc	Indicates the number of transfers that experienced excessive data transmission delay.
dtdm	Indicates the number of transfers that had their data transmission delay measured.
dtdt2	Indicates the data transmission delay time 2 in microseconds (µs).
dtdc2	Indicates the data transmission delay count 2.
lf	Indicates the link failure error count.
lsy	Indicates the loss of sync error count.
lsi	Indicates the loss of signal error count.

Table 49. **lsportstats** output (continued)	
Attribute	**Values**
pspe	Indicates the primitive sequence protocol error count.
itw	Indicates the invalid transmission word error count.
icrc	Indicates the invalid CRC error count.
bbcz	Indicates the zero buffer-buffer credit timer in microseconds (µs).
tmp	Indicates the SFP temperature in degrees Celsius.
tmpht	Indicates the SFP temperature high alarm threshold in degrees Celsius.
txpwr	Indicates the TX power in microwatts (µW).
txpwrlt	Indicates the TX power low alarm threshold in microwatts (µW).
rxpwr	Indicates the RX power in microwatts (µW).
rxpwrlt	Indicates the RX power low alarm threshold in microwatts (µW).
hsr	Indicates the slow host Read operations.
hsw	Indicates the slow host Write operations.
har	Indicates the aborted host Read operations.
haw	Indicates the aborted host Write operations.

lsportusb

Use the **lsportusb** command to display information about Universal Serial Bus (USB) ports.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

usb_port_id
(Optional) Specifies the USB port ID. Used when detailed information about a USB port is required.

Description

This command displays information about Universal Serial Bus (USB) ports.

This table provides the attribute values that can be displayed as output view data.

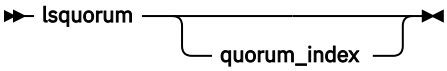
*Table 50. **lsportusb** output*	
Attribute	**Possible Values**
id	Indicates the unique ID of the USB port in the system. This ID is the `usb_port_id`. The value is a numeric `0` or greater.
node_id	Indicates the ID of the node where the USB port is. The value is a numeric string.
node_name	Indicates the name of the node where the USB port is. The value is an alphanumeric string.
node_side	Indicates the side of the node where the USB port is. The values are `front` and `rear`.
port_id	Indicates the ID of the USB port on the node side. The value is a numeric `1` or greater.
status	Indicates the status of the USB port. The values are: • `active`, which indicates that a USB flash drive is plugged in and can be used by the system. • `inactive`, which indicates that no USB flash drive is detected. • `unsupported`, which indicates that a USB device is plugged in but cannot be used.
encryption_state	Indicates the encryption status of the USB device that is attached to the port. The values are: • Blank indicates that it is not in use for encryption • `validated` indicates that encryption keys are present and validated • `missing` indicates that encryption keys were validated and were then removed, and the DMP must run to confirm the absence. • `prepared` indicates that encryption keys are prepared as part of a rekey operation. • `validated_prepared` indicates that encryption keys are validated and prepared as part of a rekey operation. • `wrong_system` indicates that encryption keys are detected on the USB device but none valid for the system. • `old` indicates that the USB device contains encryption keys that were generated for this system - but they are not the current keys. • `error` indicates that an encryption key is detected and something might be wrong with it.
encryption_filename	Indicates the name of the file in the rot directory of the USB device to which the encryption state relates. The file name can contain up to 110 characters.

Table 50. **lsportusb** output (continued)	
Attribute	**Possible Values**
service_state	Indicates the USB command status. The values are: • Blank indicates that no command is active. • `running` indicates that `satask.txt` is processing, and default USB processing is ongoing. • `complete` indicates that `satask.txt` is processing and default USB processing is complete. • `install_image` indicates that `satask.txt` processing cannot start because there is an installation image on the USB flash drive.

lsquorum

Use the **lsquorum** command to list the quorum devices that the system uses to store quorum data.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

quorum_index
(Optional) Specifies the quorum device by its index number. When you use this parameter, a detailed view of the specified device is returned. If you do not specify a device, then a concise view of all quorum devices is displayed.

Description

This command displays a concise list or a detailed view of the MDisks, drives or devices (IP Quorum Apps) that the system is using to store quorum data. This information can be used to ensure that the quorum candidates are on separate storage subsystems.

Note: For disks the object type is either MDisk or drive, but only MDisks are used to hold quorum data. If the quorum object type is a drive or device, the controller ID and name fields are blank.

This table provides the attribute values that can be displayed as output view data.

Table 51. lsquorum output	
Attribute	**Possible Values**
quorum_index	Indicates the quorum device by index number.
status	Indicates the quorum device status.
name	Indicates the name of the object that is used as the quorum device.
controller_id	Indicates the ID of the controller of an MDisk object that is used as the quorum device.
controller_name	Indicates the name of the controller of an MDisk object that is used as the quorum device.
active	Indicates whether this quorum is the active quorum device the system uses as a tie breaker. **Note:** This field is utilised for policy-based replication high availability tie-breaks and existing system tie-breaks. There can be 0 or 1 active quorum device for each application type.
object_type	Indicates the type of object the quorum device uses.
override	Indicates whether the automatic quorum selection for this quorum device was overridden.
site_id	Indicates the site value for the quorum device. This value is blank.
site_name	Indicates the site name for the quorum device (MDisks or drives). This value is blank. For the quorum application, this name identifies the site it is deployed at. By default, this name is the local host's IP address, but a custom alias can also be set.
application_type	Indicates the type of tie-break that the quorum application is used for. The field will be blank for non-app based quorum devices, **system** for single-system tie-break, or **partnership** for policy-based high availability tie-break.
metadata_backup	Indicates whether the quorum device is storing system metadata for the disaster recovery. The field will be blank for partnership application type.
partner_system_id	Indicates the system ID of the remote system that the quorum application is linked to in the policy-based replication high availability partnership. The field will be blank for all but partnership application type.
partner_system_name	Indicates the system name of the remote system in the policy-based replication high availability partnership that the quorum app is attached to. The field will be blank for all but partnership application type.

lsroute

Use the **lsroute** command to display the IP routing table.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

Description

This command displays the IP routing table. The table provides details of the gateway that is used for IP traffic to a range of IP addresses for each Ethernet port. This information can be used to diagnose configuration node accessibility problems. The **lsroute** command is equivalent to the Linux **route** command.

lstimezones

Use the **lstimezones** command to list the time zones that are available on the system. Each time zone is assigned an ID that can be used in the **settimezone** command to set the time zone.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

Description

This command displays a list of all the time zones that are available on the system. Each time zone is assigned an ID. This ID can be used in the **settimezone** command.

lssasportcandidate

Use the **lssasportcandidate** command to list the unconfigured serial-attached SCSI (SAS) ports that are logged in and available to add to the SAS worldwide port name (WWPN) or host objects.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

Description

This command returns a list of unconfigured, logged in SAS ports.

Note: The **lssasportcandidate** command presents a list of host SAS ports that are logged in to nodes. However, there are situations when the list might include host SAS ports that are no longer logged in or even part of the SAN fabric. For example, if a host SAS port is unplugged from a switch but **lssasportcandidate** shows the WWPN that is logged in to all nodes, the incorrect entry is removed when another device is plugged in to the same switch port that previously contained the removed host SAS port.

This table shows the possible output:

Table 52. lssasportcandidate output

Attribute	Description
sas_WWPN	Indicates that the SAS WWPN that is logged in is unconfigured (not assigned to a host). This value must be 16 hexadecimal characters.

lssecurity

Use the **lssecurity** command to display the current system Secure Sockets Layer (SSL) or Transport Layer Security (TLS) security settings.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

Description

This command displays the current system security settings system-wide, including the SSL or TLS and SSH security levels.

This table provides the possible values that are displayed for the **lssecurity** command.

*Table 53. **lssecurity** attribute values*	
Attribute	**Value**
sslprotocol	Indicates the current security level setting, a numeric value from 2 to 7. Use these **sslprotocol** security level settings. • 2 - Allows TLS 1.2, but disallows TLS 1.0 and TLS 1.1. • 3 - Also disallows TLS 1.2 cipher suites that are not exclusive to 1.2. • 4 - Additionally disallows RSA key exchange ciphers and static key exchange ciphers. • 5 (Compatibility mode) - Initially allows TLS 1.3, which is the preferred method of connection. If TLS 1.3 fails, TLS 1.2 is used for connections. • 6 - Allows TLS 1.3 and the five ciphers of this level. • 7 - Allows TLS 1.3 and a single FIPS cipher.

*Table 53. **lssecurity** attribute values (continued)*	
Attribute	**Value**
sshprotocol	Indicates the current security level for SSH, a numeric value of 1, 2, or 3. Use these **sshprotocol** security level settings. • 1 Allows the following key exchange methods. – curve25519-sha256 – curve25519-sha256@libssh.org – ecdh-sha2-nistp256 – ecdh-sha2-nistp384 – ecdh-sha2-nistp521 – diffie-hellman-group-exchange-sha256 – diffie-hellman-group16-sha512 – diffie-hellman-group18-sha512 – diffie-hellman-group14-sha256 – diffie-hellman-group14-sha1 – diffie-hellman-group1-sha1 – diffie-hellman-group-exchange-sha1 • 2 Allows the following key exchange methods. – curve25519-sha256 – curve25519-sha256@libssh.org – ecdh-sha2-nistp256 – ecdh-sha2-nistp384 – ecdh-sha2-nistp521 – diffie-hellman-group-exchange-sha256 – diffie-hellman-group16-sha512 – diffie-hellman-group18-sha512 – diffie-hellman-group14-sha256 – diffie-hellman-group14-sha1 • 3 Allows the following key exchange methods: – curve25519-sha256 – curve25519-sha256@libssh.org – ecdh-sha2-nistp256 – ecdh-sha2-nistp384 – ecdh-sha2-nistp521 – diffie-hellman-group-exchange-sha256 – diffie-hellman-group16-sha512 – diffie-hellman-group18-sha512 – diffie-hellman-group14-sha256
gui_timeout_mins	Indicates the number of minutes of inactivity until a browser session expires. The value is in the range 5 - 240.
cli_timeout_mins	Indicates the number of minutes of inactivity until an SSH session expires. The value is in the range 5 - 240.

*Table 53. **lssecurity** attribute values (continued)*	
Attribute	**Value**
min_password_length	Indicates the minimum number of characters that are required in a new password. The value is in the range 6 - 64.
password_special_chars	Indicates the minimum number of special characters that are required in any new passwords that are created on the system. A value of 0 means that no special characters are required. The value is in the range 0 - 3.
password_upper_case	Indicates the minimum number of uppercase characters that are required in any new passwords that are created on the system. A value of 0 means that no uppercase characters are required. The value is in the range 0 - 3.
password_lower_case	Indicates the minimum number of lowercase characters that are required in any new passwords that are created on the system. A value of 0 means that no lowercase characters are required. The value is in the range 0 - 3.
password_digits	Indicates the minimum number of digits that are required in any new passwords that are created on the system. A value of 0 means that no numbers are required. The value is in the range 0 - 3.
check_password_history	Indicates whether password history is checked to prevent a user from reusing a previous password. The value is either yes or no.
max_password_history	Indicates the number of previous passwords to compare with if **checkpasswordhistory** is enabled. A value of 0 means that the new password is compared with the current password only. The value is in the range 6 - 10.
min_password_age_days	Indicates the minimum number of days between password changes. This setting is enforced if **checkpasswordhistory** is enabled. The value is in the range 0 - 365.
password_expiry_days	Indicates the number of days before a password expires and must be changed. The value is in the range 0 - 365.
expiry_warning_days	Indicates the number of days before a password expires that a warning is raised when the user logs in. The value is in the range 0 - 30.
lockout_period_mins	Indicates the number of minutes a user is locked out for when the number of failed authentication attempts exceeds the **max_failed_logins** value. The value is in the range 0 - 10080.
max_failed_login_attempts	Indicates the number of failed logins that cause the account to become locked. The value is in the range 0 - 10.
superuser_locking	Indicates whether the user locking policy on the system applies to the superuser. The value is either **enabled** or **disabled**.
restapi_timeout_mins	Indicates the total number of minutes of activity until a RESTful API token expires. The value is in the range 10 - 120.
ssh_grace_time_seconds	Indicates the value of the LoginGraceTime field in the SSHD config. The value is in the range 15 - 1800.
ssh_max_tries	Indicates the value of the LoginGraceTime setting in the SSHD config. The value is in the range 1 - 10.
superuser_multi_factor	Indicates if the multi-factor authentication is enabled for the superuser. The value is either **yes** or **no**.
superuser_password_sshkey_required	Indicates whether superuser should provide both password and SSH public key during authentication. The value is either **yes** or **no**.

Table 53. **lssecurity** attribute values (continued)	
Attribute	**Value**
superuser_gui_disabled	Indicates whether GUI access is disabled for superuser. The value is either **yes** or **no**.
superuser_rest_disabled	Indicates whether REST-API access is disabled for superuser. The value is either **yes** or **no**.
superuser_cim_disabled	Indicates whether CIMOM access is disabled for superuser. The value is either **yes** or **no**.
two_person_integrity_enabled	Indicates whether two person integrity (TPI) is enabled on a system. The value is either yes or no. The default value is no. If **two_person_integrity_enabled** is yes and **two_person_integrity_superuser_locked** is no, the system is operating in a state that is inconsistent with TPI operations. In this case, an error event (0989051 - SS_EID_TPI_ENABLED) is logged.
two_person_integrity_superuser_locked	Indicates whether superuser is locked. It shows a value that is the same as the superuser_locked value from the **sainfo lsservicestatus** command.
ssl_protocols_enabled	Indicates the versions of the TLS protocol that are supported by the SSL protocol security level that is currently enabled.
ssl_protocol_suggested	Indicates whether the system is automatically following the suggested SSL protocol level.
ssh_protocol_suggested	Indicates whether the system is automatically following the suggested SSH protocol level.
patch_auto_update	Indicates whether the Patch Auto Updater service is enabled or disabled.

lssite

Use the **lssite** command to report the names of the sites. This command applies only to models that support multiple I/O groups.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

Description

This command reports the names of the sites.

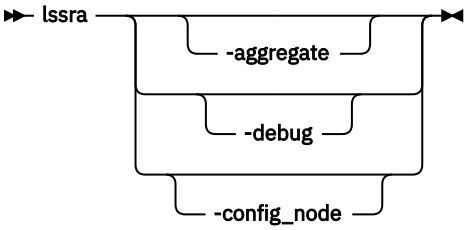
This table provides the attribute values that can be displayed as output view data.

Table 54. lssite attribute values	
Attribute	**Value**
id	Identifies the numeric value that represents the site. The value can be 1, 2, or 3.
name	Identifies the site name.

lssra

Use the **lssra** command to check both secure remote assistance status and the time of the last login.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-aggregate
(Optional) Specifies the output filtering and formatting for the GUI. This parameter is mutually exclusive with all other parameters.

-debug
(Optional) Specifies the debugging option. This parameter is mutually exclusive with all other parameters.

-config_node
(Optional) Specifies the local field values for the current configuration node. Do not specify this parameter if the specified node is part of an active clustered system. This parameter is mutually exclusive with all other parameters.

Description

This command displays both secure remote assistance status and the time of the last login.

This table provides the attribute values that can be displayed as output view data.

Table 55. **lssra** output	
Attribute	**Description**
status	Indicates whether the support assistance is enabled or not enabled (default).

*Table 55. **lssra** output (continued)*	
Attribute	**Description**
token_age_in_days	Displays the number of days that the current token exists.
active_monitor_user_count	Displays the number of support assistance monitor users that are logged in to this system.
monitor_user_last_login	Displays the time of the last monitor user login in the format YYMMDDHHMMSS. The default is blank.
active_privileged_user_count	Displays the number of secure remote access privileged users that are currently logged on.
privileged_user_last_login	Displays the date and time of the last support assistance privileged user login on the system in the format YYMMDDHHMMSS. The default is blank.
remote_support_test_status	Indicates the remote system support service state when it was last tested. The values are: • `disconnected` (default) • `connecting` • `connected` • `active` • `failure` This value is the maximum state that is reached on the configuration node, in increasing order from `disconnected` through to `failure`.
remote_support_test_time	Indicates the time stamp of the last remote system support test. The value must be in the format YYMMDDHHMMSS.
remote_support_config_changed_after_test	Indicates that a support center was added since the system was last tested. The values are `yes` or no (default).
remote_support_enabled	Indicates whether remote support is enabled. The values are `yes` or no (default).
remote_support_status	Indicates the remote support service status or state. If you do not enable remote support the status is `disconnected`. The values are: • `disconnected` (default) • `connecting` • `connected` • `active` • `failure` This value is the maximum state across all online nodes, in increasing order from `disconnected` through to `failure`.
remote_support_enabled_time	Indicates the time stamp for the last successful creation of a secure tunnel by a remote support service. The value must be in the format YYMMDDHHMMSS.
remote_support_idletimeout	Indicates the idle timeout value. The value must be a number (non-negative), and the default value is `0`.
remote_support_center_id	Indicates the support center ID (specified by using **`lssystemsupportcenter`**that is used to establish the secure tunnel. The value must be a number 0 -11, and the default is blank.

*Table 56. **lssra** output*	
Attribute	**Description**
remote_support_mode	Indicates the remote support mode.
remote_support_idletimeout	Indicates the remote support idle timeout.
remote_support_access_code	Indicates the remote support access code.
remote_support_config_changed_after_test	Indicates that remote support configuration changed.
active_service_user_count	Indicates the active service user count.
service_user_last_login	Indicates the service user last login attempt.
remote_support_status	Indicates the remote support status. The values are: • no_status_received • failure • disconnected • connecting • connected • active
remote_support_center	Indicates the remote support center.
remote_support_connected_time	Indicates the remote support connected time.
remote_support_test_time	Indicates the remote support test time.
remote_support_test_status	Indicates the remote support test status.

lsthrottle

Use the **lsthrottle** command to list throttle objects that are configured in the clustered system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

Description

This command lists throttle objects that are configured in the clustered system.

Note: The throttle limits are per node.

This table provides the attribute values that can be displayed as output view data.

Table 57. ***lsthrottle*** *output*

Attribute	Description
throttle_id	Indicates the unique ID for the throttle object. The value is a number 0 - 10144.
throttle_name	Indicates the unique name for the throttle object. The value is an alphanumeric string that is 63 characters long.
object_id	Indicates the ID of the object to which throttle is applied. The value is a number 0 - 8191.
object_name	Indicates the name of the object to which throttle is applied. The value is an alphanumeric string that is 63 characters long.
throttle_type	Indicates the type of throttle object. The values are `offload` or `vdisk`.The values are: `offload`, `vdisk`, `host`, `hostcluster`, and `mdiskgrp`.
IOPs_limit	Indicates the limit of configured IOPs. The value is a numeric string in the range 0 - 33554432. If no limit is specified the value is blank.
bandwidth_limit_MB	Indicates the bandwidth (in MBps). The value is a numeric string in the range 0 - 268435456. If no limit is specified the value is blank.

lssystem

Use the **lssystem** command to display a detailed view of a system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

Description

This command displays a detailed view of a system.

This table provides the attribute values that can be displayed as output view data.

*Table 58. **lssystem** output*	
Attribute	**Possible Values**
layer	The value can be either: • `replication` – The system can create partnerships with other replication-layer systems. – The system can use another system that is configured with storage-layer as external storage. . • `storage` – The system can create partnerships with other storage-layer systems. – The system can provide external storage for replication-layer systems.
location	The location is `local` or `remote`.
statistics status	The status is on or `off`.
auth_service_type	Indicates the type of authentication service.
auth_service_configured	Indicates if an authentication service is configured.
auth_service_enabled	Indicates if an authentication service is enabled.
email_state	The possible values are: • `running` • `stopped` • `invalid`
partnership	The possible values are: • `fully_configured` • `partially_configured_local` • `partially_configured_local_stopped` • `not_present` • `fully_configured_stopped` • `fully_configured_remote_stopped` • `fully_configured_local_excluded` • `fully_configured_remote_excluded` • `fully_configured_exceeded` • Blank

*Table 58. **lssystem** output (continued)*	
Attribute	**Possible Values**
tier	Indicates which system information is being reported. The values are: • `tier0_flash` • `tier1_flash` • `tier_enterprise` • `tier_nearline` • tier_scm
tier	Indicates which system information is being reported.
tier_capacity	Indicates the total MDisk storage in the tier.
tier_free_capacity	Indicates the amount of MDisk storage in the tier that is unused.
compression_active	Indicates whether any compressed volume copies exist in non-data reduction pools. Compressed volumes that are in data reduction pools do not count towards this value.
compression_virtual_capacity	Indicates the total virtual (provisioned) capacity for all compressed volume copies in non-data reduction pools. Compressed volumes that are in data reduction pools do not count towards this value. This value is in unsigned decimal format.
compression_compressed_capacity	Indicates the total used capacity for all compressed volume copies in non-data reduction pools. Compressed volumes that are in data reduction pools do not count towards this value. This value is in unsigned decimal format.
compression_uncompressed_capacity	Indicates the total uncompressed used capacity for all compressed volume copies in non-data reduction pools. Compressed volumes that are in data reduction pools do not count towards this value. This value is in unsigned decimal format.
physical_capacity	Indicates the total physical capacity of all standard-provisioned and thin-provisioned storage that is managed by the storage system. The value is rounded to two decimal places.
physical_free_capacity	Indicates the total free physical capacity of all standard-provisioned and thin-provisioned storage that is managed by the storage system. The value is rounded to two decimal places.

*Table 58. **lssystem** output (continued)*	
Attribute	**Possible Values**
total_reclaimable_capacity	Indicates the unused (free) capacity that will be available after data is reduced. For storage pools that are not data reduction pools, this field reports `0.00MB`. **Note: `total_reclaimable_capacity`** is a sum of all the **`lsmdiskgrp reclaimable_capacities`**, which has a +/-1% margin. Therefore, this field has up to a +/-4% margin.
used_capacity_before_reduction	Indicates the total amount of data that is written to thin-provisioned and compressed volume copies that are in data reduction storage pools - before data reduction occurs. This value does not include standard-provisioned volumes (that can be created in a data reduction storage pool) because they are not eligible for reduction.
used_capacity_after_reduction	Indicates the total amount of capacity that is used for thin-provisioned and compressed volume copies in the storage pool after data reduction occurs.
overhead_capacity	Indicates the overhead capacity consumption in all storage pools that is not attributed to data.
rc_buffer_size (deprecated)	
has_nas_key (deprecated)	
total_drive_raw_capacity	The total known capacity of all discovered drives (regardless of drive use).
email_organization	Indicates the user's organization that is shown in the call home email function.
email_machine_address	Indicates the user's mailing address that is shown in the call home email function.
email_machine_city	Indicates the user's city that is shown in the call home email function.
email_machine_state	Indicates the user's state that is shown in the call home email function.
email_machine_zip	Indicates the user's postal code that is shown in the call home email function.
email_machine_country	Indicates the user's country that is shown in the call home email function.
cache_prefetch	Indicates whether cache prefetching is enabled across the system. The values are on and `off`.
local_fc_port_mask	Indicates the Fibre Channel (FC) input/output (I/O) ports that a system can use for node-to-node communications on a local system if those FC I/O ports exist on a node. The value is 64 binary bits.

*Table 58. **lssystem** output (continued)*	
Attribute	**Possible Values**
partner_fc_port_mask	Indicates the FC I/O ports that a system can use for system-to-system communications on a partner system if those FC I/O ports exist on a node. The value is 64 binary bits.
compression_destage_mode	Indicates the cache destage mode operation for Real-time compression within the cluster.
rc_auth_status (Deprecated)	
vdisk_protection_time	Indicates the volume protection time in minutes (whether volume protection is enabled or not enabled). The value must be a number from 15 (default) to 1440.
vdisk_protection_enabled	Indicates whether volume protection is enabled (`yes`) or disabled (`no`) for a system.
product_name	Indicates the product name.
odx	Indicates whether offloaded data transfer (ODX) is enabled or disabled. The values are `on` and `off`.
easy_tier_acceleration	Indicates Easy Tier and pool balance acceleration status. The values are `on` and `off`.
max_replication_delay (Deprecated)	
partnership_exclusion_threshold (Deprecated)	
enhanced_callhome	Indicates whether to collect enhanced data in the call home report. The values are `on` or `off`. The enhanced reports include operational and event-related data and specific configuration information that is included in the inventory report. This function alerts the support center about hardware failures and potentially serious configuration or environmental issues. The support center can use the configuration information to automatically generate best practices or recommendations that are based on your actual configuration.
censor_callhome	Indicates that sensitive data is stripped from enhanced call home data. The values are `on` or `off`.
total_mdisk_capacity	Indicates the sum of **mdiskgrp** `capacity` plus the capacity of all unmanaged MDisks.
space_in_mdisk_grps	Indicates the sum of **mdiskgrp** `capacity`.
space_allocated_to_vdisks	Indicates the sum of **mdiskgrp** `real_capacity`.
total_free_space	Indicates the sum of **mdiskgrp** `free_capacity`.
total_vdiskcopy_capacity	Indicates the total provisioned capacity of all volume copies in the cluster.
total_used_capacity	Indicates the sum of **mdiskgrp** `used_capacity`.

*Table 58. **lssystem** output (continued)*	
Attribute	**Possible Values**
total_vdisk_capacity	Indicates the total provisioned capacity of volumes in the cluster.
total_allocated_extent_capacity	Indicates the total size of all extents that are allocated to volumes or otherwise in use by the system.
total_overallocation	Indicates the `total_vdiskcopy_capacity` as a percentage of `total_mdisk_capacity`. If `total_mdisk_capacity` is zero, then `total_overallocation` should display `100`.
tier0_flash_compressed_data_used	Indicates the capacity of compressed data used on the flash tier 0 storage tier.
tier1_flash_compressed_data_used	Indicates the capacity of compressed data used on the flash tier 1 storage tier.
deduplication_capacity_saving	Indicates the total amount of used capacity that is saved by data deduplication. This saving is before any compression.
compression_opportunity	Indicates the total amount of capacity of all volume-copies within data reduction pools that are enabled for compression. It does not include Real-time compression capacity or capacity that has been saved by data deduplication.
deduplication_opportunity	Indicates the total amount of `used_capacity_before_reduction` of all volume-copies within data reduction pools that are enabled for data deduplication.
host_unmap	This field displays whether SCSI Unmap support is enabled from the Host perspective. The possible values that are displayed are on or `off`.
backend_unmap	This field displays whether SCSI Unmap support for backend storage is enabled by the administrator. The possible values that are displayed are on or `off`. **Note:** If `backend_unmap` is disabled, the system may still issue **unmap** commands to drives if they are in a storage pool where the `format_required` property is yes.
quorum_mode	This field displays the mode that quorum is in for tie-breaking. The possible values that are displayed are `standard`, `preferred`, or `winner`.
quorum_site_id	This field displays the site ID that quorum uses when the quorum mode is set to **`preferred`** or **`winner`**. The possible values that are displayed are 1 or 2.
quorum_site_name	This field displays the site name that quorum uses when the quorum mode is set to **`preferred`** or **`winner`**.

*Table 58. **lssystem** output (continued)*	
Attribute	**Possible Values**
quorum_lease	This field displays the lease duration that should be used. The possible values that are displayed are `short` or `long`.
automatic_vdisk_analysis	This field displays the scheduling status of automatic volume analysis. The possible values that are displayed are on or `off`.
snapshot_policy_suspended	This field displays the scheduling status of snapshot policy suspended. The possible values that are displayed are `yes` or no. **Note:** All volume groups that have a snapshot policy associated with them can be suspended using **chsystem -snapshotpolicysuspended** `yes`. If suspended, no snapshots can be created or removed by the snapshot scheduler.
easytier	This field displays whether the IBM Easy Tier function on a system level is enabled or disabled. The values are on and `off`.
protection_provisioned_capacity	The amount of capacity provisioned to store point-in-time copies for volumes in the system. This includes snapshots and consistency protection for replication.
protection_written_capacity	The amount of capacity written to store point-in-time copies for volumes in the system before any data reduction. This includes snapshots and consistency protection for replication.
flashcopy_gui_enabled	Displays whether GUI shows the FlashCopy panels. Values can be *yes* or *no*.
console_ip	Displays either the: • Automatically populated system port 1 IP Address - Internet Protocol Version 4 (IPv4) or IPv6. • User-populated IPv4 address. The port value is always *443*, which requires the system to run by using default Hypertext Transfer Protocol Secure (HTTPS).
email_from	Displays the email address in the **from** field of the email that is sent by the system.
snapshot_preserve_parent	Indicates the snapshot and parent preservation behavior. Values can be *yes* or *no*. **Note:** If the value is set to *no*, the parent volume for snapshots goes offline when the system cannot maintain the snapshot. If the value is set to *yes*, the parent volume remains online and the snapshot is deleted.
relationship_bandwidth_limit (Deprecated)	

*Table 58. **lssystem** output (continued)*	
Attribute	**Possible Values**
anomaly_detection	Indicates whether the workload anomaly detection feature is enabled for the system. **Note:** The `anomaly_detection` is on by default on systems that have NVMe support.
flashcopy_default_grainsize	Indicates the default grainsize being used when creating new FlashCopy relationships.
storage_insights_control_access	Indicates whether the storage insights can submit a storage partition migration request or not.
flash_grid_system_uuid	Indicates a globally unique ID for the system, used to identify it within a FlashSystem grid.
flash_grid_uuid	Indicates the globally unique ID for the FlashSystem grid that the system is part of. If the ID is not part of a FlashSystem grid, it remains blank.
flash_grid_name	Indicates the name of the FlashSystem grid that the system is part of. If the name is not part of a FlashSystem grid, it remains blank.
flash_grid_owner_uuid	Indicates the flash_grid_system_uuid that owns the FlashSystem grid of which the system is a member of.
auto_drive_download	Indicates whether the auto drive download feature is enabled for the system or not.

lssystemlimits

Use the **lssystemlimits** command to display a list of object consumption and limits for the system.

Syntax



Parameters

None.

Description

This table provides the attribute values that can be displayed as output view data.

*Table 59. **lssystemlimits** output*	
Attribute	**Description**
type	Indicates the name of the object type.
configured	Indicates the number of objects that currently exist in the system.
max	Indicates the maximum number of objects that can be created on this platform.

Table 59. **lssystemlimits** output (continued)	
Attribute	**Description**
available	Indicates the maximum number of additional objects that can be created on this platform.
warning threshold	Indicates the number of available objects below which a warning event will be logged.
cloud_enabled_volumes	Indicates the number of volumes enabled for cloud backup.

lssystemcert

Use the **lssystemcert** command to list information about the current system Secure Sockets Layer (SSL) certificate.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

Description

This command lists information about the current system Secure Sockets Layer (SSL) certificate and indicates whether an outstanding certificate request exists.

Table 60 on page 199 provides attribute values that can be displayed as output view data.

Table 60. lssystemcert output	
Attribute	**Possible Values**
certificate	Indicates a readable version of the current SSL certificate.
certificate export	Indicates an encoded version of the SSL certificate.
certificate_request_outstanding	Indicates that an unfinished certificate request exists (if the value is yes) and installs a signed certificate. The value is yes or no.
certificate_type	Indicates the certificate type. The values can be self_signed, system_signed, and externally_signed.
automatic_renewal	Indicates whether the certificate is automatically renewed before it expires. Values can be yes or no.

lssystemethernet

Use the **lssystemethernet** command to list system-wide Ethernet priority tag details for host attachment, system, and storage traffic.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

Description

This command lists information about the current system Ethernet priority values for host attachment, system, and storage traffic.

This table provides attribute values that can be displayed as output view data.

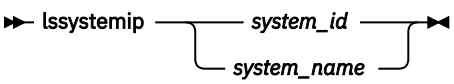
Table 61. ***lssystemethernet*** *output*

Attribute	Possible Values
host_attach_cos	Displays the system-wide host attachment priority tag that is used by all host attachment traffic over iSCSI, or NVMeF. The value is in the range 0 (default) - 7.
system_cos	Indicates the system-wide priority tag that is used by system traffic. The value is in the range 0 to 7, where 0 is the default value.
storage_cos	Displays the system-wide storage priority tag that is used by storage traffic. The value is in the range 0 (default) - 7.

lssystemip

Use the **lssystemip** command to display a list of the system management IP addresses configured for each port.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

system_id | system_name
(Required) Specifies the name or ID of a system.

Description

This command displays a list of the system management IP addresses configured for each port.

This table provides the attribute values that can be displayed as output view data.

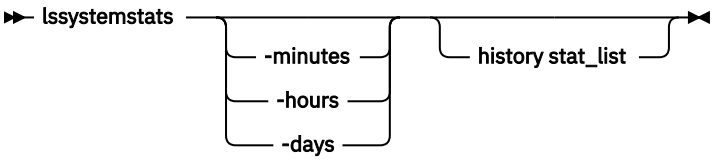
Table 62. lssystemip output

Attribute	Possible Values
cluster_id	Indicates the ID of the system.
cluster_name	Indicates the name of the system.
location	Indicates the system location.
port_id	Indicates the following: • Port id 1 represents the default IP address for the outbound traffic. • Port id 2 represents the optional IP address. **Note:** Port id does not indicate the physical port id. To view the physical port id, use the lsip command.
IP_address	Indicates the Internet Protocol Version 4 (IPv4) address.
subnet_mask	Indicates the IPv4 subnet mask.
gateway	Indicates the IPv4 gateway.
IP_address_6	Indicates the Internet Protocol Version 6 (IPv6) address.
gateway_6	Indicates the IPv6 gateway.
prefix_6	Indicates the IPv6 prefix.

lssystemstats

Use the **lssystemstats** command to display the most recent values of all node statistics in a system, or to display a history of values for a specified subset of available statistics across all nodes in a system. This command also can be used to display a history of values for a specified subset of available statistics.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-minutes
(Optional) Displays stats at minutes granularity.

-hours
(Optional) Displays stats at hours granularity.

-days
(Optional) Show stats at days granularity.

Note:

- Flags are mutually exclusive. In absence of any the live (5 sec) samples are displayed and used with or without **-history** parameter.
- Each granularity will show up to 60 samples per stat when used with **-history** parameter.

-history stat_list
Provides the most recent node statistical values, specific node statistical values, or historical data for any node.

Description

This command returns one set of statistics for all the nodes in the system. The statistical values are determined by using samples that are received from each node.

Note: Values are rounded to the nearest integer when appropriate (for example, between one and ninety-nine for percentages).

This table provides the attribute values that can be displayed as output view data.

Table 63. ***lssystemstats*** *attribute values*

Attribute	Value
stat_current	The current value of the statistic field. The value is accurate up to 3 decimal places (microseconds).
stat_list	The system history of the reported statistics.
stat_name	The name of the statistic field.
stat_peak	The peak value of the statistic field in the last 5 minutes. The value is accurate up to 3 decimal places (microseconds).
stat_peak_time	The time that the peak occurred.
sample_time	The time of the sample occurrence.
stat_value	The statistical value at the epoch interval.

Remember: Filtering is supported on the `stat_name` field by using the concise view.

This table provides the possible values that are applicable to the values that are displayed for the **stat_name** attribute.

*Table 64. **Stat_name** field values*	
Value	**Description**
compression_cpu_pc (Deprecated)	Displays the percentage of allocated CPU capacity that is used for compression.
cpu_pc	Displays the percentage of allocated CPU capacity that is used for the system.
fc_mb	Displays the total number of megabytes transferred per second (MBps) for Fibre Channel traffic on the system. This value includes host I/O and any bandwidth that is used for communication within the system.
fc_io	Displays the total input/output (I/O) operations that are transferred per seconds for Fibre Channel traffic on the system. This value includes host I/O and any bandwidth that is used for communication within the system.
sas_mb	Displays the total number of megabytes transferred per second (MBps) for serial-attached SCSI (SAS) traffic on the system. This value includes host I/O and bandwidth that is used for background RAID activity.
sas_io	Displays the total I/O operations that are transferred per second for SAS traffic on the system. This value includes host I/O and bandwidth that is used for background RAID activity.
iscsi_mb	Displays the total number of megabytes transferred per second (MBps) for iSCSI traffic on the system.
iscsi_io	Displays the total I/O operations that are transferred per second for iSCSI traffic on the system.
write_cache_pc	Displays the percentage of the write cache usage for the node.
total_cache_pc	Displays the total percentage for both the write and read cache usage for the node.
vdisk_mb	Displays the average number of megabytes transferred per second (MBps) for read and write operations to volumes during the sample period. **Note:** Only the write operation value is displayed.
vdisk_io	Displays the average number of I/O operations that are transferred per second for read and write operations to volumes during the sample period.
vdisk_ms	Displays the average amount of time in microseconds that the system takes to respond to read and write requests to volumes over the sample period.
mdisk_mb	Displays the average number of megabytes transferred per second (MBps) for read and write operations to MDisks during the sample period.
mdisk_io	Displays the average number of I/O operations that are transferred per second for read and write operations to MDisks during the sample period.
mdisk_ms	Displays the average amount of time in microseconds that the system takes to respond to read and write requests to MDisks over the sample period.
drive_mb	Displays the average number of megabytes transferred per second (MBps) for read and write operations to drives during the sample period.
drive_io	Displays the average number of I/O operations that are transferred per second for read and write operations to drives during the sample period.
drive_ms	Displays the average amount of time in microseconds that the system takes to respond to read and write requests to drives over the sample period.

*Table 64. **Stat_name** field values (continued)*	
Value	**Description**
vdisk_w_mb	Displays the average number of megabytes transferred per second (MBps) for read and write operations to volumes during the sample period.
vdisk_w_io	Displays the average number of I/O operations that are transferred per second for write operations to volumes during the sample period.
vdisk_w_ms	Displays the average amount of time in microseconds that the system takes to respond to write requests to volumes over the sample period.
mdisk_w_mb	Displays the average number of megabytes transferred per second (MBps) for write operations to MDisks during the sample period.
mdisk_w_io	Displays the average number of I/O operations that are transferred per second for write operations to MDisks during the sample period.
mdisk_w_ms	Displays the average amount of time in microseconds that the system takes to respond to write requests to MDisks over the sample period.
drive_w_mb	Displays the average number of megabytes transferred per second (MBps) for write operations to drives during the sample period.
drive_w_io	Displays the average number of I/O operations that are transferred per second for write operations to drives during the sample period.
drive_w_ms	Displays the average amount of time in microseconds that the system takes to respond write requests to drives over the sample period.
vdisk_r_mb	Displays the average number of megabytes transferred per second (MBps) for read operations to volumes during the sample period.
vdisk_r_io	Displays the average number of I/O operations that are transferred per second for read operations to volumes during the sample period.
vdisk_r_ms	Displays the average amount of time in microseconds that the system takes to respond to read requests to volumes over the sample period.
mdisk_r_mb	Displays the average number of megabytes transferred per second (MBps) for read operations to MDisks during the sample period.
mdisk_r_io	Displays the average number of I/O operations that are transferred per second for read operations to MDisks during the sample period.
mdisk_r_ms	Displays the average amount of time in microseconds that the system takes to respond to read requests to MDisks over the sample period.
drive_r_mb	Displays the average number of megabytes transferred per second (MBps) for read operations to drives during the sample period
drive_r_io	Displays the average number of I/O operations that are transferred per second for read operations to drives during the sample period.
drive_r_ms	Displays the average amount of time in microseconds that the system takes to respond to read requests to drives over the sample period.
power_w	Displays the total system power in watts.
temp_c	Displays the ambient temperature in Celsius.
temp_f	Displays the ambient temperature in Fahrenheit.

*Table 64. **Stat_name** field values (continued)*	
Value	**Description**
iplink_mb	Displays the average number of megabytes requested to be transferred per second (MBps) over the IP partnership link during the sample period. This value is calculated before any compression of the data takes place. This value does not include iSCSI host input/output (I/O) operations.
iplink_io	Displays the total input or output (I/O) operations transferred per second for IP Replication traffic on the system. This value does not include iSCSI host I/O.
iplink_comp_mb	Displays the average number of compressed megabytes transferred per second (MBps) over the IP Replication link during the sample period. This value is calculated after any compression of |the data takes place. This value does not include iSCSI host I/O operations. **Note:** If compression is disabled, the `iplink_mb stats` ID value is displayed instead.
cloud_up_mb	Displays the average number of megabytes transferred per second (Mbps) for upload operations to a cloud account during the sample period.
cloud_up_ms	Displays the average amount of time (in microseconds) it takes for the system to respond to upload requests to a cloud account during the sample period.
cloud_down_mb	Displays the average number of Mbps for download operations to a cloud account during the sample period.
cloud_down_ms	Displays the average amount of time (in microseconds) it takes for the system to respond to download requests to a cloud account during the sample period.
nvme_rdma_mb	Displays the total number of megabytes transferred per second (MBps) for NVMe over RDMA traffic on the system.
nvme_rdma_io	Displays the total I/O operations that are transferred per second for NVMe over RDMA traffic on the system.
nvme_tcp_mb	Displays the total number of megabytes transferred per second (MBps) for NVMe over TCP traffic on the system.
nvme_tcp_io	Displays the total I/O operations that are transferred per second for NVMe over TCP traffic on the system.

lstargetportfc

Use the **lstargetportfc** command to generate the list of host count and login count for each worldwide port names (WWPNs) required to set up Fibre Channel (FC) zoning and to display the current failover status of host IO ports.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

Description

This command generates lists of worldwide port names (WWPNs) required to set up Fibre Channel (FC) zoning. This command also displays the current failover status of host I/O ports.

This table provides the attribute values that can be displayed as output view data.

Table 65. ***lstargetportfc*** *output*

Attribute	Description
id	Indicates the ID of the port.
WWPN	Indicates the WWPN of the port. The value is hexadecimal.
WWNN	Indicates the worldwide node name (WWNN) of the port. The value is hexadecimal.
port_id	Indicates the system port ID. The value is the same as the **lsportfc** `port_id` field.
owning_node_id	Indicates the ID of the node that owns the port. **Note:** This node can be offline whether the port is online or offline.
current_node_id	Indicates the ID of the node on which this port is active. The value is blank if the port is not active on any node
nportid	Indicates the `nportid` hexadecimal value.
host_io_permitted	Indicates whether host I/O operations can run on the port. The values are `yes` and `no`.
virtualized	Indicates whether it is a virtualized port. The values are `yes` and `no` (which indicates that this port cannot be online on any node other than the owning node).
protocol	Indicates the protocol supported by the port. The values are **scsi** and **nvme**. **Note:** Some Fibre channel adapters and platforms do not support NVMe over Fibre Channel.
fc_io_port_id	Indicates the Fibre Channel IO port ID.
portset_count	Indicates the number of portsets the WWPN is added to.
host_count	Indicates the number of hosts configured through WWPN through the same portset.
active_login_count	Indicates the number of active login counts through WWPN. For a configured host, blocked login will not be included.

lscompatibilitymode

Use the **lscompatibilitymode** command to display the information on features that are enabled for compatibility.

Syntax



Parameters

Parameters are not applicable for this command.

Description

This table lists the attributes that are displayed.

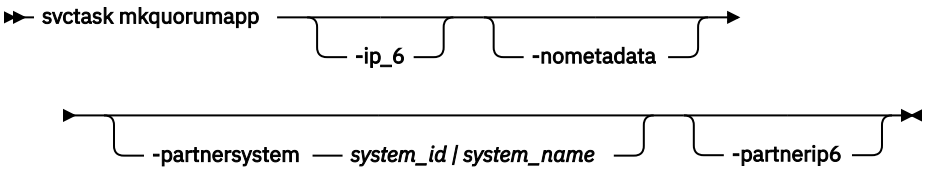
*Table 66. **lscompatibilitymode** output*

Attribute	Value
feature	Indicates the name of the feature.
enabled	Indicates whether the feature is enabled on the system.
in_use	Indicates whether the feature is in use. **Note:** A feature cannot be disabled if it is in use.

mkquorumapp

Use the **mkquorumapp** command to generate a Java™ application to use for quorum.

Syntax



Parameters

-ip_6
(Optional) Specifies that the quorum application uses IPv6 service addresses to connect to the system that generated the quorum application. If this parameter is not specified, IPv4 is used to connect this system.

-nometadata
(Optional) This parameter is used if you do not require metadata that stores configuration data for node recovery operations with the IP quorum application. This parameter cannot be specified with the **-partnersystem** parameter.

-partnersystem
(Optional) Specifies the system ID or system name for the remote system in the policy-based replication high availability partnership to provide quorum.

Note: If this parameter is specified, the quorum app will only be used for policy-based high availability tie-break. If it is not specified, the quorum app will only be used for single-system tie-break and (optionally) metadata storage.

-partnerip6
(Optional) Specifies that the quorum application uses IPv6 service addresses to connect to the partner system. If this parameter is not specified, IPv4 is used to connect to the partner system.

Description

This command generates a Java application to use as quorum for this system or for high availability with a partnered system.

An invocation example

This example creates the file `/dumps/ip_quorum.jar` to use for IP quorum on an IPv4 network.

```
mkquorumapp
```

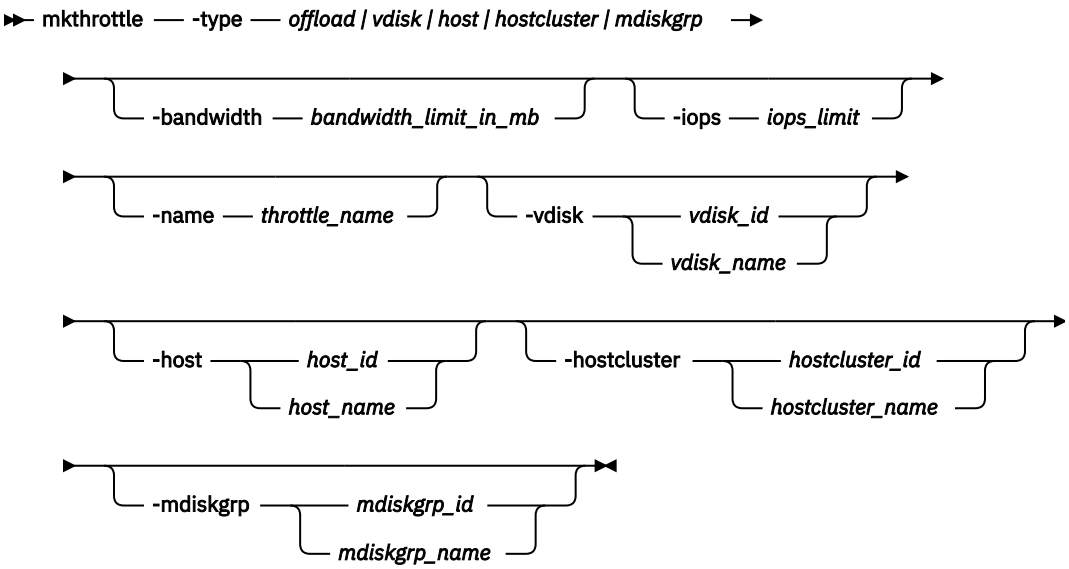
The detailed resulting output:

```
No feedback
```

mkthrottle

Use the **mkthrottle** command to create a new throttle object and associate it with an object (such as a volume). You can also create offloaded I/O throttling (which is a single clustered system throttle).

Syntax



Parameters

-type ***offload | vdisk | host | hostcluster | mdiskgrp***
(Required) Specifies the type of throttle.

-bandwidth ***bandwidth_limit_in_mb***
(Optional) Specifies the bandwidth in MBps. This must be a numeric value 0 - 268435456.

-iops ***iops_limit***
(Optional) Specifies the I/O operations limit. This must be a numeric value 0 - 33554432.

-name ***throttle_name***
(Optional) Specifies the throttling object's name. This value must be an alphanumeric string up to 63 characters long.

-vdisk *vdisk_id | vdisk_name*
(Optional) Specifies the volume ID or name of the volume to throttle. The value must be a numeric or alphanumeric string.

Note: This keyword must be specified when you specify `-type vdisk`.

This parameter is mandatory for volume throttling but cannot be used for offload throttling.

-host *host_id | host_name*
(Optional) Specifies the host ID or name to throttle.

-hostcluster *hostcluster_id | hostcluster_name*
(Optional) Specifies the host cluster ID or name to throttle.

-mdiskgrp *mdiskgrp_id | mdiskgrp_name*
(Optional) Specifies the MDisk group (storage pool) or name to throttle. This applies to parent or child storage pools.

Description

This command creates a new throttle object and associates it with an object (such as a volume).

Note:

- The throttle I/O limit is per node.
- A throttle object cannot be defined for a host if it is a part of host cluster that already has a host cluster throttle object defined for it.
- If a host cluster does not have a throttle object defined, its member hosts can have individual host throttles defined.
- If a volume has multiple copies then throttling is done for the storage pool serving primary copy. Throttling is not applicable for secondary pools that are part of mirrored volumes.

An invocation example for creating a volume throttle of 10000 IOPs and a bandwidth limit of 500 MBps for volume vdisk0

```
mkthrottle -type vdisk -iops 10000 -bandwidth 500 -vdisk vdisk0
```

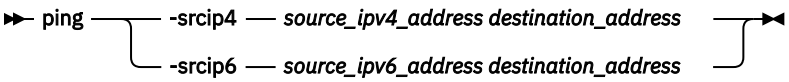
The detailed resulting output:

```
No feedback
```

ping

Use the **ping** command to diagnose IP configuration problems. The command checks whether the specified IP address is accessible from the node on which the command is run by using the specified IP address.

Syntax



Parameters

-srcip4 *source_ipv4_address destination_address*
(Required if **-srcip6** is not specified) Specifies the IPv4 address that sends the ping packet. The IPv4 address must already be bound to a port on the node on which the command is issued.The source address must be in IPv4 format. The destination address can either be IPv4 format or a fully

qualified domain name (FQDN) (max 256 characters). If you do not specify this parameter, you must specify **srcip6**.

-srcip6 ***source_ipv6_address destination_address***
(Required if **-srcip4** is not specified) Specifies the IPv6 address that sends the ping packet. The IPv6 address must already be bound to a port on the node on which the command is issued.The source address must be in IPv6 format. The destination address can either be IPv6 format or a fully qualified domain name (FQDN) (max 256 characters). If you do not specify this parameter, you must specify **srcip4**.

Description

This command checks whether the specified IP address is accessible from the node on which the command is run by using the specified IP address.

Use this command to ping from any port on any node when you are logged on to the service assistant on that node.

An invocation example

```
ping -srcip4 192.168.1.51 192.168.1.30
```

The resulting output

```
PING 192.168.1.51 (192.168.1.51)PING 9.20.136.11 (9.20.136.11) 56(84) bytes of data.
64 bytes from 192.168.1.51: icmp_seq=1 ttl=249 time=0.690 ms
64 bytes from 192.168.1.51: icmp_seq=2 ttl=249 time=0.382 ms
64 bytes from 192.168.1.51: icmp_seq=3 ttl=249 time=0.311 ms

PING 192.168.1.30 (192.168.1.30)PING 9.20.136.11 (9.20.136.11) 56(84) bytes of data.
64 bytes from 192.168.1.30: icmp_seq=1 ttl=249 time=0.690 ms
64 bytes from 192.168.1.30: icmp_seq=2 ttl=249 time=0.382 ms
64 bytes from 192.168.1.30: icmp_seq=3 ttl=249 time=0.311 ms
```

rmiscsistorageport

Use the **rmiscsistorageport** command to remove established Internet Small Computer Systems Interface (iSCSI) sessions between system nodes and backend iSCSI target.

Syntax



Parameters

lsiscsistorageport_row_id
(Required) Specifies the row ID of the selected row in the output of **lsiscsistorageport** command.

Description

Use this command to remove path groups that are established after you specify **addiscsistorageport** (not individual paths).

Any dependencies must be removed before you specify this command. The ID of the session that is listed (after you specify **lsiscsistorageport**) is used to identify the sessions that must be removed.

A detailed invocation example

First, specify **addiscsistorageport** for discovery and **lsiscsistorageport** to show any added session. Then, specify **rmiscsistorageport** to remove sessions that are indicated by view ID 0. No *tgt_user_name* or *target_chap* is required for discovery or session establishment.

```
rmiscsistorageport 0
```

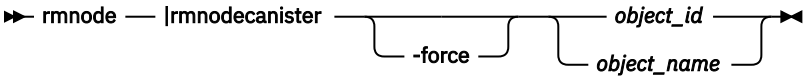
The following detailed output is displayed:

```
No feedback
```

rmnode

The **rmnode** / **rmnodecanister** command deletes a node (or spare node) from the clustered system. You can enter this command anytime after a clustered system is created. This command makes the node a candidate that is ready to be added back into this clustered system or another system.

Syntax



Parameters

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

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

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

object_id | object_name
(Required) Specifies the object name or ID that you want to modify. The variable that follows the parameter is either:

- The object name that you assigned when you added the node to the clustered system
- The object ID that is assigned to the node (not the worldwide node name)

Note: Only a user with the Security Administrator or superuser role is permitted to use **-force** parameter to delete a node from the clustered system when there are Safeguarded snapshots or Safeguarded backups.

Description

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

Attention: When you run the **rmnode** command to remove the configured hardware for a node:

- Small Computer System Interface-3 (SCSI-3) reservations and registrations (through that node) are removed.

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

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

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

Prerequisites:

Before you issue the **rmnode** / **rmnodecanister** command, remember the following points:

Removing the last node

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

Removing the configuration node

- If you are removing the configuration node or node canister, the **rmnode** / **rmnodecanister** command causes the configuration node or node canister to move to a different node or node canister within the clustered system. This process can take a short time, typically less than a minute. The clustered system IP address remains unchanged, but any SSH client that is attached to the configuration node or node canister might need to reestablish a connection. The management GUI reattaches to the new configuration node or node canister transparently.
- If this is the last node or node canister in the clustered system or if it is currently assigned as the configuration node, all connections to the system are lost. The user interface and any open CLI sessions are lost if the last node or node canister in the clustered system is deleted. A time-out might occur if a command cannot be completed before the node or node canister is deleted.
- When **rmnode** command is used to remove a node, it can no longer use the system password so it uses the default password instead.

Removing the spare node

When the **rmnode** command is used against a spare node, the output is determined by the state of that spare node.

- If the **rmnode** command is called against an inactive spare, then the spare is deactivated and deleted from the clustered system. The node is transitioned to a "candidate" state and the **-force** parameter is not required.
- If the **rmnode** command is called against an active spare, then the spare is deactivated and deleted from the clustered system. But the **-force** parameter is required if checks indicate that a volume is

offline or can be corrupted if this action takes place. The node is transitioned to a "candidate" state and is no longer a spare node.

An invocation example

```
rmnode 1
```

The resulting output:

```
No feedback
```

An invocation example

```
rmnodecanister 1
```

The resulting output:

```
No feedback
```

rmnodecanister

Use the **rmnodecanister** command to delete a node canister from the system. Enter this command at any time after a system has been created.

Syntax



Parameters

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

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

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

object_id | object_name
(Required) Specifies the object name or ID that you want to modify. The variable that follows the parameter is either:

- The object name that you assigned when you added the node to the clustered system
- The object ID that is assigned to the node - not the worldwide node name (WWPN)

Description

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

Attention: When you run the **rmnodecanister** command to remove the configured hardware for a node canister:

- Small Computer System Interface-3 (SCSI-3) reservations (through that node canister) are removed
- Small Computer System Interface-3 (SCSI-3) registrations (through that node canister) are removed

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

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

Run the **rmnodecanister** command with the **force** parameter to take the specified node canister offline immediately without flushing the cache (while ensuring that there is no data loss).

Attention: These prerequisites should be considered:

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

Attention:

1. Removing the last node canister in the system destroys the system. Before you delete the last node canister in the system, ensure that you want to destroy the system.
2. If you are removing a single node canister and the remaining canister in the I/O group is online, the data can be exposed to a single point of failure if the remaining canister fails.
3. This command might take some time to complete since the cache in the I/O group for that node canister is flushed before the canister is removed. If the **force** parameter is used, the cache is not flushed and the command completes more quickly. However, if the deleted canister is the last canister in the I/O group, using the **force** option results in the write cache for that canister being discarded rather than flushed, and data loss can occur. Use the **force** option with caution.
4. If both nodes in the I/O group are online and the volumes are already degraded before deleting the node, redundancy to the volumes is already degraded 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 **rmnodecanister** command causes the configuration node canister to move to a different canister within the system. This process might take a short time: typically less than a minute. The system IP address remains unchanged, but any SSH client attached to the configuration node canister might need to reestablish a connection. The management GUI reattaches to the new configuration node canister transparently.
2. If this is the last node canister in the system 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 canister in the system is deleted. A time-out might occur if a command cannot be completed before the node canister is deleted.
3. When **rmnodecanister** command is used to remove a node, it can no longer use the system password so it uses the default password instead.

An invocation example

```
rmnodecanister 1
```

The resulting output

```
No feedback
```

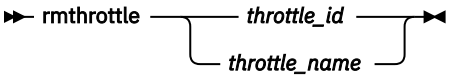
rmportip (Deprecated)

The **rmportip** command has been deprecated. Use the **rmip** command instead.

rmthrottle

Use the **rmthrottle** command to remove the throttle object associated with any volume.

Syntax



Parameters

throttle_id | throtle_name
(Required) Specifies the throttle object ID or name. The value must be a numeric or alphanumeric string up to 63 characters long.

Description

This command removes the throttle object associated with a specified volume.

An invocation example to remove a throttle object with the ID 2

```
rmthrottle 2
```

The detailed resulting output:

```
No feedback
```

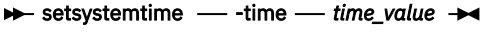
setclustertime (Discontinued)

Attention: The **setclustertime** command has been discontinued. Use the **setsystemtime** command instead.

setsystemtime

Use the **setsystemtime** command to set the time for the system.

Syntax



Parameters

-time ***time_value***
(Required) Specifies the time to which the system must be set. This value must be in the following format (where M is month, D is day, H is hour, m is minute, and Y is year):

MMDDHHmmYYYY

Description

This command sets the time for the system.

An invocation example

```
setsystemtime -time 040509142003
```

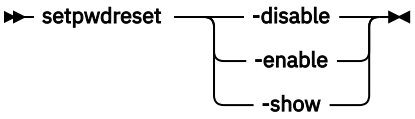
The resulting output

```
No feedback
```

setpwdreset

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

Syntax



Parameters

-disable
Disables the password-reset feature that is available through the USB flash drive interface.

-enable
Enables the password-reset feature that is available through the USB flash drive interface.

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

Description

The system provides an option to reset the system superuser password to the default value. For all systems, this option can be done by using the USB flash drive interface. If this feature remains enabled, make sure that there is adequate physical security to the system hardware. If this feature is disabled, it is not possible to reset the superuser password by using a USB flash drive. You can view or change the status of this feature.

Note: The superuser can enable the *password and key required* feature only when the password-reset feature is enabled. Conversely, the password-reset feature cannot be disabled as long as the *password and key required* feature remains enabled.

An invocation example

```
setpwdreset -show
```

The resulting output:

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

This output means that the password or reset feature is enabled. If the password status is `[0]`, this feature is disabled.

settimezone

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

Syntax



Parameters

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

Description

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

Set the time zone to use when formatting the event log that is produced by issuing `dumperrlog`

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

An invocation example

```
settimezone -timezone 5
```

The resulting output:

```
No feedback
```

showtimezone

Use the **showtimezone** command to display the current time zone settings for the cluster.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

Description

This command displays a single time zone and its associated ID. This is the current time zone setting for the cluster. A list of available time-zones can be viewed by running the **lstimezones** command. The time zone can be changed by running the **settimezone** command.

An invocation example

```
showtimezone -delim :
```

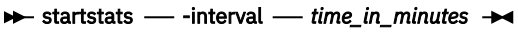
The resulting output:

```
id:timezone
522:UTC
```

startstats

Use the **startstats** command to modify the interval at which per-node statistics for volumes, managed disks (MDisks), and nodes are collected.

Syntax



Parameters

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

Statistics must be set to 1 or 5 minute intervals when sending data to Storage Insights.

Description

Running the **startstats** command resets the statistics timer to zero, and give it a new interval at which to sample. Statistics are collected at the end of each sampling period as specified by the **-interval** parameter. These statistics are written to a file, with a new file created at the end of each sampling period. Separate files are created for MDisks, volumes, and node statistics.

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

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

```
Nm_stats_nodepanelname_date_time
Nv_stats_nodepanelname_date_time
Nn_stats_nodepanelname_date_time
```

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

These files can be listed by using the **lsdumps** command.

The following naming convention is used for these files:

*stats_type*_stats_*nodepanelname_date_time*

Where:

- The value for *stats_type* is `Nm` for MDisks, `Nv` for volumes, and `Nn` for node statistics.
- The value for *nodepanelname* is the current configuration node panel name.
- The value for *date* is in the format of `yymmdd`.
- The value for *time* is in the format of `hhmmss`.

This an example of:

- An MDisk statistics file name: `Nm_stats_000229_031123_072426`
- A volume statistics file name: `Nv_stats_000229_031123_072426`
- 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 is read and written during the sample period
- Per MDisk, cumulative read and write external response times in milliseconds
- Per MDisk, cumulative read and write queued response times

Statistics are collected for each volume and recorded in the `Nv_stats_`*`nodepanelname_date_time`* file, including the following statistical information:

- The total number of processed SCSI read and write commands
- The total amount of read and written data
- Cumulative read and write response time in milliseconds
- Statistical information about the read or write cache usage
- Mirroring statistics that include latency

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

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

An invocation example

```
startstats -interval 25
```

The resulting output:

```
No feedback
```

stopstats (Deprecated)

The **stopstats** command is deprecated. You can no longer disable statistics collection.

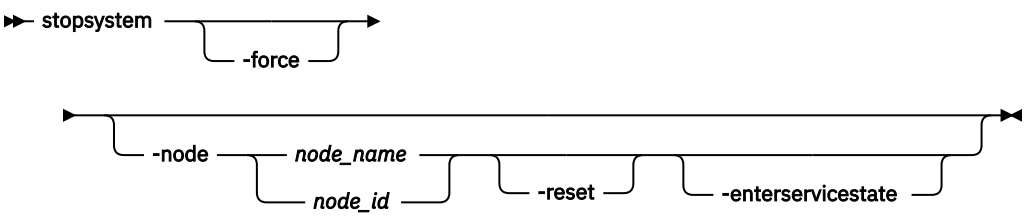
stopcluster (Discontinued)

Attention: The **stopcluster** command is discontinued. Use the **stopsystem** command instead.

stopsystem

Use the **stopsystem** command to shut down a single node or the entire system in a controlled manner. When you issue this command, you are prompted with a confirmation of intent to process the command.

Syntax



Parameters

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

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

If you use the **-force** parameter as a result of an error about volumes going offline, you force the node to shut down, even if it is the last online node in the I/O group.

Remember: The **-force** parameter should always be used with caution.

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

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

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

-reset
(Optional) Soft-resets a specified node. Must be used with -node *node_id/node_name*.

-enterservicestate
(Optional) Puts the specified node into service state.

Description

Use this command to shut down a single node or the entire clustered system in a controlled manner. You are prompted with a confirmation of intent (to process the command) when you specify this command.

For IBM Storage Virtualize for Public Cloud on Amazon Web Services (AWS), you must stop the EC2 instances on AWS after you use this command.

If you need to shut down the system or a single node, use this command instead of removing the power to the node or node canister, or powering off the main power supplies to the system.

Before shutting down a node or system, complete the following requirements:

1. Quiesce all I/O operations that are destined for this node or system. If you do not quiesce these, failed I/O operations might be reported to your host operating systems.
2. Stop all FlashCopy and data migration operations.
3. Ensure that all asynchronous deletion operations have completed.

Using this command to shut down a single node fails if shutting down the node makes any volumes inaccessible, or if it is the last node in an I/O group. If you still need to shut down the node, you can use the **-force** option to override these checks.

Check the Status LEDs on the node to ensure it has shutdown completely before moving the hardware.

An invocation example

```
stopsystem
```

The following confirmation prompt is displayed:

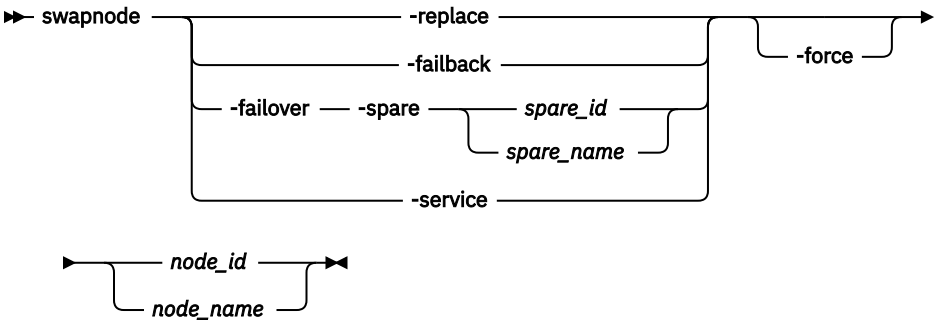
```
 Are you sure that you want to continue with the shut down?
```

Select yes to confirm or no to cancel.

swapnode

Use the **swapnode** command to facilitate maintenance actions that maintain availability without interruptions by using virtualized target ports on a hot spare node. This command is used during hardware replacement procedures, or as alternative to place a node in service state.

Syntax



Parameters

-service | -failover | -failback | -replace
(Required) Specifies the type of service action is being performed on a selected node. The values are:

- **-service:** puts an active online node into service state after the system triggers a failover of N_Port ID Virtualization (NPIV) ports to the partner node in the I/O group. This command can succeed whether or not spare protection is available.
- **-failover:** fails over a clustered node to a standby spare, even if the clustered node is online. The specified node moves into service and the standby spare becomes an online_spare. The command can succeed with or without the optional **-spare** parameter. When used with the **-spare** parameter, it defines the standby spare to be used; otherwise, the best fit standby spare node is chosen.
- **-failback:** fails back an online_spare to a standby spare, and the node it protects exits the service mode and returns to the cluster. You must specify the node object ID or name of the online_spare for this command.
- **-replace:** replaces a specified node that is logically offline or in a service state with a candidate node whose WWNN matches the specified node.

  Note: This **-replace** parameter is used in hardware upgrade procedures or when directed by (and along with) IBM support to recover node that is reporting in error code 578 (loss of hardened data).

-force
(Optional) Specifies that a node be removed, even if disruption to host system I/O might occur as a result.

Important: Specifying **-force** might result in a loss of access. Use it only under the direction of your product support information.

-spare *node_id | node_name*
(Optional) Specifies the ID or name of the standby spare to be brought online as a temporary stand-in. This parameter can be used only with **-failover** parameter.

Note: The value for the ID must be greater than 1 because the spare is never the first node in the clustered system.

node_id | node_name
(Required) Specifies the node ID or name that is being swapped or serviced.

Description

Use the **swapnode** command to facilitate the maintenance actions by failing over a specified node to a hot spare and maintaining availability without interruptions on its virtualized target ports, failing back an online spare to a standby spare, and returning the original node to the cluster. This command can also be used during hardware replacement procedures to complete tasks that require the **rmnode** and **addnode** commands or as an alternative to placing a node in service state.

The **-replace** parameter simplifies the process (commands) required to replace an offline or service node, which is protected by an online_spare. It uses the I/O group and site values that are associated with the offline node as configuration parameters for the replacement. The replacement node must be a candidate node with the same WWNN value as the offline node.

Note: The WWNN of a new replacement candidate node must be changed only to the WWNN of a logically offline node, or node in service state, when that node is physically powered off and disconnected from the system.

Important: NVMe virtualized target ports are not supported on some Fibre Channel adapters. Before you use the CLI, check the spare node configuration. Allows the user to download the cloud usage history.

An invocation example

The following example lists the steps to put a node into service mode.

1. Use the **lsnode** command to list all the nodes in the system. In the following example of **lsnode** output, the node "node2" with ID 2 is in online state.

   ```
   lsnode
   ```

   The 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
   panel_name enclosure_id canister_id enclosure_serial_number site_id site_name
   1  node1                 500507680C000128 online       0
   io_grp0       yes                        SV1      iqn.1986-03.com.ibm:2145.mcr-cay-
   cluster-23.node1            G71H00P                  1
   2  node2                 500507680C000130 online       0
   io_grp0       no                         SV1      iqn.1986-03.com.ibm:2145.mcr-cay-
   cluster-23.node2            G71H00M                  1
   3  spare1                500507680C000138 online_spare 1
   io_grp1       no                         SV1      iqn.1986-03.com.ibm:2145.mcr-cay-
   cluster-23.node3            G71H00X                  1
   ```

2. Issue the **swapnode -service** command with the ID or name of the node. The following example changes the node with ID 2 to the service mode.

   ```
   swapnode -service 2
   ```

   The detailed resulting output:

   ```
   No feedback
   ```

3. Use **lsnode** command to verify that the node is in the service state. The following example shows the node "node 2" with ID 2 is in service state.

   ```
   lsnode
   ```

   The 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
   panel_name enclosure_id canister_id enclosure_serial_number site_id site_name
   1  node1                 500507680C000128 online       0
   io_grp0       yes                        SV1      iqn.1986-03.com.ibm:2145.mcr-cay-
   cluster-23.node1            G71H00P                  1
   2  node2                 500507680C000130 service      0
   io_grp0       no                         SV1      iqn.1986-03.com.ibm:2145.mcr-cay-
   ```

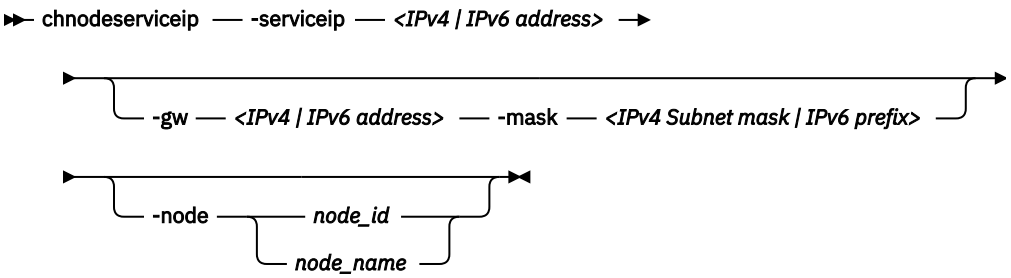
```
cluster-23.node2             G71H00M                 1
3  spare1                 500507680C000138 online_spare 1
io_grp1       no                        SV1      iqn.1986-03.com.ibm:2145.mcr-cay-
cluster-23.node3             G71H00X                 1
```

chnodeserviceip

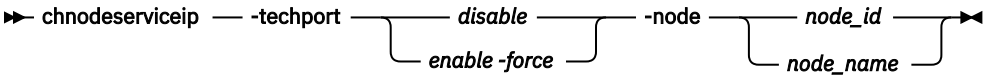
Use the chnodeserviceip command to set the service address for a specific node. This command enables a non-superuser to set the service address for a specific node.

Syntax

Setting static service IP



Setting Techport



Parameters

-serviceip ***<IPv4 | IPv6 address>***
(Required) Specifies the new IP address for the service assistant. The type of the IP address indicates which service IP will be updated. The Internet Protocol Version 6 (IPv6) and Internet Protocol Version 4 (IPv4) address is used for the service assistant.

-gw ***<IPv4 | IPv6 address>***
(Optional) Specifies the IPv4 or IPv6 address to access the gateway with the new service IP.

-mask ***<IPv4 Subnet mask | IPv6 prefix>***
(Optional) Specifies the new IPv4 subnet mask or IPv6 prefix to with the new service IP. If **-mask** is specified, **-gw** must be specified.

-node ***node_name | node_id***
(Optional) Specifies the node ID or node name to modify with the command.

-techport ***enable | disable***
(Required) Specifies whether techport is enabled or disabled.

-force
(Optional) Forces in enabling the technician port. The parameter is not required for disabling the technician port.

Description

This command sets the service assistant IP address for a specific node. If the node is part of a system, then the system gateway, subnet, and prefix are used unless specified otherwise. If the node is a candidate node, the subnet, prefix, and gateway must be specified. If you specify an IPV4 or IPV6 address, but do not provide a gateway, mask, or prefix, then the existing gateway, mask, and prefix values are preserved.

An invocation example

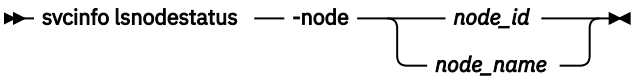
```
chnodeserviceip -serviceip 1.2.3.4 -gw 1.2.3.1 -mask 255.255.255.0 -node 6
```

```
No feedback.
```

svcinfo lsnodestatus

Use the **svcinfo lsnodestatus** command to view the status of a node.

Syntax



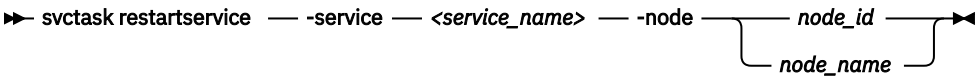
Parameters

-node *node_id | node_name*
(Required) Specifies the node ID or node name.

svctask restartservice

Use the `svctask restartservice` command to restart a named service. This command enables security administrator and administrator to restart a service.

Syntax



Parameters

-service *<service_name>*
(Required) Specifies the name of the service that you want to restart. The following services are supported:

- syslogd
- cron
- sshd
- java_tomcat
- nginx
- keymgr
- cloudd
- cfrest
- rscd
- tshistd
- ec_rngdaemon
- orchestrator
- sgcd
- reclaim_se

-node
Specifies the node ID or node name.

An invocation example

```
svctask restartservice -service java_tomcat -node node14
```

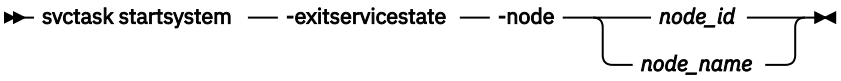
The resulting output:

```
No feedback.
```

svctask startsystem

Use the `svctask startsystem` command to take a single node out of service state. This command only enables security administrator and administrator to start a node.

Syntax



Parameters

-exitservicestate
(Required) Takes the specified node out of service state. The node must be a member of the cluster.

-node *node_name | node_id*
(Required) Specifies the node that you want to start. You can specify one of the following values:

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

An invocation example

```
svctask startsystem -node node42 -exitservicestate
```

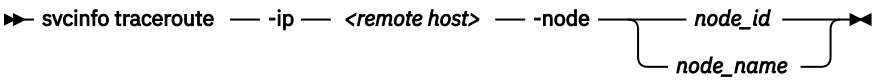
The resulting output:

```
No feedback.
```

svcinfo traceroute

Use the **svcinfo traceroute** command to print the wide area network (WAN) route packets to a specified host.

Syntax



Parameters

-ip *<remote host>*
(Required) Specifies the system hostname. The destination address can either be IPv4 / IPv6 or a fully qualified domain name (FQDN).

-node ***node_id | node_name***
(Required) Specifies the node ID or node name from which to perform the traceroute.

Description

This command prints WAN print route packets to a specified host.

An invocation example

```
svcinfo traceroute -ip hymir
```

The resulting output:

```
traceroute to hymir (x.xx.xx.xx), 30 hops max, 60 byte packets
 1  x.xx.xx.x  0.349 ms  0.428 ms  0.508 ms
 2  x.xx.xx.xx  0.245 ms  0.224 ms  0.224 ms
```

Chapter 10. Clustered system diagnostic and service-aid commands

Use the clustered system diagnostic and service-aid commands to diagnose and find system problems.

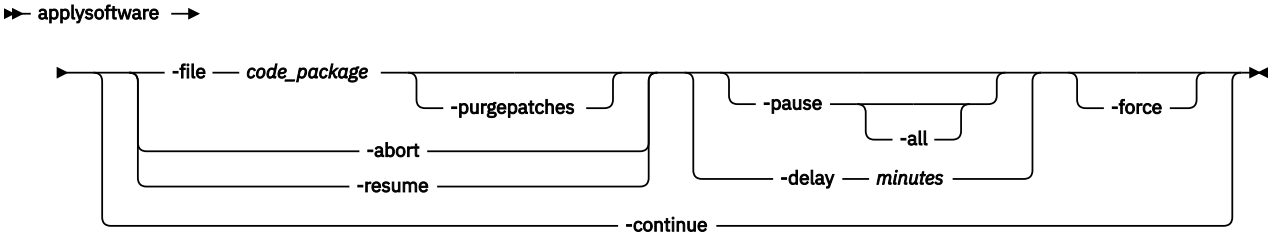
The system enables you to perform service activity, such as problem determination and repair activities, with a limited set of command-line tools. When you are logged in under the administrator role, all command-line activities are permitted. When you are logged in under the service role, only those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

applysoftware

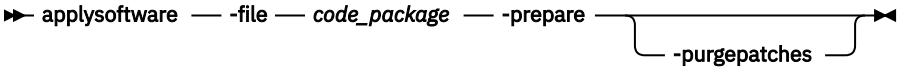
Use the **applysoftware** command to update the system to a new level of code or install a patch on the system.

Syntax

Syntax to start and manage an automatic code level update



Syntax to start a manual code level update



Syntax to install a patch



Parameters

-file *code_package*
(Required) Specifies the file name of the installation update package. Copy the update package onto the configuration node before you run the **applysoftware** command.

Note:

- Do not to specify a path if you use this parameter.

-purgepatches
When a code package is installed on a system, all the existing patches are reapplied. When you specify the **-purgepatches** parameter, all the existing patches on the node are removed.

-pause
(Optional) Specifies that the upgrade of a system is paused at the halfway point. This parameter must be specified with **-file** or **-resume**.

© Copyright IBM Corp. 2003, 2025

-all
(Optional) Specifies that the upgrade of a system is paused before each node goes offline for upgrade. This parameter must be specified with **-pause**, and **-file** or **-resume**.

-prepare
(Optional) Prepares the system for a manual code level update.

-delay ***minutes***
(Optional) Specifies that you can overwrite the default 10-minutes delay at the half way point of a software update. This parameter must not be specified with the **-pause**.

-abort
(Required for stopping an update) Specifies that a stalled or prepared update is stopped, returning the system to the original code level.

The **abort** parameter can also be used when the **lsupdate** command reports a status of:

- `prepare_failed`
- `prepared` (if all nodes are online)

-resume
(Required for resuming an update) Resumes a stalled automatic update, update cancel, or update completion process by retrying the step that stalled.

-enclosure ***enclosure_id***
(Optional) Specifies the enclosure to which to apply the software updates.

-force
(Optional) Specifies that the update or abort must proceed even if there is a lack of redundancy in the system. Disabling redundancy checking might cause loss of data, or loss of access to data. Use the **force** parameter with the **abort** parameter if one or more nodes are offline.

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

-continue
(Optional) Specifies that the concurrent upgrade continue.

patch_name

Specifies the name of the patch file that is uploaded on the configuration node.

Description

Note: The latest software upgrade test utility must be run before upgrading the software. To know more about this test, refer to Software Upgrade Test Utility.

The **applysoftware** command applies a new level of software or patches to the system. The command initiates the process that starts the installation of updated version and continues to the next stage of the upgrade that is automatic but pauses between nodes.

The **applysoftware** command cannot be used in service state, which means the system must be running for the command to be used and be successful. This command is synchronous and therefore reports success or failure. The actual update completes asynchronously.

The code package as specified by the file name must first be copied onto the current configuration node in the `/home/admin/update` directory; use the PuTTy secure copy (scp) application to copy the file.

If the **applysoftware** command is successful, the **lsupdate** command reports the status is prepared. If the **applysoftware** command fails, the **lsupdate** command reports the status is inactive.

If specified, the **prepare** parameter must succeed in order to successfully update. It is recommended to use the same package for the prepare as the actual update. The **prepare** parameter can be canceled by using the abort parameter (even after the system is prepared) as long as the **lsupdate** command reports the status as prepared.

Important: The **prepare** parameter might timeout. If a timeout occurs, the **prepare** parameter causes an asynchronous condition, and the **lsupdate** command reports the prepare status as `preparing`. If such scenario, wait until **lsupdate** reports the update as prepared before you proceed with the manual update process.

The command completes when the update process is successful. The command fails and the update package is deleted if:

- The specified package fails an integrity check due to corruption.
- Any node in the system has a hardware type that is not supported by the new code.
- The new code level does not support updates from the currently installed code.
- The code level of a remote system is incompatible with the new code.
- There are volumes that are dependent on the status of a node.

  Note: The **force** parameter can be used to override these scenarios if you are prepared to lose access to data during the update. Before proceeding, use the **lsdependentvdisks** command with the **node** parameter to list the node-dependent volumes at the time the command is run. If the command returns an error, move the quorum disks to MDisks that are accessible through all nodes. Rerun the command until the errors stop.

You can use the `lsdumps -prefix /home/admin/update` command to view the contents of the `/home/admin/update` directory.

An invocation example

The following command initiates an upgrade that runs to completion with no user intervention. The upgrade pauses for 10 minutes at the halfway point.

```
applysoftware -file INSTALL_6.4.0.0
```

The resulting output:

```
No feedback.
```

caterrlog (Deprecated)

The **caterrlog** command has been deprecated. Use the **lseventlog** command instead.

caterrlogbyseqnum (Deprecated)

The **caterrlogbyseqnum** command has been deprecated. Use the **lseventlog** command instead.

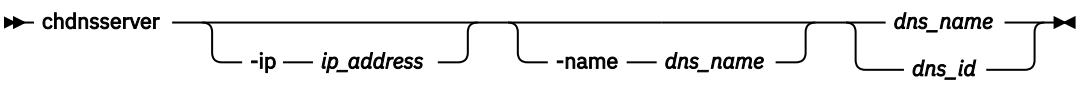
cherrstate (Deprecated)

The **cherrstate** command has been deprecated. Use the **cheventlog** command instead.

chdnsserver

Use the **chdnsserver** command to change a Domain Name System (DNS) server Internet Protocol (IP) address or name on a system.

Syntax



Parameters

-ip *ip_address*
(Optional) Specifies the DNS server IP address. The value must be in standard IPv4 or IPv6 format.

-name *DNS_name*
(Optional) Specifies a unique name for the system DNS server that is being changed.

dns_name | dns_id
(Required) Specifies DNS unique name or ID associated with the DNS server that is being changed.

Description

This command changes DNS server parameters (such as an IP address or name).

An invocation example

```
chdnsserver -ip 192.168.48.220 1
```

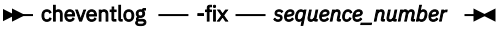
The resulting output:

```
No feedback
```

cheventlog

Use the **cheventlog** command to modify events in the event log.

Syntax



Parameters

-fix *sequence_number*
(Required) Mark an unfixed event as fixed.

Description

Important: You must specify the **-fix** parameter.

An invocation example to mark an event fixed

```
cheventlog -fix 120
```

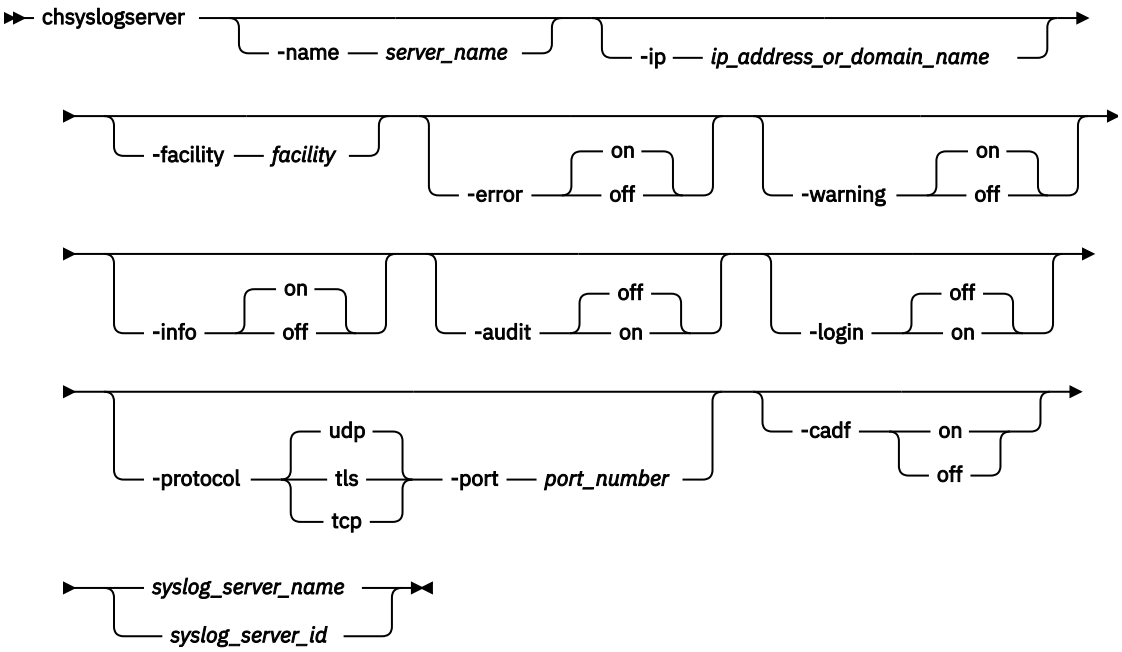
The resulting output:

```
No feedback
```

chsyslogserver

Use the **chsyslogserver** command to modify the parameters of an existing syslog server.

Syntax



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_or_domain_name*
(Optional) Specifies an IP address or a fully qualified domain name to assign to the syslog server. The IP address must be a valid IPv4 or IPv6 address or a fully qualified domain name. If you specify a domain name, a DNS server must be configured on your system. You can use the **mkdnsserver** command to configure DNS servers.

-facility *facility*
(Optional) Specifies a facility number to identify the origin of the message to the receiving server. Servers configured with facility values of `0` - 3 receive syslog messages in concise format. Servers configured with facility values of 4 - 7 receive syslog messages in fully-expanded format. This parameter is mutually-exclusive with **-cadf**.

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

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

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

-audit on | off
(Optional) Specifies whether the server receives CLI audit logs. The default value is `off`. These notifications are always sent as facility level 3 messages.

-login on | off
(Optional) Specifies whether the server receives authentication logs. The default value is `off`. These notifications are always sent as facility level 3 messages.

-protocol udp|tls|tcp
(Optional) Specifies the communication protocol that is used by this server. The default value is udp.

-port port_number
(Optional) Specifies the communication port that is used by this server. You cannot use this parameter unless **-protocol** is specified. This number must be in the range of 1 - 65535. The default value is 514 for udp and 6514 for `tcp`.

-cadf on | off
(Optional) Specifies that Cloud Auditing Data Federation (CADF) data reporting be turned on or `off`. Only cloud account create, update, removal and cloud backup activity syslog notifications sent to the server are formatted to the CADF standard. This parameter is mutually-exclusive with **-facility**.

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

Description

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

-login messages are sent to the authpriv facility on the remote syslog server and **-audit** messages are sent to facility.notice level (where facility is what's being selected when you create the syslog server and level is always `notice`).

For syslog servers, SNMP servers, and email users with type set to `local`, the valid combinations of notification types are:

- **-error**, **-warning**, and **-info**
- **-error** and **-warning**

If you disable CADF notifications for a syslog server that has CADF notification enabled, the facility value must be set to `0`.

An invocation example

```
chsyslogserver -facility 5 2
```

The resulting output:

```
No feedback
```

cleardumps

Use the **cleardumps** command to clear (or delete) the various dump directories on a specified node or node canister.

Syntax

Parameters

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

- /dumps (clears all files in all subdirectories)
- /dumps/cimom
- /dumps/cloud
- /dumps/configs
- /dumps/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 cleared.

Note: The following rules apply to the use of wildcards when using the CLI:

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

```
>cleardumps -prefix "/dumps/elogs/*.txt"
```

canister_id
(Optional) Specifies the canister to be cleared.

If you specify no canister to be cleared, the command executes on the configuration node rather than the canister whose ID is provided.

That optional parameter is just a number 1 or 2 (canister 1 or canister 2).

Description

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

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

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

You can list the contents of these directories on the given node or node canister by using the **lsxxxxdumps** commands.

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

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

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

An invocation example

```
cleardumps -prefix /dumps/configs
```

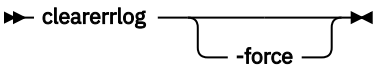
The resulting output:

```
No feedback
```

clearerrlog

Use the **clearerrlog** command to clear all entries from the event log including status events and any unfixed errors.

Syntax



Parameters

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

Description

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

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

An invocation example

```
clearerrlog -force
```

The resulting output:

```
No feedback
```

cpfabricdumps (Discontinued)

The **cpfabricdumps** command is discontinued. Use the **cpdumps** command instead.

dumperrlog

Use the **dumperrlog** command to dump the contents of the event log to a text file.

Syntax



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 system event log to a file using a system-supplied prefix of `errlog`, which includes the node ID and time stamp. When a file name prefix is provided, the same operation is performed but the details are stored in the dumps directory within a file with a name that starts with the specified prefix.

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

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

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

Note: The DMP family is printed at the end of all events in the event log.

An invocation example

```
dumperrlog -prefix testerrorlog
```

The resulting output:

```
No feedback
```

finderr

Use the **finderr** command to analyze the event log for the highest severity unfixed event.

Syntax



Parameters

None

Description

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

You can use this command to determine the order in which to fix the logged event.

An invocation example

```
finderr
```

The resulting output

```
Highest priority unfixed event code is [1010]
```

setevent (Discontinued)

Attention: The **setevent** command is discontinued. SNMP notification can be configured using the following commands: **svctask mksnmpserver**, **svctask chsnmpserver**, **svctask rmsnmpserver**, and **svcinfo lssnmpserver**.

lscimomdumps (Deprecated)

The **lscimomdumps** command is deprecated. Use the **lsdumps** command to display a list of files in a particular dumps directory.

lscopystatus

Use the **lscopystatus** command to determine whether any file copies are currently in progress.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

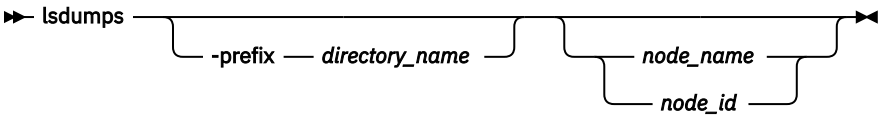
Description

This command displays an indicator that shows if a file copy is in progress. Only one file can be copied in the clustered system at a time.

lsdumps

Use the **lsdumps** command to display a list of files in a particular dumps directory on one of the nodes in the system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-prefix ***directory_name***
(Optional) Specifies the name of the directory to list files for. The default is the /dumps directory. Valid directory names:

- /dumps
- /dumps/audit
- /dumps/cimom
- /dumps/cloud
- /dumps/easytier
- /dumps/elogs
- /dumps/feature
- /dumps/iostats
- /dumps/iotrace
- /dumps/mdisk
- /home/admin/update
- /dumps/drive
- /dumps/enclosure
- /dumps/configs
- /dumps/debug
- /dumps/mdisk
- /dumps/syslogs
- /dumps/tejas

node_name | node_id
(Optional) Specifies the node ID or name to list the available dumps for. If you do not specify a node, the available dumps on the configuration node are listed.

Description

This command displays a list of files that are detected by a node. You can specify the name of the directory to list files for, and the node ID or name. If you do not specify a directory, the /dumps directory is used.

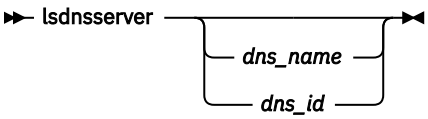
The files are listed in order of time that is created, with the oldest files listed first.

Use the **lsdumps** command with the optional **prefix** parameter to specify a directory. If you do not specify a directory, /dumps is used as the default. Use the optional *node_id_or_name* parameter to specify the node to list the available dumps. If you do not specify a node, the available dumps on the configuration node are listed.

lsdnsserver

Use the **lsdnsserver** command to list information for any Domain Name System (DNS) servers in the system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

dns_name | dns_id
(Optional) Specifies the DNS server name or ID for which to display details. The value for the ID must be a number and the value for the name must be an alphanumeric string.

Description

This command lists information for any DNS servers in the system.

This value provides the attribute values that can be displayed as output view data.

Table 67. ***lsdnsserver*** *output*

Attribute	Description
id	Indicates the DNS server ID. The value must be a number.
name	Indicates the DNS server name. The value must be an alphanumeric string.
type	Indicates the DNS server Internet Protocol (IP) address type. The value must be a standard IPv4 or IPv6 address.
IP_address	Indicates the IP address of the DNS server. The value must be a standard IPv4 or IPv6 address.
status	Indicates the status of the DNS server. Possible values are: • active • failed • untried • reachable • unresponsive

lserrlogbyfcconsistgrp (Deprecated)

The **lserrlogbyfcconsistgrp** command has been deprecated. Use the **lseventlog** command instead.

lserrlogbyfcmap (Deprecated)

The **lserrlogbyfcmap** command has been deprecated. Use the **lseventlog** command instead.

lserrlogbyhost (Deprecated)

The **lserrlogbyhost** command has been deprecated. Use the **lseventlog** command instead.

lserrlogbyiogrp (Deprecated)

The **lserrlogbyiogrp** command has been deprecated. Use the **lseventlog** command instead.

lserrlogbymdisk (Deprecated)

The **lserrorlogbymdisk** command has been deprecated. Use the **lseventlog** command instead.

lserrlogbymdiskgrp (Deprecated)

The **lserrlogbymdiskgrp** command has been deprecated. Use the **lseventlog** command instead.

lserrlogbynode (Deprecated)

The **lserrlogbynode** command has been deprecated. Use the **lseventlog** command instead.

lserrlogbyrcconsistgrp (Deprecated)

Attention: The **lserrlogbyrcconsistgrp** command has been deprecated. Use the **lseventlog** command instead.

lserrlogbyrcrelationship (Deprecated)

The **lserrlogbyrcrelationship** command has been deprecated. Use the **lseventlog** command instead.

lserrlogbyvdisk (Deprecated)

The **svcinfo lserrlogbyvdisk** command has been deprecated. Use the **svcinfo lseventlog** command instead.

lserrlogdumps (Deprecated)

The **lserrlogdumps** command is deprecated. Use the **lsdumps** command to display a list of files in a particular dumps directory.

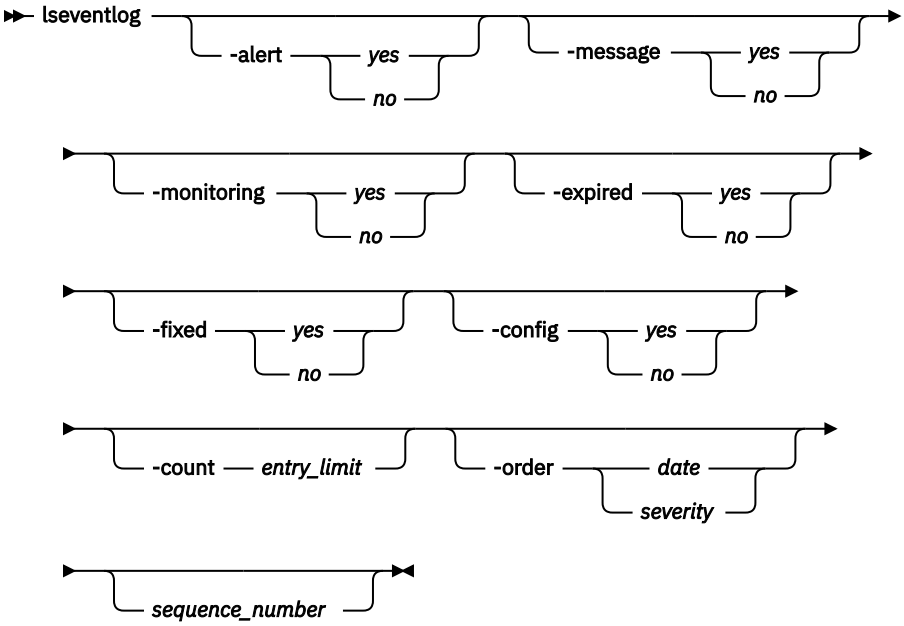
lsfeaturedumps (Deprecated)

The **lsfeaturedumps** command is deprecated. Use the **lsdumps** command to display a list of files in a particular dumps directory.

lseventlog

Use the **lseventlog** command to display a concise view of the system event log, or a detailed view of one entry from the log.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-alert *yes | no*
(Optional) Includes (or excludes) events with *alert* status.

-message *yes | no*
(Optional) Includes events with *message* status.

-monitoring *yes | no*
(Optional) Includes events with *monitoring* status.

-expired *yes | no*
(Optional) Includes (or excludes) events with *expired* status.

-fixed *yes | no*
(Optional) Includes (or excludes) events with *fixed* status.

-config *yes | no*
(Optional) Includes (or excludes) events with *config* status.

-count *entry_limit*
(Optional) Indicates the maximum number of events to display.

-order ***date | severity***
(Optional) Indicates what order the events must be in. Ordering by *date* displays the oldest events first. Ordering by severity displays the events with the highest severity first. If multiple events have the same severity, then they are ordered by date, with the oldest event displayed first.

The following list shows the order of severity, starting with the most severe:

1. Unfixed alerts (sorted by error code; the lowest error code has the highest severity)
2. Unfixed messages
3. Monitoring events (sorted by error code; the lowest error code has the highest severity)
4. Expired events
5. Fixed alerts and messages

sequence_number
(Optional) Indicates whether the command must display a full view of the event.

Description

This command displays a concise view of the system event log, or a detailed view of one entry from the log. You can sort the events and entries by severity or age.

The default values for included events are:

- alert=yes
- config=no
- expired=no
- fixed=no
- message=yes
- monitoring=no

This table provides the attribute values that can be displayed as output view data.

Table 68. ***lseventlog*** *output*

Attribute	Description	Value
machine_type	Indicates the node machine type and model number.	The value is an alphanumeric string up to 7 characters long.
serial number	Indicates the node serial number.	The value is an alphanumeric string up to 7 characters long.
sequence_number	Indicates the sequence number of the event.	The value is numeric from `0` to `8000000`.
first_timestamp	Indicates when the event was added to the log.	The value is in the format `YYMMDDHHMMSS`.
first_timestamp_epoch	Indicates when the event was added to the log (in seconds) after the epoch occurs.	The value is a numeric 32-bit value.
last_timestamp	Indicates when the event was most recently updated.	The value is in the format `YYMMDDHHMMSS`.
last_timestamp_epoch	Indicates the most recent update (in seconds) after an epoch for an event.	The value is a numeric 32-bit value.
fixed_timestamp	Indicates the time stamp when event is fixed.	The value is in the format `YYMMDDHHMMSS`.

*Table 68. **lseventlog** output (continued)*		
Attribute	**Description**	**Value**
fixed_timestamp_epoch	Indicates the time stamp (in seconds) when an event is fixed after an epoch occurs.	The value is a numeric string.
fru	Indicates the field-replaceable unit (FRU) for error or event; this field contains probable FRUs (separated by commas).	The value is an ASCII string up to 255 characters long.
object_type	Indicates the type of the object the event is logged against.	The values are: • `mdisk` • `mdiskgrp` • `volume` (or volume copy) • `node` • `host` • `io_grp` (`iogroup` in **dumperrlog**) • `fc_consist_grp` (`fcgrp`in **dumperrlog**) • `rc_consist_grp`(`rcgrp` in **dumperrlog**) • `fc_map` (`fcmap` in **dumperrlog**; *flash* in **caterrlog**) • `rc_relationship` (`rcmap`in **dumperrlog**; *remote* in **caterrlog**) • `cluster` • `controller` (`device`in **caterrlog** and **dumperrlog**) • `quorum` • `migrate` • `email_server` (*emailserver* in **caterrlog** and **dumperrlog**) • `enclosure` • `drive`
object_id	Indicates the ID of the object the event is logged against.	The value is a numeric 64-bit value. It is displayed in decimal for all object types other than clustered systems. For a clustered system, this value is hexadecimal but is blank for events with `cluster` object types.

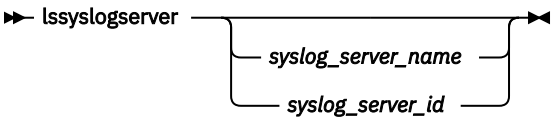
*Table 68. **lseventlog** output (continued)*		
Attribute	**Description**	**Value**
object_name	Indicates the name of the object the event is logged against.	This value is the object name format and is blank if the object was deleted or does not have a name.
copy_id	Indicates the volume copy ID the event is logged against.	The value is a numeric value `0` to `1`; it is blank if not a `vdiskcopy` event.
reporting_node_id	Indicates the ID of the node that reported the event.	The value is a numeric 64-bit value that is blank if the event is reported by the clustered system.
reporting_node_name	Indicates the name of the node that reported the event.	This value is the object name format and is blank if node is deleted or event is reported by the clustered system.
root_sequence_number	Indicates the sequence number of the root or causal event **Important:** If the event is directly caused by another event, the sequence_number of the related event is shown here.	The value is a numeric value from `1` to `8000000`; blank if there is no root or if the event is not directly caused by another event.
event_count	Indicates the number of reported events that are combined into this event	The value is a numeric 32-bit value.
status	Indicates the event category.	The values are: • `alert` • `message` • `monitoring` • `expired`
fixed	Indicates whether the event was marked fixed (for an alert) or read (for a message).	The values are: • `yes` • no (for events that cannot be fixed, or are not fixed)
auto_fixed	Indicates whether event is marked fixed by the code.	The values are: • `yes` • no (for events that cannot be fixed, or are not fixed)
notification_type	Indicates the type of event notification.	The values are: • `error` • `warning` • `informational` • `none`

Table 68. **lseventlog** output (continued)		
Attribute	**Description**	**Value**
event_id	Indicates the event ID.	The value is a 6-digit numeric value.
event_id_text	Indicates the description that is associated with the event ID.	The value is a text value with a maximum of 200 bytes. This appears in CLI requested language.
error_code	Indicates the error code that is associated with this event.	The value is a 4-digit numeric value but is blank if there is no error code.
error_code_text	Indicates the description that is associated with the error code.	The value is a text value with a maximum of 200 bytes that is blank if there is no error code. This value appears in the language that is requested by the CLI.
description	Indicates the description that is associated with the event. If the event has an error code, this value is the same as the `error_code_text` field; otherwise, it is the same as the `event_id_text` field	Text (maximum of 200 bytes).
callhome_type	Indicates the type of Call Home that is associated with the event.	The values are `hardware`, `software`, and `none`.
sense1	Indicates the sixteen bytes of hex-encoded sense data; least significant byte is on the left.	The value is 16 2-character hexadecimal numbers that are separated by spaces.
sense2		
sense3		
sense4		
sense5		
sense6		
sense7		
sense8		

lssyslogserver

Use the **lssyslogserver** command to return a concise list or a detailed view of syslog servers that are configured on the clustered system.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

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 clustered system.

This table provides the attribute values that can be displayed as output view data.

Table 69. ***lssyslogserver*** *output*

Attribute	Description
id	Indicates the syslog server ID. The value must be a number.
name	Indicates the syslog server name. The value must be an alphanumeric string.
IP_address	Indicates the syslog server Internet Protocol (IP) address. The value must be a valid IP address.
facility	Indicates the syslog server facility value. The value must be a number 0 - 7 but is blank for a CADF notification-enabled server.
error	Indicates whether error messages are on. The values are on or off.
warning	Indicates whether warning messages are on. The values are on or off.
info	Indicates whether informational messages are on. The values are on or off.
cadf	Indicates whether CADF data reporting is enabled or not for the syslog server. The values are on or off.
audit	Indicates whether this server receives CLI audit logs. The values are on or off (default).
login	Indicates whether this server receives authentication logs. The values are on or off (default).
protocol	Indicates the communication protocol that is used by this server (or blank for default). The values are tcp or udp (default).
port	Indicates the communication port that is used by this server (or blank for default). The value is an integer in the range of 1 - 65535. A value of 514 indicates UDP protocol and a value of 6514 indicates TCP protocol.

lssoftwaredumps (Deprecated)

The **lssoftwaredumps** command is deprecated. Use the **lsdumps** command to display a list of files in a particular dumps directory.

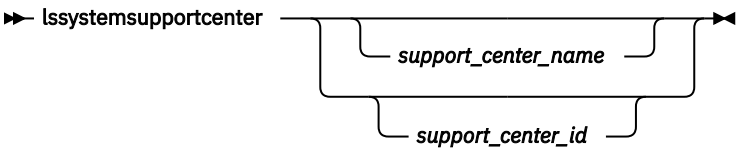
lssoftwareupgradestatus (Deprecated)

The **lssoftwareupgradestatus** command is deprecated. Use the **lsupdate** command instead.

lssystemsupportcenter

Use the **lssystemsupportcenter** command to list details about remote support servers.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

support_center_name* | *support_center_id
(Optional) Specifies a name or ID for a remote support server in the server index. The value for the ID must be a number (integer) and the value for the name must be an alphanumeric string. This parameter displays a full view of any configured name or ID values.

Description

This command lists details about remote support servers.

This command returns a concise or detailed view of the remote support servers defined.

This table provides the attribute values that can be displayed as output view data.

Table 70. ***lssystemsupportcenter*** *output*

Attribute	Description
id	Indicates the support center or proxy server ID. The value must be a number (integer).
name	Indicates the support center or proxy server name. The value must be an alphanumeric string.
IP_address	Indicates the Internet Protocol Version 4 (IPv4) or Version 6 (IPv6) address for the new support center or proxy server. The value must be a valid IPv4 or IPv6 address.
port	Indicates the port number for the configured support center or the proxy server. The value must be a number (integer).
proxy	Indicates that the target server is a proxy server (and not the support center). The values are yes or no.

lsupdate

Use the **lsupdate** command to display a system's machine code (code) upgrade status.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

Description

The **applysoftware** command is used to update the system code. The following are the upgrade status states reported by the **lsupdate** command:

status
Indicates the overall update-related status of the system. The values are:

- `success`, which indicates that all updating activity is complete.
- `system_preparing`, which indicates that the system is preparing a manual update.
- `system_prepared`, which indicates that the system is ready to starting a manual update.
- `system_prepare_failed`, which indicates that the system could not start a manual update.

  Note: Check the event log.
- `system_initializing`, which indicates that the system is readying nodes for an update.
- `system_updating`, which indicates that the nodes are being updated.
- `system_updating_pausing`, which indicates that the system is pausing before it continues to update the nodes.

  Note: If you specify `applysoftware -continue` the status changes to `system_updating`.
- `system_committing`, which indicates that all nodes are updated and the system is readying the new code.
- `system_stalled`, which indicates that an update is stalled because of unexpected node problems.
- `system_stalled_non_redundant`, which indicates that an update is stalled because of dependent volumes.
- `system_stalled_during_committing`, which indicates that a T2 occurred during the switching to the new CSM and that the system moved into a stalled state in order to abort and roll back the upgrade.
- `system_restoring`, which indicates that a stalled update is canceled by the user.

  Note: The previous code version is being restored.
- `system_restoring_pausing`, which indicates that the system is pausing before it continues to restore the nodes.

  Note: If you specify `applysoftware -continue` the status changes to `system_updating_restoring`.
- `system_restoring_stalled_non_redundant`, which indicates that an update is canceled and then stalled because of dependent volumes.
- `system_manual_update`, which indicates that a manual update is in progress.

- `system_completion_required`, which indicates that all nodes are updated and management functions are available, but further system changes are necessary to complete the update.

  Note: Check the event log.
- `system_completing`, which indicates that an automatic update completion is in progress.
- `system_completing_pausing`, which indicates that automatic update completion is paused.
- `system_completing_paced`, which indicates that a paced update completion is in progress.
- `system_completing_stalled`, which indicates that an automatic update completion stalled because of an unexpected problem.
- `enclosures`, which indicates that enclosure firmware is being updated.
- `enclosures_stalled`, which indicates that an enclosure firmware update stalled because of an unexpected problem or the enclosure has a lack of redundancy.
- `enclosures_restoring`, which indicates a stalled update has been canceled by the user. The previous code version is being restored to the enclosure canisters.
- `drives`, which indicates that drive firmware is being updated.

event_sequence_number
Indicates an event that describes any current problem with the code update. The value must be a numeric string in decimal format (or blank).

progress
Indicates the completion percentage of the current update activity in terms of number of objects updated (rather than time elapsed). The value must be a numeric string (decimal) in the range 0 - 100.

estimated_completion_time
Indicates estimated completion time of current update activity. It is valid only if the current update activity is automatic and is not stalled. The value must be in the format YYMMDDHHMMSS (or blank).

suggested_action
Indicates the actions that help the update progress. The value must be:

- `cancel`, which indicates to abort and roll back the CSM upgrade after a T2 occurs and the system is in a stalled state.
- `complete`, which indicates the system update is complete and update completion must be issued. Nodes are online.
- `continue`, which indicates that the concurrent upgrade is paused and you must `applysoftware -continue` to continue the concurrent upgrade.
- `fix`, which indicates that an update cannot continue because of a problem. Check the event log, specifically the **event_sequence_number** output. Nodes are offline.
- `manual`, which indicates that a manual update is in progress.
- `pacednext`, which indicates that the paced update is in progress and the next node must be scheduled for updating. Nodes are online.
- `resume_cancel`, which indicates that the update is stalled but can be resumed or canceled. Nodes are online.
- `resume`, which indicates that the update completion is stalled but can be resumed. Nodes are online.
- `start`, which indicates that the system is ready for a new update to start. No update is in progress or prepared and all nodes are online.
- `wait`, which indicates that the system is busy (no action is required) because an update is in progress.

system_new_code_level
Indicates that a new level of code is being updated. The value must be the build version (or blank if not updating or restoring the system).

system_forced
Indicates any current node-related activity in forced mode (ignoring dependent volumes). The values are `yes` or `no`.

system_next_node_status
Indicates the status of the next node in the current node-related update activity. The values are:

- `none`, which indicates that there is no node to update.
- `paused`, which indicates that the current node is paused during a concurrent upgrade, and you must `applysoftware -continue` to continue the concurrent upgrade.
- `waiting`, which indicates that the node is ready for updating and that the system is waiting (typically for multipathing failover).
- `ready`, which indicates that the node is ready for updating, and the update activity is paced. You must start the update manually.
- `updating`, which indicates that the node is updating.
- `stalled`, which indicates that the node is going to be updated next, but the update is stalled.

system_next_node_time
Indicates the time that the next node update will start. It is valid only if the **system_next_node_status** is `waiting`. The value must be in the format YYMMDDHHMMSS x (or blank).

system_next_node_id
Indicates the ID of the next node in the current node-related update. The value must be a numeric string (or blank).

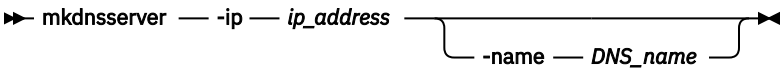
system_next_node_name
Indicates the name of the next node in the current node-related update. The value must be an alphanumeric string (or blank).

mkdnsserver

Use the **mkdnsserver** command to configure a new Domain Name System (DNS) server for a system.

Syntax

Note: This command is not supported by FlashSystem 840 AE1.



Parameters

-ip ***ip_address***
(Required) Specifies the DNS server Internet Protocol (IP) address. The value must be in standard IPv4 or IPv6 format. Depending on the format you specify the system will validate the format to make sure it is correct.

-name ***DNS_name***
(Optional) Specifies a unique name for the system DNS server being created. If a DNS server name is not specified, a unique name is generated and then assigned to the DNS server.

Description

This command configures a new DNS server for a system.

Use the **-ip** parameter to specify the DNS server IP address. You can enter any valid IPv4 or IPv6 address. The system validates the format of specified IP address to make sure that it is correct.

An invocation example

```
mkdnsserver -ip 192.168.44.34
```

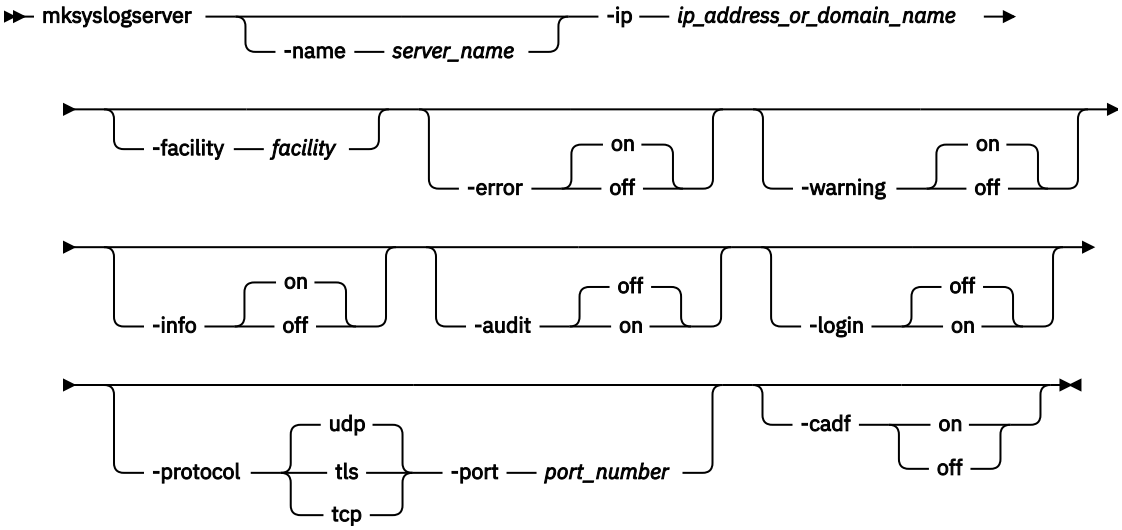
The resulting output:

```
DNS Server id [0] successfully created
```

mksyslogserver

Use the **mksyslogserver** command to create a syslog server to receive notifications.

Syntax



Parameters

-name *server_name*
(Optional) Specifies a unique name to assign to the syslog server. If a name is not specified, then a system default of `syslog`*n* is applied, where *n* is the ID of the server. When specifying a server name, `syslog` is a reserved word.

-ip *ip_address_or_domain_name*
(Required) Specifies the Internet Protocol (IP) address or domain name of the syslog server. The IP address must be a valid IPv4 or IPv6 address, or a fully qualified domain name. If you specify a domain name, a DNS server must be configured on your system. You can use the **mkdnsserver** command to configure DNS servers.

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

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

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

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

-audit on | off
(Optional) Specifies whether the server receives CLI audit logs. The default value is `off`.

-login on | off
(Optional) Specifies whether the server receives authentication logs. The default value is `off`.

-protocol udp |tls|tcp
(Optional) Specifies the communication protocol that is used by this server. The default value is udp.

-port port_number
(Optional) Specifies the communication port that is used by this server. You cannot use this parameter unless **-protocol** is specified. This number must be in the range of 1 - 65535. The default value is 514 for udp and 6514 for `tcp`.

-cadf on | off
(Optional) Specifies that Cloud Auditing Data Federation (CADF) data reporting be turned on or `off`. Only cloud account create, update, removal and cloud backup activity syslog notifications sent to the server are formatted to the CADF standard. This parameter is mutually-exclusive with **-facility**.

Description

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

-login messages are sent to the authpriv facility on the remote syslog server and **-audit** messages are sent to facility.notice level (where facility is what's being selected when you create the syslog server and level is always `notice`).

For syslog servers, SNMP servers, and email users with type set to `local`, the valid combinations of notification types are:

- `error`, `warning`, and `info`
- `error` and `warning`

An invocation example

```
mksyslogserver -ip 1.2.3.4
```

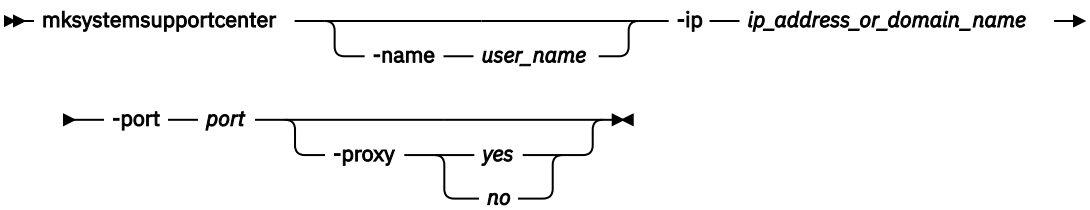
The resulting output:

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

mksystemsupportcenter

Use the **mksystemsupportcenter** command to add a support center or Remote Support Proxy server to your remote support configuration.

Syntax



Parameters

-name *user_name*
(Optional) Specifies the unique name of the support center or proxy to be defined. If a center with that name is defined, the command fails. The value must be an alphanumeric string that:

- Cannot start with the string `default_support_center`
- Cannot start with a hyphen or number
- Cannot begin or end with a space
- Must be 1 - 64 characters long (using numbers, letters, spaces, periods, or an underscore)

Additionally, the names `SupportCenter` and `Proxy` cannot be used as shown. Do not specify these names unless you use all lower-case letters - for example, `supportcenter` or `proxy`.

Note: If you do not specify a name, a system default of `supportservern` is used, where n is the object index.

-ip *ip_address_or_domain_name*
(Required) Indicates the IP address or fully qualified domain name for the new support center or Remote Support Proxy server. For IP addresses, the value must be a valid IPv4 or IPv6 address. If you specify a domain name, a DNS server must be configured on your system. You can use the **mkdnsserver** command to configure DNS servers.

Note:

- Remote Support Proxy servers are deprecated. You can configure HTTP proxy servers instead by using **mkproxy** or **chproxy** command. You can configure your firewall to allow traffic to pass directly from the system or you can route traffic through an HTTP proxy server within your environment. Do not configure an HTTP proxy server by using the **mksystemsupportcenter** command. For more information, see HTTP proxy server.
- With the addition of the HTTP proxy support, Remote Support Proxy servers are no longer necessary, but they are still supported for existing configurations. Alternately, a Remote Support Proxy can be configured to consolidate firewall traffic from a number of storage systems. Remote upgrades cannot be completed through the Remote Support Proxy server.

-port *port*
(Required) Specifies the port number for the new connection. The value must be a number in the range 1 - 65535. The default value is 1025.

-proxy *yes | no*
(Optional) Specifies whether a target server is a Remote Support Proxy server (and not a support center). The values are `yes` or `no`.

Description

This command creates a support center or Remote Support Proxy server to your remote support configuration. The maximum number of Remote Support Proxy servers that you can define is six.

If you configure a Remote Support Proxy server, remote support assistance is available only by using Remote Support Proxy server (which means that no direct connections are tried). If you configure multiple support centers (or proxies), the system cycles through each support center before it enables remote support assistance which means that it tries all support centers at least three times before it ends in `failure` state. Use the set of default support centers that are configured if needed.

Important: Unless your support team recommends it, do not configure any new support centers. However, a Remote Support Proxy server can be configured if required. If you are routing Remote Support connections of a system with all the nodes configured with a service IPv6 address, then you must use a Remote Support Proxy server. The Remote Support Proxy server should listen on an IPv6 address. It can additionally listen on IPv4 address. This Remote Support Proxy server IPv6 address needs to be specified on your system using this command (with the -proxy parameter).

For remote access, port 443 is suggested as it provides more security. If the connection from port 443 remote access is unsuccessful, then port 22 is used as a backup. When you configure the firewall rules with port 443, the fallback of remote access from port 443 to port 22 is not applicable.

An invocation example

```
mkemailserver -ip servername.example.com -port 1025
```

The resulting output:

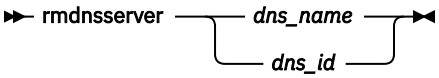
```
Emailserver id [2] successfully created
```

rmdnsserver

Use the **rmdnsserver** command to remove a Domain Name System (DNS) server from a system.

Syntax

Note: This command is not supported by FlashSystem 840 AE1.



Parameters

dns_name | dns_id
(Required) Specifies the ID or name of the Domain Name System (DNS) server to be removed from the system. The value must be a number for the ID and an alphanumeric string for the name.

Description

This command removes a DNS server from a system.

Note: The last DNS server cannot be removed if a named proxy server is in use.

An invocation example

```
rmdnsserver 1
```

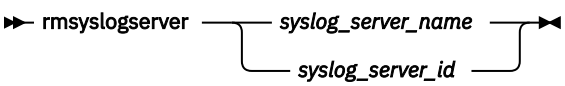
The resulting output:

```
No feedback
```

rmsyslogserver

Use the **rmsyslogserver** command to delete the specified syslog server.

Syntax



Parameters

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

Description

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

An invocation example

```
rmsyslogserver 2
```

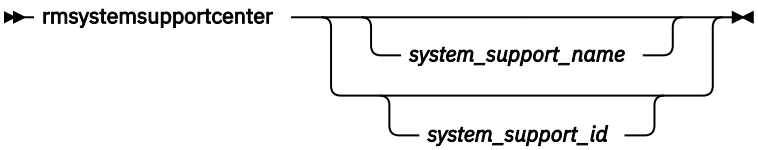
The resulting output (if the command is successful):

```
No feedback
```

rmsystemsupportcenter

Use the **rmsystemsupportcenter** command to delete a configured support center or proxy server.

Syntax



Parameters

system_support_name | system_support_id
(Required) Specifies a name or ID for a remote support server in the server index that is to be removed. The value for the ID must be a number (integer) and the value for the name must be an alphanumeric string. The name or ID must correspond to the entry that is displayed when you specify **lssystemsupportcenter**.

Description

This command deletes a configured support center or proxy server.

Note: You cannot remove default support centers that are automatically configured for the system.

An invocation example

```
rmsystemsupportcenter 0
```

The resulting output:

```
No feedback
```

setlocale

Use the **setlocale** command to change the locale setting for the system. It also changes command output to the chosen language.

Syntax



Parameters

-locale *locale_id*
(Required) Specifies the locale ID. The value must be a numeric value depending on the wanted language. The list in the command description shows the possible values.

Description

This command changes the language in which error messages are displayed as output from the command-line interface. Then, 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).

Specify the **setlocale** command to change the locale setting for the system; all interface output is changed to the chosen language. For example, to change the language to Japanese, type the following on the command line:

```
setlocale -locale 3
```

Where 3 is the value for Japanese. The following values are supported:

- 0 US English (default)
- 1 Simplified Chinese
- 2 Traditional Chinese
- 3 Japanese
- 4 French
- 5 German
- 6 Italian
- 7 Spanish
- 8 Korean
- 9 Portuguese (Brazilian)
- 10 Russian

An invocation example (where 3 is Japanese)

```
setlocale -locale 3
```

The resulting output:

```
No feedback
```

svqueryclock

Use the **svqueryclock** command to return the date, time, and current time-zone of the system.

Syntax



Parameters

None

Description

This command returns the date, time and current time-zone of the system.

An invocation example

```
svqueryclock
```

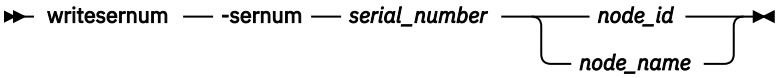
The resulting output:

```
Mon Nov 25 14:59:28 GMT 2013
```

writesernum

Use the **writesernum** command to write the node serial number into the planar NVRAM.

Syntax



Parameters

-sernum *serial_number*
(Required) Specifies the serial number to write to the nonvolatile memory of the system planar.

node_id | node_name
(Required) Specifies the node where the system planar is located. The serial number is written to this system planar. This name is not the worldwide node name (WWNN).

Description

This command writes the node serial number into the planar NVRAM and then reboots the system. You can find the serial number at the front of the node without having to remove it from the rack. The seven-digit alphanumeric serial number is located on a label on the front of the node. The serial number on the label might contain a hyphen. Omit this hyphen when typing the serial number with the **writesernum** command.

Note: Once you have written the serial number to the planar NVRAM, you can issue the **lsnodevpd** command to verify that the number is correct. The `system_serial_number` field contains the serial number.

An invocation example

```
writesernum -sernum 1300027 node1
```

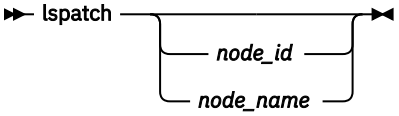
The resulting output:

```
No feedback
```

lspatch

Use the **lspatch** command to view a list of all the patches on a specific node in the system.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

Parameters

node_id | node_name
(Optional) Specifies the node ID or node name of the canister.

Note: If node ID or node name is not provided, the patch details of configuration node are displayed.

Description

Use node ID or node name to get list of patches of a particular node. This table shows the possible outputs.

Table 71. ***lspatch*** *output*

Attribute	Values
id	Indicates the number of the patch in the total count of all patches.
node_id	Indicates the node canister ID.
node_name	Indicates the name of the node or node canister.
panel_name	Indicates the panel name of the node.
patch_name	Indicates the name of the patch.
patch_status	Indicates the installation state of the patch: • `Installed` indicates that the patch is installed on the system. • `Obsolete` indicates that the patch was installed in the previous update and the new update does not require it. • `Error` indicates that the patch failed to install.
auto_update	Indicates whether the Patch Auto Update process is enabled or disabled for this patch.

lssystempatches

Use the **lssystempatches** command to display a list of all system patches.

Syntax

Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

Description

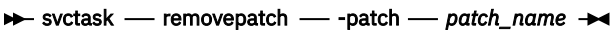
Use this command to verify that the system patches are installed on all the nodes in the system.

*Table 72. **lssystempatches** output*	
Attribute	**Values**
id	Indicates the number of the patch in the total count of all patches.
node_id	Indicates the node canister ID.
node_name	Indicates the name of the node or node canister.
panel_name	Indicates the panel name of the node.
patch_name	Indicates the name of the patch.
patch_status	Indicates the installation state of the patch: • `Installed` indicates that the patch is installed on the system. • `Obsolete` indicates that the patch was installed in the previous update and the new update does not require it. • `Error` indicates that the patch failed to install.
auto_update	Indicates whether the Patch Auto Update process is enabled or disabled for this patch.

svctask removepatch

Use the **svctask removepatch** command to remove a specified patch from all nodes in the system.

Syntax



Parameters

-patch *patch_name*
(Required) Specifies the name of the patch file.

An invocation example

```
svctask removepatch -patch patch_name
```

The resulting output:

```
No feedback.
```

svctask purgepatches

Use the **svctask purgepatches** command to remove all patches from all the nodes in the system.

Syntax



Parameters

-force
(Required) Specifies that the service action removes all installed patches from all clustered nodes.

An invocation example

```
svctask purgepatches -force
```

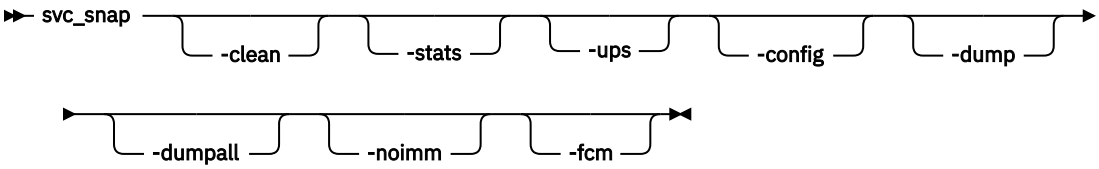
The resulting output:

```
No feedback.
```

svc_snap

Use the **svc_snap** command to create a snap file containing diagnostic information from the system.

Syntax



Parameters

-clean
(Optional) Erases all snap files on the configuration node.

-stats
(Optional) Gathers performance statistics files from all system nodes.

-ups
(Optional) Gathers the UPS logs from all system nodes.

-config
(Optional) Generates a new backup for the configuration node.

-dump
(Optional) Gathers the latest dump and/or livedump information from the configuration node.

-dumpall
(Optional) Gathers the latest dump and/or livedump information from every system node.

-noimm
(Optional) Suppresses the collection of the Integrated Management Module (IMM) dump.

-fcm
(Optional) Generates a new drive dump for each FCM in the system.

Description

This command creates a new snap file to collect diagnostic information from the system.

If this command is run from the command line, the snap file will be created in the /dumps directory on the configuration node.

An invocation example

```
svc_snap
```

The resulting output:

```
No feedback
```

Important: The name of the output file (placed on the configuration node) is `snap.nodeid.date.time.tgz`.

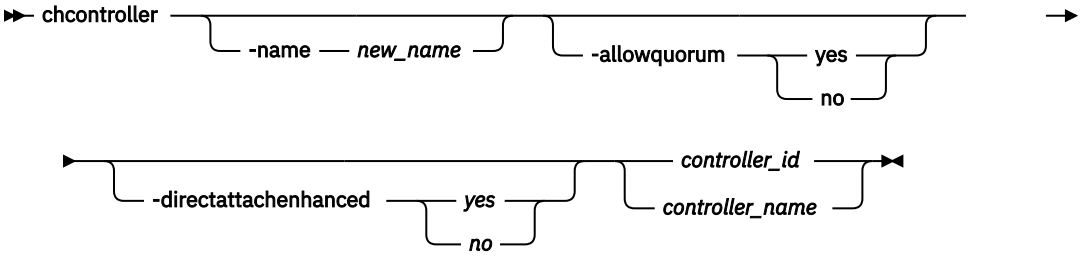
Chapter 11. Controller commands

Use the controller commands to modify the name of a storage controller.

chcontroller

Use the **chcontroller** command to modify the attributes of a controller.

Syntax



Parameters

-name *new_name*
(Optional) Specifies the new name to be assigned to the controller.

-allowquorum yes | no
(Optional) Specifies that the controller is allowed or is not allowed to support quorum disks. A value of yes enables a suitable controller to support quorum disks. A value of no disables a controller from supporting quorum disks, provided that the specified controller is not currently hosting a quorum disk.

-directattachenhanced yes | no
(Optional) Specifies that the controller has enhancements enabled for operation using a direct cable attachment between the system nodes and the controller. The value can be yes or no. The default value is set to no.

controller_id | controller_name
(Required) Specifies the controller to modify. Use either the controller name or the controller ID.

Description

This command changes the name of the controller that is specified by the *controller_id | controller_name* variable to the value that you specify with the **-name** parameter.

If any controller that is associated with an MDisk shows the **allow_quorum** attribute set to no with the **lscontroller** command, the set quorum action fails for that MDisk. Before using the **chcontroller** command to set the **-allowquorum** parameter to yes on any disk controller, check the following website to see whether the controller supports quorum.

www.ibm.com/support

You can add a new disk controller system to your SAN at any time. Follow the switch zoning guidelines in the section about switch zoning. Also, ensure that the controller is set up correctly for use with the system.

To add a new disk controller system to a running configuration, ensure that the system has detected the new storage MDisks by issuing the **detectmdisk** command. The controller has automatically been assigned a default name. If you are unsure of which controller is presenting the MDisks, issue the **lscontroller** command to list the controllers. The new controller is listed with the highest numbered default name. Record the controller name and follow the instructions in the section about determining a disk controller system name.

© Copyright IBM Corp. 2003, 2025

Give this controller a descriptive name by issuing the following command:

`chcontroller -name newname oldname`

List the unmanaged MDisks by issuing the following command:

`lsmdisk -filtervalue mode=unmanaged:controller_name=newname`

These MDisks correspond to the RAID arrays or partitions that you have created. Record the field controller LUN number. The field controller LUN number corresponds with the LUN number that you assigned to each of the arrays or partitions.

An invocation example

```
chcontroller -name newtwo 2
```

The resulting output:

```
No feedback
```

An invocation example

```
chcontroller -directattachenhanced yes
```

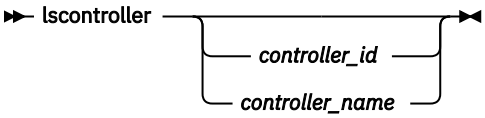
The resulting output:

```
No feedback
```

lscontroller

Use the **lscontroller** command to display a concise list or a detailed view of controllers that are visible to the system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

controller_id | controller_name
(Optional) Specifies the name or ID of a controller.

Description

This command returns a concise list, or a detailed view, of controllers visible to the system.

The following values are applicable to the data in the output views:

`degraded no, yes`

To differentiate the name of a storage controller from the name that is shown on the system, list the storage controllers by issuing the **lscontroller** command. Record the controller name or ID for the controller that you want to determine. For the controller in question, issue the **lscontroller** *controller name | id* command, where *controller name | id* is the controller name or ID. Record the worldwide node name (WWNN) for the controller. You can use the WWNN to determine the actual storage controller by launching the native controller user interface, or by using the command line tools that it provides to verify the actual controller that has the WWNN.

Notes:

1. The *mdisk_link_count* value is the number of MDisks currently associated with this storage controller.
2. The *max_mdisk_link_count* value is the highest value that the *mdisk_link_count* reaches since it was last reset to the *mdisk_link_count* value.

   Remember: This value is reset by specific maintenance procedures or when the event log is cleared.
3. A SAN connection from a node or node canister port to a controller port for a single MDisk is a path. The controller port *path_count* value is the number of paths that are currently being used to submit input/output (I/O) data to this controller port.
4. The storage controller *max_path_count* value is the highest value that the storage controller *path_count* reaches since it was last reset to the *path_count* value. This value is reset by specific maintenance procedures or when the system error log is cleared.

   Important: The *max_path_count* value is the highest value that the *path_count* reaches since it was last reset to the *path_count* value.

   Remember: This value is reset by specific maintenance procedures or when the event log is cleared.
5. The *allow_quorum* value identifies if the controller is enabled to support quorum disks. Quorum support is either enabled or disabled depending on the controller hardware type.
6. The *ctrl_s/n* value is the controller serial number.

   Important: This data comes from vendor-controlled sources and might not be available.

This table provides the attribute values that can be displayed as output view data.

*Table 73. **lscontroller** output*

Attribute	Possible Values
id	Indicates the controller ID.
name	Indicates the controller name.
WWNN	Indicates the worldwide node name (WWNN). This field is blank for iSCSI controllers.
mdisk_link_count	Indicates the MDisk link count.
max_mdisk_link_count	Indicates the maximum MDisk link count.
degraded	Indicates whether the controller has degraded MDisks.
vendor_id	Indicates the vendor identification name or number.
product_id_low	Indicates the product identification.
product_id_high	Indicates the product identification.
product_revision	Indicates the product revision.
ctrl_s/n	Indicates the controller serial number.

*Table 73. **lscontroller** output (continued)*	
Attribute	**Possible Values**
allow_quorum	Indicates that the controller can support quorum disks.
WWPN	Indicates the worldwide port name (WWPN). This field is blank for iSCSI controllers.
path_count	Indicates the number of paths that are currently being used to submit input/output (I/O) data to the controller port.
max_path_count	Indicates the maximum number of paths that are currently being used to submit input/output (I/O) data to the controller port.
site_id	Indicates the site value for the controller. This numeric value is 1, 2, 3, or blank.
site_name	Indicates the site name for the controller. This value is an alphanumeric value or is blank.
fabric_type	Indicates a Fibre Channel (FC) or SAS controller. • `fc` indicates an FC controller • `sas_direct` indicates an SAS direct-attached controller • `multiple` indicates multiple controllers (either FC, SAS, or both)
iscsi_port_id	Indicates the I/O port identifier, which is the same as the WWPN value from the FC domain. This value shows the iSCSI port ID for the iSCSI controller and is blank for other controllers. This value must be a numeric value. This ID refers to the row number in output from the **lsiscsistorageport** command. The **lsiscsistorageport** can be used to find the controller IQN.
ip	Indicates the IP address that is associated with the `iscsi_port_id`. This value shows the IP value for the iSCSI controller and is blank for other controllers. This value must be an IPv4 or IPv6 address.
physical_capacity	Indicates the physical capacity of the controller. This value is always blank or empty for controllers that do not report physical capacity information.
direct_attach_enhanced	Indicates whether the direct attach enhancement feature is enabled. The value can be `yes` or `no`. **Note:** This field is only displayed in the **lscontroller** detailed view.

lscontrollerports

Use the **lscontrollerports** command to list the controller port information for the IBM Storage Insights.

Syntax



Description

This table provides the attribute values that can be displayed as output view data.

*Table 74. **lscontrollerports** attribute values*

Attribute	Value
controller_id	Indicates the ID of the controller.
controller_name	Indicates the name of the controller.
port_id	Indicates the controller port ID.
WWPN	Indicates the worldwide port name (WWPN) of the fibre-channel port.
iscsi_port_id	Indicates the port ID of the iSCSI port.
ip	Indicates the user defined IP address (IPv4 or IPv6) for the iSCSI port.

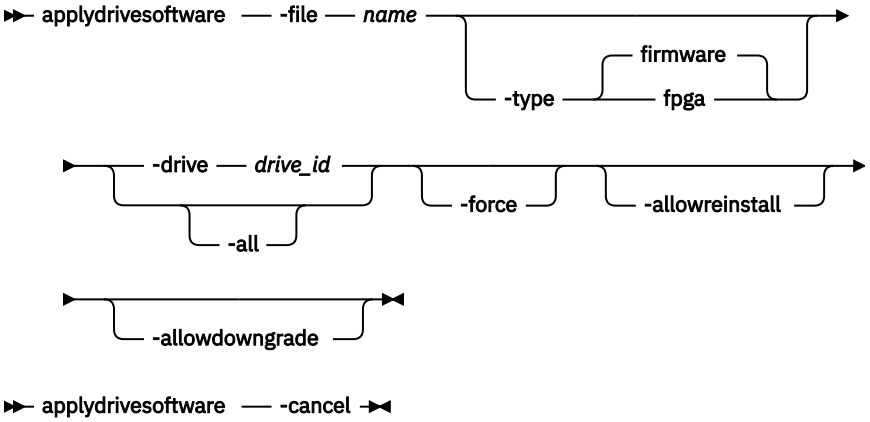
Chapter 12. Drive commands

Use the drive commands to capture information to assist with managing drives.

applydrivesoftware

Use the **applydrivesoftware** command to update drives.

Syntax



Parameters

-file *name*
(Required) Specifies the firmware update file name that must be copied to the `/home/admin/update/` directory on the configuration node.

-type fpga | firmware
(Optional) Specifies the type of drive firmware to update. Drive firmware updates can be performed online, concurrently with I/O. However, `fpga` updates require the drive to be taken offline, which means target drives must be made `candidate` before issuing the **applydrivesoftware** command. The default value is `firmware`. See the **chdrive** command for more details.

-all
(Optional) Specifies that the drive firmware should be applied to every drive in the system, as long as that drive is online and has `use member`, `use spare`, or `use candidate`.

This does not apply to:

- Drives that have dependent volumes
- Drives that are members of non-redundant arrays

Drives hosting quorum qualify, but there is risk. To avoid this risk use **-drive** and make sure the quorum is moved in between **applydrivesoftware** invocations. Use the **chquorum** command to avoid updating a drive that is hosting quorum.

If you specify **-all** you must specify the **-type** as `firmware`.

Remember: The **-all** parameter differs from the **-drive** parameter because unsuitable drives are not added to the list of drives scheduled for update when you use **-all**.

-drive *drive_id*
(Optional) Specifies one drive ID or a list of drive IDs (separated by a colon, [:]) to be updated. The maximum number of IDs is 128. If you have more than 128, use **-all** or multiple **applydrivesoftware** invocations to complete the update.

© Copyright IBM Corp. 2003, 2025

Remember: The **-drive** parameter differs from the **-all** parameter because if you specify an unsuitable drive using the **-drive** parameter, the **applydrivesoftware** command fails. Additionally, if you specify all three drives as hosting quorum, the command fails. If you use the **-drive** option to specify a single drive, and that drive has use=unused, it is updated.

-force
(Optional) Bypasses the dependent volume check. By default **applydrivesoftware** cannot run if any volumes dependent on the drive. Specifying **-force** bypasses this check, allowing the drive software update to proceed. Drive-dependent volumes generally result from non-redundant or degraded RAID arrays.

Note: Restore redundancy to the system (where possible) instead of using the **-force** parameter.

Important: Using the **-force** parameter might result in a data loss. Use it only under the direction of your product support information, or if you are willing to accept the risk of data loss in the array or pool to which the drive belongs.

-allowreinstall
(Optional) Specifies to make the system install the current level (again) onto drives that contain a file in the package.

Remember: Using this parameter is not recommended.

-allowdowngrade
(Optional) Specifies to allow the system to downgrade the firmware on a drive (that contains a file in the package).

Remember: Use this parameter only under the direction of your product support information.

-cancel
(Optional) Specifies that the command be stopped.

Description

Note: From 8.4.2 release onwards, before upgrading drive firmware, the latest software upgrade test utility must be installed and run using the following command:

```
svcupgradetest -w <firmware_package>
```

To know more about the software upgrade test utility, see https://www.ibm.com/support/pages/software-upgrade-test-utility.

Use this command to update the firmware of drives that are managed by the system.

There are two types of drive software that can be updated using this command:

- firmware
- fpga

Drive firmware updates can be performed online while the drive is in use. When used on an array member drive **applydrivesoftware** checks for volumes that are dependent on the drive and refuses to run if any are found. Drive dependent volumes are usually caused by non-redundant or degraded RAID arrays. Where possible you should restore redundancy to the system by replacing any failed drives before using the **applydrivesoftware** command. When this is not possible, for example on drives that are members of a RAID-0 array, you can either add redundancy to the volume by adding a second copy in another pool, or use the **-force** parameter to bypass the dependant volume check.

Remember: Only use **-force** if you are willing to accept the risk of data loss on dependent volumes (if the drive fails during the firmware update).

Drive firmware updates occur asynchronously, and conclude after the **applydrivesoftware** command completes. To see the status of the updates, use the **lsdriveupgradeprogress** command.

Drive fpga updates might require the drive to be taken offline for several minutes. Drives must be changed to the candidate state before **applydrivesoftware** can be used to update fpga software.

The fpga updates occur asynchronously, continuing in the background after the **applydrivesoftware** command has returned. You must check the FPGA_level field in **lsdrive** *N*, where *N* is the *drive_id*, to see whether or not the update completed successfully.

Remember: Interrupting an fpga update by removing power from the drive or enclosure might make the drive unusable. Only one drive's fpga can be updated per **applydrivesoftware** invocation. Make sure that the update is complete before unseating the drive or removing power from the enclosure.

An invocation example

```
applydrivesoftware -file DRIVE_XXXXXXXX -type firmware -drive 4
```

The resulting output:

```
No feedback
```

An invocation example

```
applydrivesoftware -cancel
```

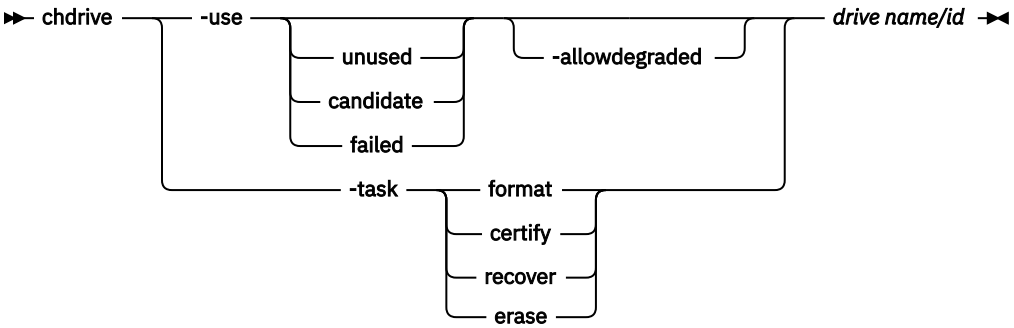
The resulting output:

```
No feedback
```

chdrive

Use the **chdrive** command to change the drive properties.

Syntax



Parameters

-use unused | candidate | failed
Describes the role of the drive:

- **unused** indicates that the drive is not in use and will not be used as a spare.
- **candidate** indicates that the drive is available for use in an array.
- **failed** indicates that the drive has failed.

Note: To create member drives, add the drives to (new) arrays by using the **mkarray** command.

If a drive fails for a distributed array, the array remains associated with the failed drive while it is in the **failed** state.

Note: If **chdrive -use candidate** is run on a compressed drive, it causes the drive to perform a format. The format must complete before the drive becomes online and available for use.

-allowdegraded
(Optional) Permits permission for a change of drive use to continue, even if a hotspare drive is not available for the array that the drive is a member of. You cannot specify **-allowdegraded** and **-task** together. You cannot specify **-allowdegraded** when an array expansion is in progress.

Important: Avoid using **-allowdegraded**.

-task format | certify | recover | eraseformat | certify | erase
Causes the drive to perform a task:

- **format** indicates that a drive will be formatted for use in an array; only permitted when drive is a candidate or has failed validation.

  Note: For NVMe drives, formatting the drive completes a cryptographic erase of the drive.
- **certify** indicates that the drive will be analyzed to verify the integrity of the data it contains; permitted for any drive that is a member.

  Note: certify is not supported in NVMe drives.
- **recover** recovers an offline flash drive without losing data; permitted when the drive is offline because a build is required, or when the drive has failed validation.

  Note: recover is not supported in NVMe drives.
- **erase** causes a sanitize disk erase operation. The type of erase is either cryptographic, block, or overwrite. The type of erase depends on the drive manufacturer, drive type, and firmware.

  Note: This parameter is not available on cloud systems.

drive_id
The identity of the drive.

Description

Use this command to change the drive role, or to start long running drive tasks.

You can use **lsdriveprogress** to display progress (percentage) and estimated completion time of ongoing drive tasks.

When a drive that is associated with a distributed array is changed from **member** to **failed**, if the distributed array does not have available rebuild space then it is degraded. If **-allowdegraded** is not specified, the command fails because of insufficient rebuild areas. If the **-allowdegraded** parameter is specified, the command succeeds and the array no longer uses the drive for I/O operations. If a drive is changed from **failed** to another configuration, the information for that drive is removed and the distributed array creates a missing member that belongs in the member table. Use the **charraymember** command to replace the missing member.

An invocation example to certify drive 23

```
chdrive -task certify 23
```

The resulting output:

```
No feedback
```

lsdrive

Use the **lsdrive** command to display configuration information and drive vital product data (VPD).

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-bytes
(Optional) The size (capacity) of the drive in bytes.

drive_id
(Optional) The identity of the drive.

Description

Use this command to display configuration information and drive VPD.

Note: Filtering should be permitted on all concise fields.

Note: A backup of system state is stored on quorum disks, which contributes to the physically used space. The amount of data that is used for this backup depends on the system configuration, but is typically a few MB and cannot exceed 1 GB.

This table describes possible outputs.

Table 75. ***lsdrive*** *output*

Attribute	Value
id	Indicates the ID of the drive.
status	Indicates the summary status of the drive: • `online`, which indicates that the drive is available through all drive ports. • `degraded`, which indicates that the drive is available but not through all drive ports. • `offline`, which indicates that the drive is unavailable.
error_sequence_number	Indicates the error sequence number that describes the cause of the drive status: • `online`, which is blank. • `degraded`, which is populated if associated with an error. • `offline`, which must be populated. **Note:** Error sequence numbers indicate an entry in the event log. This value includes entries that are both errors, and informational messages (for example, the drive is formatting).
use	Indicates the current role of the drive: • `unused` if the drive is not configured to be used by anything. • `candidate` if the drive is available to be configured. • `member` if the drive is configured as a member of an array. • `failed` if the drive is rejected and is no longer available for use.
UID	Indicates that the unique ID reported by the drive.

*Table 75. **lsdrive** output (continued)*	
Attribute	**Value**
tech_type	Indicates the drive technology used. The values are: • unsupported indicates that the drive is not supported in this platform (contact your support team). • tier_scm • tier0_flash • tier1_flash • tier_enterprise • tier_nearline
replacement_date	Indicates the date of a potential drive failure. The format must be YYMMDD. This value is blank for non-SSD drives.
capacity	Indicates the capacity of the disk, excluding quorum area. For compressed drives, this output represents the logical capacity of the drive, which is greater than the physical capacity, and is determined by the number of addressable logical blocks that are presented by the drive. For non-compressed drives, the capacity value is the same for both logical and physical capacity.
block_size	Indicates the block size of the disk.
vendor_id	Indicates the manufacturer of the drive.
product_id	Indicates the product ID of the drive.
FRU_part_number	Indicates the FRU part number of the drive.
FRU_identity	Indicates the 11S number that combines manufacturing part number and serial number.
RPM	Indicates the specified RPM of the disk.
firmware_level	Indicates the firmware level of the disk; blank if unknown.
FPGA_level	Indicates the FPGA level, if applicable; blank if not applicable or unknown.
mdisk_id	Indicates the ID of the array MDisk that the drive is a member of.
mdisk_name	Indicates the name of the MDisk that the drive is a member of.
member_id	Indicates the ID of the MDisk array member.
enclosure_id	Indicates whether the: • Drive is contained in an enclosure (not a node) and the slot position is known. This value is the ID of the enclosure in which the drive is located. • Drive is contained in a node (not an enclosure). This value is blank. • Enclosure ID is not determined yet. This value is blank.
slot_id	Indicates the **slot_id** of the drive in the enclosure or node. It can be referred to as the drive bay or location. This value can be blank.
node_name	Indicates the node name where the drive is located for a drive that is contained within a node. For a drive contained within an enclosure, it is blank.

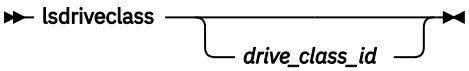
*Table 75. **lsdrive** output (continued)*	
Attribute	**Value**
node_id	Indicates the node ID where the drive is located for a drive that is contained within a node. For a drive contained within an enclosure, the value is blank.
quorum_id	Indicates the ID of quorum disk. The value is blank if not quorum disk.
port_1_status	Indicates the connectivity status of the drive's first port. The values are `online`, `offline`, or `excluded`. **Note:** Port 1 is attached to the node that has a panel name that ends in `-1`.
port_2_status	Indicates the connectivity status of the drive's second port. The values are `online`, `offline`, or `excluded`. **Note:** Port 2 is attached to the node that has a panel name that ends in `-2`.
health_state	Indicates the drive health status: • `good` indicates good health. • `warning` indicates a warning that must be examined. • `critical` indicates something critical that must be addressed. • `unknown` indicates that the status is unknown.
encrypted	Indicates the state of the drive encryption.
rekey	Indicates the state of the rekey of each drive: • `committing` indicates that the rekey is being committed. • `prepared` indicates that the rekey is being prepared. • `error` indicates that there is an error with the rekey. • `blank` indicates that there is no status.
auto_manage	Indicates whether the `auto_manage` process is running (`active`) or not running (`idle`).
drive_class_id	Indicates which drive class the drive is part of.
write_endurance_used	Indicates the drive writes per day (DWPD). This value is blank for drives that are not SSD drives. The value must be a number `0` - `255`. This value indicates the percentage of life that is used by the drive. The value `0` indicates that full life remains, and `100` indicates that the drive is at or past its end of life. **Note:** The drive must be replaced when the value exceeds `100`. This value is blank for drives that are either: 1. Not SSDs. 2. SSDs that predate support of the endurance indicator. This value also applies to drives that are yet to be polled, which can take up to 24 hours.

Table 75. **lsdrive** output (continued)	
Attribute	**Value**
write_endurance_usage_rate	Indicates the DWPD usage rate. The values are: • measuring • high • marginal • low This value is blank for non-SSD drives. **Note:** This field displays a value only when the write_endurance_used value changes.
date_of_manufacture	Indicates the manufacture date for SAS drives (YYMMDD).
fips_enabled	Indicates if the drive follows FIPS standards. Value can be yes or no.

lsdriveclass

Use the **lsdriveclass** command to display all drive classes in the system. Flash cards are not displayed in this view because they do not support distributed RAID levels.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

drive_class_id
(Optional) The identity of the drive class.

Description

This command displays all drive classes in a system. Drives are displayed if they are managed.

Note: Only SAS drives have classes that appear in this view.

This table provides the attribute values that can be displayed as output view data.

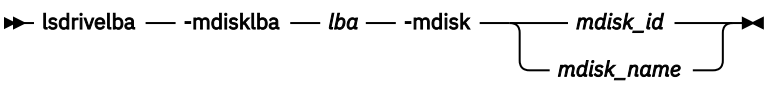
Table 76. **lsdriveclass** output	
Attribute	**Possible Values**
id	Indicates the drive class ID.
RPM	Indicates the speed of the drive class.

Table 76. **lsdriveclass** output (continued)	
Attribute	**Possible Values**
capacity	Indicates the capacity of the drive class.
IO_group_id	Indicates the I/O group ID associated with the drive class.
IO_group_name	Indicates the I/O group name that is associated with the drive class.
tech_type	Indicates the technology type of the drive class. • tier0_flash • tier1_flash • tier_enterprise • tier_nearline • tier_scm
block_size	Indicates the block size of the drive class.
candidate_count	Indicates the number of drives in the drive class that are in `candidate` state.
superior_count	Indicates the total number of drives in this class and the drives that count as superior. (This value applies to distributed arrays created by using **mkdistributedarray**.)
total_count	Indicates the total number of drives in the drive class. The drive state is irrelevant.
transport_protocol	Indicates the transport protocol of the drive. The possible values are: • sas • nvme
compressed	Indicates whether the drive is compressed.
physical_capacity	For compressed drives, this output represents the total amount of physical storage on the drive. This value can be smaller than the logical capacity presented by the **capacity** value. For non-compressed drives, the **physical_capacity** is the same as the logical capacity.

lsdrivelba

Use the **lsdrivelba** command to map array MDisk logical block address (LBA) to a set of drives.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-mdisklba *lba*
(Required) The logical block address (LBA) on the MDisk. The LBA must be specified in hex, with a `0x` prefix.

-mdisk*mdisk_id* **|** *mdisk_name*
(Required) The ID or name of the MDisk.

Description

This command maps the array MDisk LBA to a set of drives.

Note: This command applies to SAS drives.

The system provides volumes that have LBAs for 512-byte block sizes, but back-end disks that have a block size of either 512 or 4096 bytes can also be used. Drives are listed in their physical size.

Use the **lsdrive** command to display the drive block size, and use the **lsdrive** or `lsarray` command to list each object (the drive and the MDisk).

This table describes possible outputs.

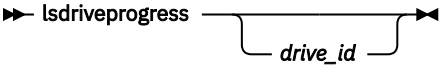
Table 77. lsdrivelba output

serial-attached SCSI | **| | (SAS)** **Attribute**	**Value**
drive_id	The ID of drive; blank if no configured array member exists (for example, in a degraded array).
type	The type of information on the disk: • `parity`, in which LBA range contains `parity` (RAID levels 5 and 6 only) • `qparity`, in which LBA range contains `qparity` (RAID level 6 only) • `data`, in which LBA range contains data
drive_lba	The LBA on the drive.
drive_start	The start of range of LBAs (strip) on the drive.
drive_end	The end of range of LBAs (strip) on the drive.
mdisk_start	The start of range of LBAs (strip) on the array MDisk.
mdisk_end	The end of range of LBAs (strip) on the array MDisk.

lsdriveprogress

Use the **lsdriveprogress** command to view the progress of various drive tasks.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

Description

The following outputs are possible:

drive_id
Indicates the ID for the drive with the active task.

task
Indicates the type of task:

- format
- certify
- erase
- recover

progress
Indicates the percentage completion for a job.

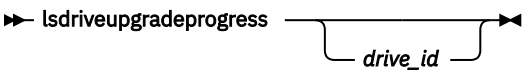
estimated_completion_time
Indicates the estimated completion time, in the format YYMMDDHHMMSS, where:

- Y is the year.
- (The first) M is the month.
- D is the day.
- H is the hour.
- (The second) M is the minute.
- S is the second.

lsdriveupgradeprogress

Use the **lsdriveupgradeprogress** command to view the status or progress of drives with pending downloads. This command applies to serial-attached SCSI (SAS) drives.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

drive_id
(Optional) Specifies the update status or progress for a single drive. If not specified, the update status for all scheduled drives is displayed.

Note: If you specify this parameter, **lsdriveupgradeprogress** displays the update status of this drive. If you do not specify this parameter, **lsdriveupgradeprogress** displays the update status of all requested drives.

Description

The **lsdriveupgradeprogress** command completes whether the original **applydrivesoftware** command was addressed to a single drive (synchronous command) or multiple drives (asynchronous command).

The following outputs are possible:

id
Indicates the identity of the active drive.

status
Indicates the drive status. Each of the following values has a specific meaning:

- `progressing` indicates all scheduled drives have completed; wait 270 seconds before you issue **applydrivesoftware** again.
- `completed` indicates a successful firmware download.
- `updating` indicates the update is ongoing.
- `scheduled` indicates the update is in the download list, waiting to download.
- `not_scheduled` indicates the drive is not scheduled.

  Note: This means the corresponding drive was not scheduled when **applydrivesoftware -all** was last issued.
- `canceled` indicates the update has been canceled, either by command or by a system change that stops the **applydrivesoftware** command from running properly.
- `invalid` indicates the drive download status is invalid, also the initial state.
- If the field is blank, that indicates that the download is not scheduled (for example, the **applydrivesoftware** command has not been issued).

estimated_completion_time
Indicates the estimated completion time (YYMMDDHHMMSS), where:

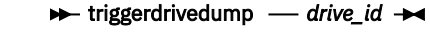
- Y is year
- (The first) M is month
- D is day
- H is hour
- (The second) M is minute
- S is second

The value is blank if the status is either canceled or blank.

triggerdrivedump

Use the **triggerdrivedump** command to collect support data from a disk drive. This data can help you understand problems with the drive and does not contain any data that applications write to the drive.

Syntax



Parameters

drive_id
(Required) The ID of the drive to dump.

Description

Use this command to collect internal log data from a drive and store the information in a file in the /dumps/drive directory. This directory is on one of the nodes connected to the drive. The system limits the number of drive dump files in the directory to 24 per node.

An invocation example

```
triggerdrivedump 2
```

The resulting output:

```
Dump file for drive [2] created
```

Note: The system chooses the node on which to run the statesave.

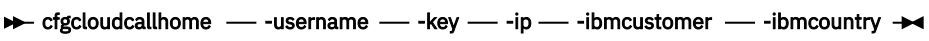
Chapter 13. Call home and event notification commands

Use the call home and event notification commands to enable your system to send notifications.

cfgcloudcallhome

Use the **cfgcloudcallhome** command to configure email and metering features by using an Internet Protocol (IP) quorum server as a Simple Mail Transfer Protocol (SMTP) server on a system.

Syntax



Parameters

-username
(Required) Specifies the IBM Cloud application programming interface (API) user name.

-key
(Required) Specifies the IBM Cloud API key.

-ip
(Required) Specifies the IP address of the IP quorum server.

-ibmcustomer
(Required) Specifies the customer number that is assigned when a software license is automatically added to the entitlement database. The value must be a 7 - 10-digit number.

-ibmcountry
(Required) Specifies the country ID used for entitlement and the call home system. The value is either a 3-digit number or blank.

Description

This command configures email and billing features by using an Internet Protocol (IP) quorum server as a Simple Mail Transfer Protocol (SMTP) server on a system.

An invocation example

```
# cfgcloudcallhome -username callhome1@de.ibm.com -key xxxxx -ip 192.168.0.1 -ibmcustomer
12345678 -ibmcountry 886
```

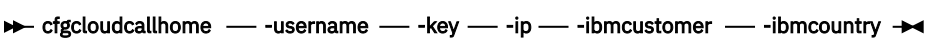
The following output is displayed:

```
None
```

cfgcloudcallhome

Use the **cfgcloudcallhome** command to configure email and metering features by using an Internet Protocol (IP) quorum server as a Simple Mail Transfer Protocol (SMTP) server on a system.

Syntax



© Copyright IBM Corp. 2003, 2025

Parameters

-username
(Required) Specifies the IBM Cloud application programming interface (API) user name.

-key
(Required) Specifies the IBM Cloud API key.

-ip
(Required) Specifies the IP address of the IP quorum server.

-ibmcustomer
(Required) Specifies the customer number that is assigned when a software license is automatically added to the entitlement database. The value must be a 7 - 10-digit number.

-ibmcountry
(Required) Specifies the country ID used for entitlement and the call home system. The value is either a 3-digit number or blank.

Description

This command configures email and billing features by using an Internet Protocol (IP) quorum server as a Simple Mail Transfer Protocol (SMTP) server on a system.

An invocation example

```
# cfgcloudcallhome -username callhome1@de.ibm.com -key xxxxx -ip 192.168.0.1 -ibmcustomer
12345678 -ibmcountry 886
```

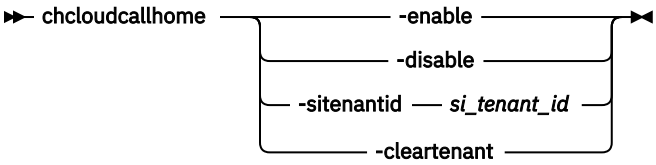
The following output is displayed:

```
None
```

chcloudcallhome

Use the **chcloudcallhome** command to enable or disable the Cloud Call Home function.

Syntax



Parameters

-enable
(Optional) Call Home data is sent directly to the Cloud. This parameter is the default setting.

-disable
(Optional) Call Home data is not sent to the Cloud.

-sitenantid
(Optional) Adds the IBM Storage Insights service name.

-cleartenant
(Optional) Removes the IBM Storage Insights service name.

Description

This command allows configuration changes to Cloud Call Home.

Note:

- The **-enable** and **-disable** parameters are mutually exclusive.
- Call Home **sitenantid** is a combination of special characters (- , _ , and .) and alphanumeric characters. When you run the **chcloudcallhome** command with **sitenantid** parameter, it sets the statistics frequency to five minutes by default, unless it is already set to one or five minutes.

An invocation example

```
chcloudcallhome -enable
```

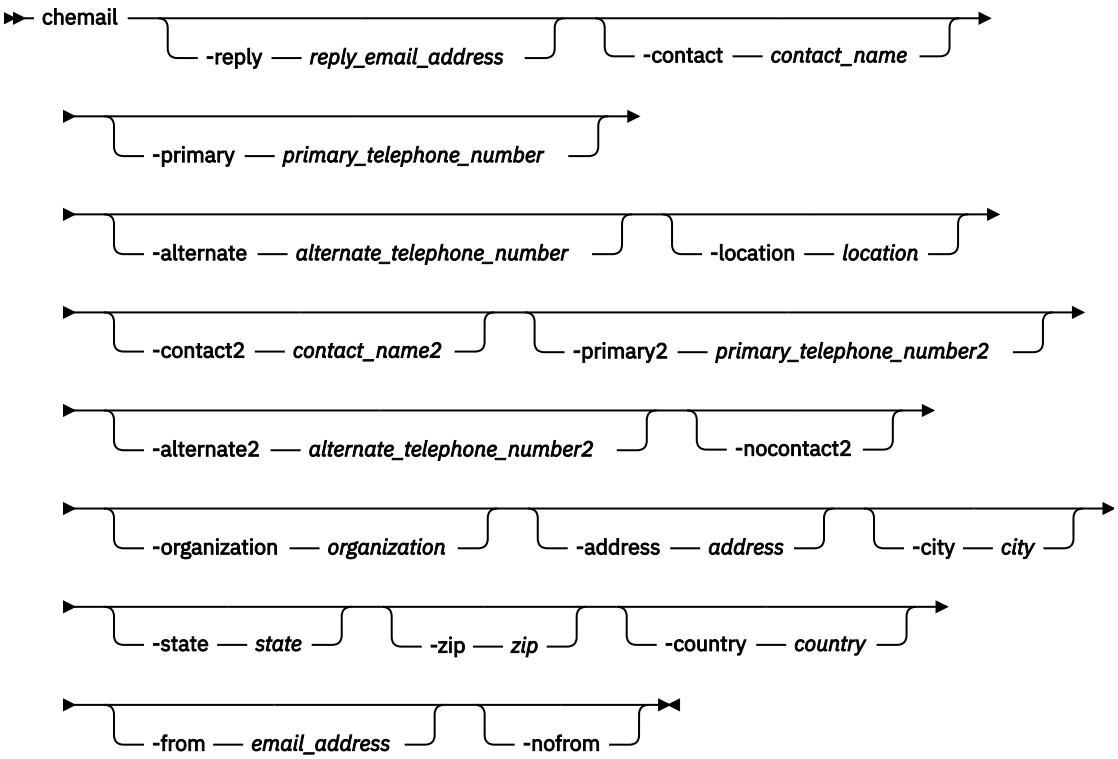
The following output is displayed:

```
No feedback.
```

chemail

Use the **chemail** command to set or modify contact information for email event notifications. At least one of the parameters must be specified to modify settings.

Syntax



Parameters

-reply ***reply_email_address***
(Optional) Specifies the email address to which a reply is sent.

-contact ***contact_name***
(Optional) Specifies the name of the person to receive the email.

-primary ***primary_telephone_number***
(Optional) Specifies the primary contact telephone number.

Note: For machine types 2071 and 2072 (in the United States and Canada), the value entered must be exactly ten decimal digits. For machines types 2071 and 2072 (in other countries) the value entered can be five to nineteen decimal digits. Otherwise, there can be up to nineteen characters.

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

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

-contact2 *contact_name2*
(Optional) Specifies the name of the second contact person to receive the email.

For machine types 2071 and 2072 the maximum number of characters is 30. For other machine types the maximum number of characters is 72.

-primary2 *primary_telephone_number2*
(Optional) Specifies the primary contact telephone number for the second contact person.

Note: For machine types 2071 and 2072 (in the United States and Canada), the value entered must be exactly ten decimal digits. For machines types 2071 and 2072 (in other countries) the value entered can be five to nineteen decimal digits. Otherwise, there can be up to nineteen characters.

-alternate2 *alternate_telephone_number2*
(Optional) Specifies the alternate contact telephone number for the second contact person.

-nocontact2
(Optional) Removes all the contact details for the second contact person.

-organization *organization*
(Optional) Specifies the user's organization as it should appear in Call Home emails.

-address *address*
(Optional) Specifies the first line of the user's address as it should appear in Call Home email.

-city *city*
(Optional) Specifies the user's city as it should appear in Call Home email.

-state *state*
(Optional) Specifies the user's state as it should appear in Call Home email. This is a two-character value such as NY for New York.

-zip *zip*
(Optional) Specifies the user's zip code or postal code as it should appear in Call Home email.

-country *country*
(Optional) Specifies the country in which the machine resides as it should appear in Call Home email. This is a two-character value such as US for United States.

For machine types 2071 and 2072 this value cannot be US or CA if the value for **primary** or **primary2** telephone number is not blank or exactly 10 digits.

-from *email_address*
(Optional) Specifies the email address that is set in the **from:** field of the email that is sent by the system.

-nofrom
(Optional) Removes the email address that is set in the **from:** field of the email that is sent by the system.

Description

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

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

Remember: When considering e-mail addresses:

- Alphanumeric characters and additionally underscore (_), at symbol (@), and dot (.) characters are permitted.
- There must be exactly one @ character in the string, and the @ characters must not start or end the string.
- A plus (+) character is permitted before the @ character.

The email notification system cannot be started unless the following fields are set (in addition to the existing four: reply, contact, primary, location):

- Organization
- Address
- City
- State
- Zip
- Country

The state field has a default value of XX, which is for systems that are not in the USA, Canada, Puerto Rico, the Virgin Islands, or Guam.

Remember: The user must enter values for these fields in order to enable the email notification system.

These fields do not have to be set to start the email notification system, but if the new fields are set they are included in the email event notifications.

An invocation example

```
chemail -reply jane.doe@ibm.com
 -contact 'Jane Doe'
 -primary 01234 567890
 -location 'Server Room'
 -organization IBM
 -address '1 Orchard Rd'
 -city 'New York'
 -zip 10504
 -country US
```

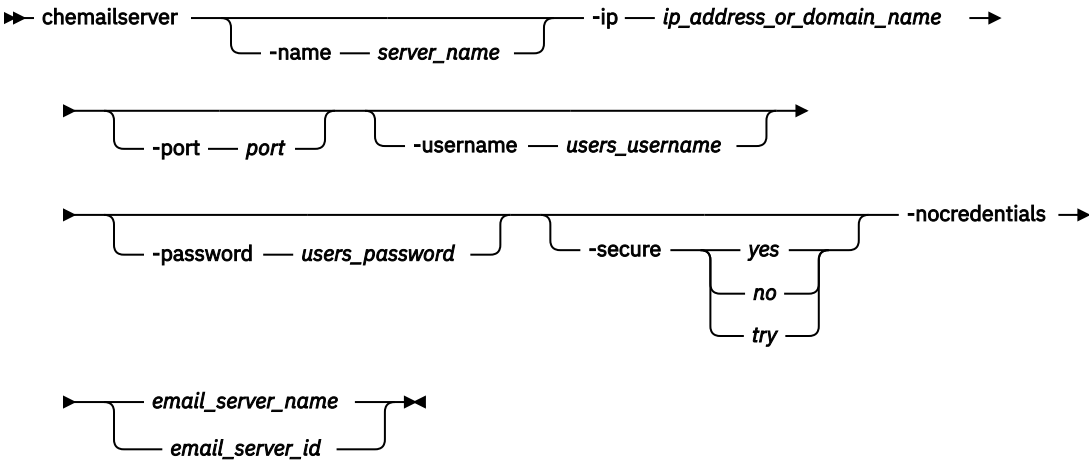
The resulting output:

```
No feedback.
```

chemailserver

Use the **chemailserver** command to modify the parameters of an existing email server object.

Syntax



Parameters

-name *server_name*
(Optional) Specifies a unique name to assign to the email server object. The name must be a 1- through 63-character string, and cannot start with a hyphen or number. When specifying a server name, emailserver is a reserved word.

-ip *ip_address_or_domain_name*
(Optional) Specifies the IP address or domain name of a remote. The IP address must be a valid IPv4 or IPv6 address or a fully qualified domain name. IPv6 addresses can be zero compressed. If you specify a fully qualified domain name, a DNS server must be configured on your system. You can use the **mkdnsserver** command to configure DNS servers.

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

-username *users_username*
(Optional) Specify the username for authenticating to the email server. The maximum number of characters for the username is 64.

-password *users_password*
(Optional) Specify the password for authenticating to the email server. The maximum number of characters for the password is 64.

Note: If the username and password are not already set, you must specify both values in the command.

-secure *yes|no|try*
(Optional) Specify whether to establish a TLS connection to the email server. The following options are available:

- `try` - if the email server does not support TLS connection, then continue and send the email insecurely.
- `yes` - if the email server does not support TLS, stop the connection. Emails fail to send.
- no - do not attempt to establish a TLS connection to the email server. Emails are sent insecurely.

-nocredentials
(Optional) Removes the credentials that are used for authenticating to the email server .

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

Description

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

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

An invocation example

```
chemailserver -name newserver 0
```

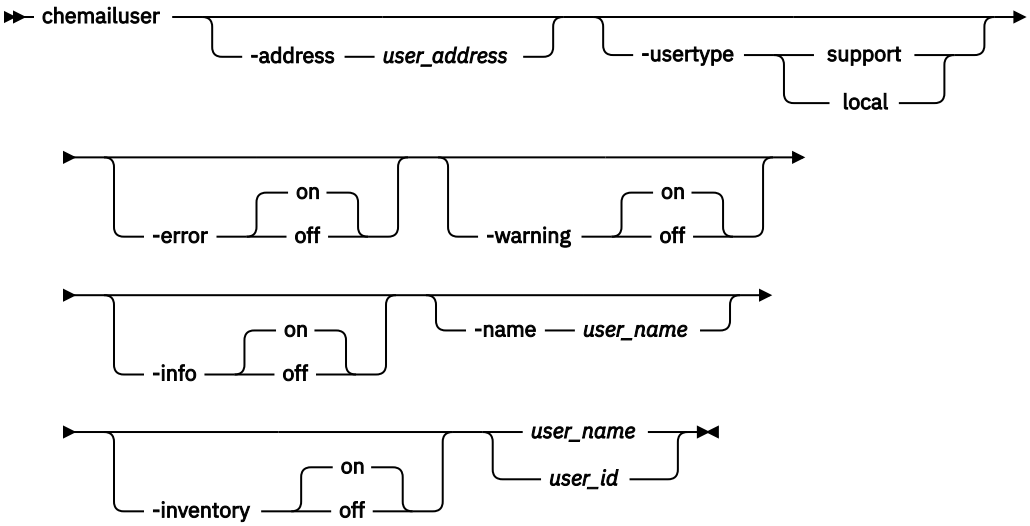
The resulting output:

```
No feedback.
```

chemailuser

Use the **chemailuser** command to modify the settings that are defined for an email recipient.

Syntax



Parameters

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

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

support
Address of the support organization that provides vendor support.

local
All other addresses.

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

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

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

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

-inventory on | off
(Optional) Specifies whether this recipient receives inventory email notifications.

user_name | user_id
(Required) Specifies the email recipient for whom you are modifying settings.

Description

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

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

Remember: When considering e-mail addresses:

- Alphanumeric characters and additionally underscore (_), at symbol (@), and dot (.) characters are permitted.
- There must be exactly one @ character in the string, and the @ characters must not start or end the string.
- A plus (+) character is permitted before the @ character.

For syslog servers, Simple Network Management Protocol (SNMP) servers, and email users with type set to local, the valid combinations of notification types are:

- **error**, **warning**, and **info**
- **error** and **warning**

An invocation example

The following example modifies email settings for email recipient manager2008:

```
chemailuser -usertype local manager2008
```

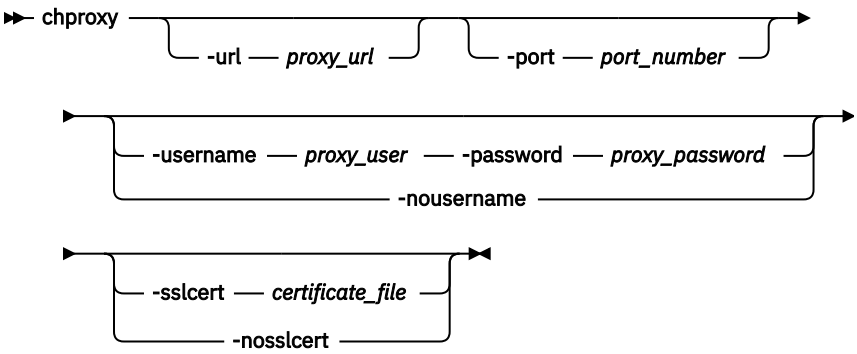
The resulting output:

```
No feedback
```

chproxy

Use the **chproxy** command to update the system-wide web proxy server configuration.

Syntax



Parameters

-url *proxy_url*
(Optional) Specifies the proxy server URL with a protocol prefix in fully qualified domain name format.

-port *port_number*
(Optional) Specifies the proxy server port number. The value must be an integer (1-65535).

-username *proxy_user*
(Optional) Specifies the basic authentication username. The username cannot start or end with a blank. The username can consist of a string of 1 - 64 ASCII characters except for the following characters `% : " , * '`.

-password *proxy_password*
(Optional) Specifies the 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. Do not specify the **-password** parameter if you want to have an empty password.

-nousername
(Optional) Specifies the removal of the user name and password values.

-sslcert *certificate_file*
(Optional) Specifies the certificate file path and file name for the proxy server.

-nosslcert
(Optional) Specifies the removal of the certificate file.

Description

This command updates a system proxy server configuration.

Invocation examples

```
chproxy -username puffin

chproxy -nosslcert

chproxy -nousername

chproxy -url http://654321 -port 8080
CMMVC6523E The URL that you have entered is not valid.

chproxy -port XYZ
CMMVC5706E An invalid argument has been entered for the [port] parameter.

chproxy -sslcert /upgrade/no-such-file.txt
CMMVC8031E file not found.
```

```
chproxy -sslcert /upgrade/not-a-cert.txt
CMMVC8794E Invalid certificate file.
```

The resulting output:

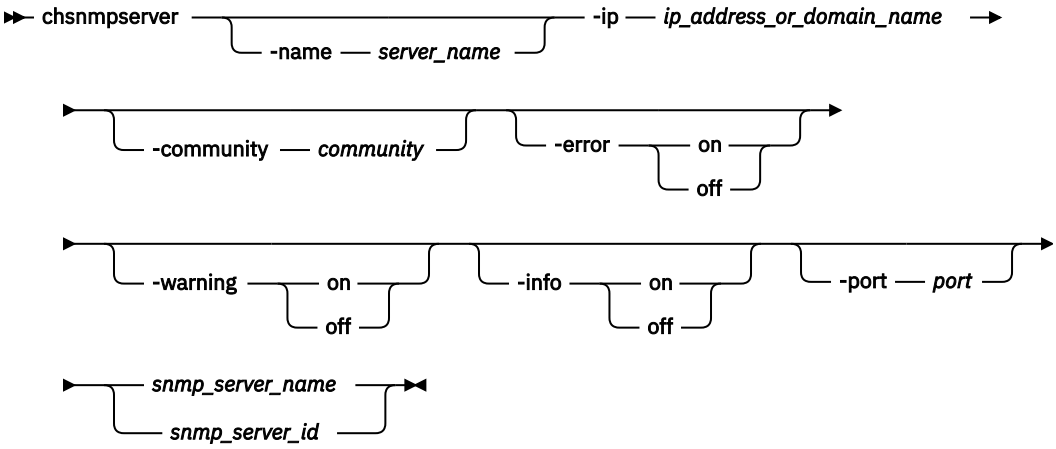
```
No feedback
```

chsnmpserver

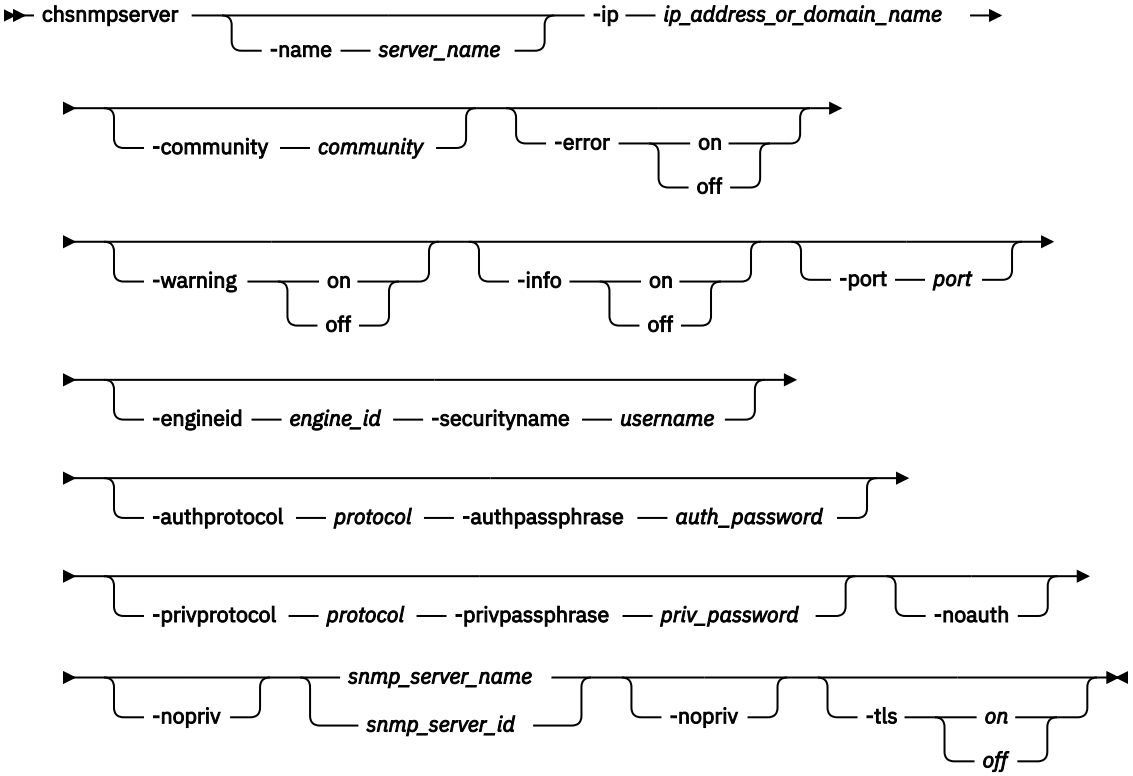
Use the **chsnmpserver** command to modify the parameters of an existing SNMP server.

Syntax

This diagram is for SNMPv2.



This diagram is for SNMPv3.



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_or_domain_name***
(Optional) Specifies an IP address or domain name to assign to the SNMP server. The IP address must be a valid IPv4 or IPv6 address, or a fully qualified domain name. If you specify a domain name, a DNS server must be configured on your system. You can use the **mkdnsserver** command to configure DNS servers.

-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.

-engineid ***engine_id***
(Optional) Specifies the engine ID that identifies the local system to the SNMPv3 server. This value has a maximum of 32 bytes, and must meet the requirements outlined in RFC 3411. You must specify **-securityname** with this parameter.

-securityname ***username***
(Optional) Specifies the security name for the SNMPv3 server. This value has a max of 32 characters. You must specify **-engineid** with this parameter.

-authprotocol ***protocol***
(Optional) Specifies the authentication protocol for the SNMPv3 server. Available values are **sha1** or **md5, sha224**, **sha256**, **sha384** or **sha512**. You must specify **-authpassphrase** with this parameter.

-authpassphrase ***auth_password***
(Optional) Specifies the authentication password for the SNMPv3 server. This value must be in the range of 8 - 64 characters. You must specify **-authprotocol** with this parameter.

-privprotocol ***protocol***
(Optional) Specifies the privacy protocol for the SNMPv3 server. Available values are **aes128** or **des, aes192**, **aes256**. You must specify **-privpassphrase** with this parameter.

-privpassphrase ***priv_password***
(Optional) Specifies the privacy password for the SNMPv3 server. This value must be in the range of 8 - 64 characters. You must specify **-privprotocol** with this parameter.

-tls on | off
(Optional) Specifies whether SNMP traps sent over a TLS connection. Values can be on or off. Default is set to off.

-noauth
(Optional) Removes all authentication and privacy settings from the SNMPv3 server. It is not possible to have privacy settings without authentication.

-nopriv
(Optional) Removes the privacy protocol and passphrase settings from the SNMPv3 server.

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

Description

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

To switch from an SNMPv3 server to an SNMPv2 server, you must specify the **-community** parameter.

For syslog servers, SNMP servers, and email users with type set to `local`, the valid combinations of notification types are:

- **error**, **warning**, and **info**
- **error** and **warning**

Note: The port defaults to 10162 instead of 162 if tls is set to on and no other port is specified.

An invocation example to change the name of snmpserver 0

```
chsnmpserver -name newserver 0
```

The resulting output:

```
No feedback
```

lscloudcallhome

Use the **lscloudcallhome** command to view the status of the Call Home information that is sent to a server in the Cloud.

Syntax



Parameters

No parameters are available for this command.

Description

This command displays the status of Call Home information that is sent directly to a server in the Cloud. This table shows the possible output:

*Table 78. **lscloudcallhome** output*

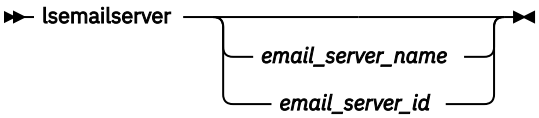
Attribute	Description
status	Displays the status of the Cloud Call Home function. The possible values are `unsupported`, `disabled`, or `enabled`.
connection	Displays the connection state to the server. The possible values are: • `active` - A good connection exists to the servers in the Cloud. • `error` - A connection error exists. The error value that is displayed in the `error_sequence_number` field details the event log number that contains further information about the connection error. • `untried` - A value of `untried` occurs directly after enabling cloud call home when the system is waiting for the results of a connection test.

Table 78. **lscloudcallhome** output (continued)	
Attribute	**Description**
error_sequence_number	Indicates the error sequence number that describes the cause of the connection error.
last_success	Displays the date and time of the last successful message that was sent. The format is YYMMDDhhmmss.
last_failure	Displays the date and time of the last message failure. When there is a working active connection failure that is older than 24 hours, this field is blank. The format is YYMMDDhhmmss.
si_tenant_id	Displays the IBM Storage Insights service ID.

lsemailserver

Use the **lsemailserver** command to display a concise list or a detailed view of email servers that are configured on the system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

email_server_name | email_server_id
(Optional) Specifies the name or ID of an existing email server that must be listed.

Description

Use this command to display a concise list or a detailed view of email servers that are configured on the system.

This table provides the attribute values that can be displayed as output view data.

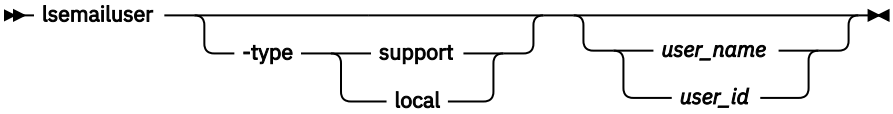
Table 79. **lsemailserver** output	
Attribute	**Description**
id	Indicates the ID of the email server.
name	Indicates the name of the email server.
IP_address	Indicates the IP address or domain name of the email server.
port	Indicates the port of the email server.

Table 79. **lsemailserver** output (continued)	
Attribute	**Description**
status	Indicates the state of the email server as reported by the last periodic health check or email send attempt. The possible values are: • untried (default) • failed • failed_temporary • active
username	Indicates the username for authenticating to the email server.
password_set	Indicates whether a password is used for authenticating to the email server.
secure	Indicates whether to establish a TLS connection to the email server. The following options are available: • `try` - if the email server does not support TLS connection, then continue and send the email insecurely. • `yes` - if the email server does not support TLS, stop the connection. Emails fail to send. • `no` - do not attempt to establish a TLS connection to the email server. Emails are sent insecurely.

lsemailuser

Use the **lsemailuser** command to generate a report that lists the email event notification settings for all email recipients, an individual email recipient, or a specified type (local or support) of email recipient.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-type support | local
(Optional) Specifies the types of email recipients you want to view, either customer-based or support-based as determined by the following definitions:

support
Address of the support organization that provides vendor support.

local
All other addresses.

user_name | user_id
(Optional) Specifies the user ID or user name of the email event recipient for whom you want to see the email notification settings.

Description

When you issue this command, a report is displayed that lists the email event notification settings for all email recipients, an individual email recipient, or a specified type (local or support) of email recipient. The concise and detailed views report the same information.

lsproxy

Use the **lsproxy** command to display the system-wide web proxy server configuration.

Syntax



Description

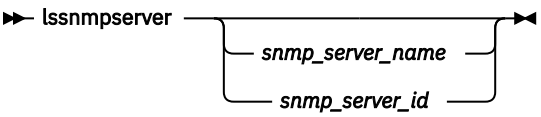
This table provides possible values for the attributes that are displayed as data in the output views.

*Table 80. **lsproxy** command output values*	
Attribute	**Value**
enabled	Displays whether a proxy server is configured. The value is either **yes** or **no**.
url	Displays the proxy server URL with the protocol prefix in fully qualified domain name format.
port	Displays the proxy server port number.
username	Displays the basic authentication user name.
password_set	Displays whether a basic authentication proxy password is defined. The value is either **yes** or **no**.
certificate	Displays a human-readable description of the whole server SSL certificate (as generated by openssl x509 -text). This field is blank if a certificate is not set.

lssnmpserver

Use the **lssnmpserver** command to return a concise list or a detailed view of SNMP servers that are configured on the system.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

snmp_server_name | snmp_server_id
(Optional) Specifies the name or ID of an existing SNMP server that must be listed.

Description

Use this command to display a concise list or a detailed view of SNMP servers that are configured on the system. This table describes possible outputs.

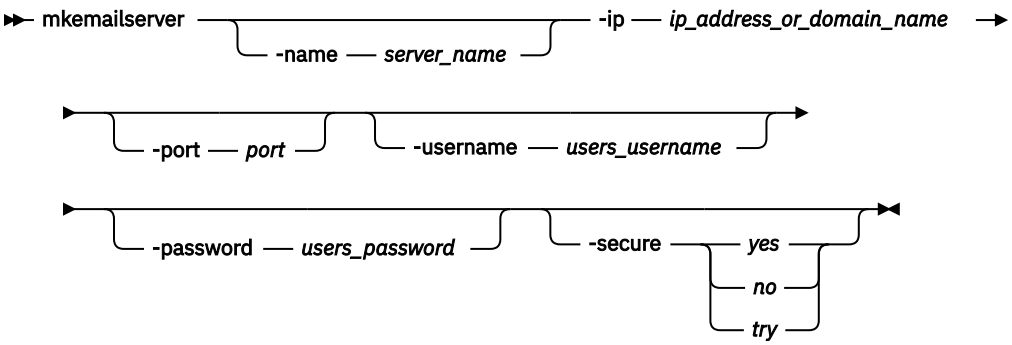
*Table 81. **lssnmpserver** output*

Attribute	Value
id	Indicates the ID of the SNMP server.
name	Indicates the name of the SNMP server.
IP_address	Indicates the IP address or domain name of the SNMP server.
error	Indicates whether the server receives error notifications. The values are on or off.
warning	Indicates whether the server receives warning notifications. The values are on or off.
info	Indicates whether the server receives information notifications. The values are on or off.
port	Indicates the remote port number for the SNMP server. This value is in the range of 1 - 65535.
community	Indicates the community name for the SNMPv2 server. This value is blank for SNMPv3 servers.
version	Indicates the SNMP server version. The values are 2 or 3.
securitylevel	Indicates the SNMPv3 security level. The values are blank, no_auth_no_priv, auth_no_priv, or auth_priv.
engineid	Indicates the SNMPv3 server engine ID. The values are blank or a 32-byte character string with the leading 0x removed.
securityname	Indicates the SNMPv3 server security name. The values are blank or a max 32-character string.
authprotocol	Indicates the SNMPv3 server authentication protocol. The values are blank, sha, or md5.
privprotocol	Indicates the SNMPv3 server privacy protocol. The values are blank, aes, or des.
authpassphrase_set	Indicates whether the SNMPv3 server authentication passphrase has been set. The values are yes or no.
privpassphrase_set	Indicates whether the SNMPv3 server privacy passphrase has been set. The values are yes or no.
tls	Indicates whether the SNMPv3 server is configured to use TLS over TCP. The values are on or off.

mkemailserver

Use the **mkemailserver** command to create an email server object that describes a remote Simple Mail Transfer Protocol (SMTP) email server.

Syntax



Parameters

-name ***server_name***
(Optional) Specifies a unique name to assign to the email server object. The name must be a 1- through 63-character string, and cannot start with a hyphen or number. If a name is not specified, then a system default of `emailserver`*n* is applied, where *n* is the object ID. When specifying a server name, emailserver is a reserved word.

-ip ***ip_address_or_domain_name***
(Required) Specifies the IP address or domain name of a remote. This must be a valid IPv4 or IPv6 address or a fully qualified domain name. IPv6 addresses can be zero compressed. If you specify a fully qualified domain name, a DNS server must be configured on your system. You can use the **mkdnsserver** command to configure DNS servers.

-port ***port***
(Optional) Specifies the port number for the email server. This must be a value of 1 - 65535. The default value is 25.

-username ***users_username***
(Optional) Specify the username for authenticating to the email server. The maximum number of characters for the username is 64.

-password ***users_password***
(Optional) Specify the password for authenticating to the email server. The maximum number of characters for the password is 64.

-secure ***yes|no|try***
(Optional) Specify whether to establish a TLS connection to the email server. The following options are available:

- `try` - if the email server does not support TLS connection, then continue and send the email insecurely.
- `yes` - if the email server does not support TLS, stop the connection. Emails fail to send.
- no - do not attempt to establish a TLS connection to the email server. Emails are sent insecurely.

Description

This command creates an email server object that represents the SMTP server. The system uses the email server to send event notification and inventory emails to email users. It can transmit any combination of error, warning, and informational notification types.

The system supports multiple email servers to provide redundant access to the external email network. The email servers are used in turn until the email is successfully sent. The attempt is successful when the system gets a positive acknowledgment from an email server that the email has been received by the server.

If you are using a cloud-based email server, authentication can be achieved by setting up the username and password.

You can specify **-secure** flag and the system proceeds based on whether the email server supports TLS connection.

An invocation example

```
mkemailserver -ip servername.example.com -port 78
```

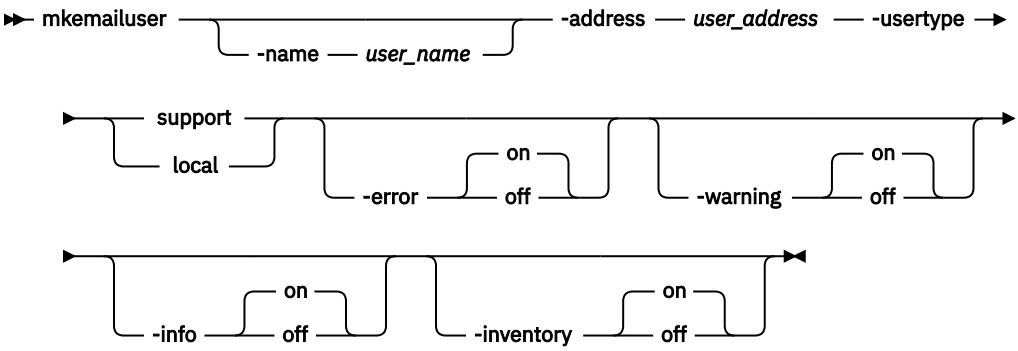
The resulting output:

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

mkemailuser

Use the **mkemailuser** command to add a recipient of email event and inventory notifications to the email event notification facility. Add up to twelve recipients (one recipient at a time).

Syntax



Parameters

-name *user_name*
(Optional) Specifies the name of the person who is the recipient of email event notifications. The *user_name* value must be unique, must not contain spaces, and must not contain only numbers. If you do not specify a user name, the system automatically assigns a user name in the format of `emailusern`, where *n* is a number beginning with 0 (`emailuser0`, `emailuser1`, and so on).

The name emailuser*n*, where *n* is a number, is reserved and cannot be used as one of your user names.

-address *user_address*
(Required) Specifies the email address of the person receiving the email event or inventory notifications, or both. The *user_address* value must be unique.

-usertype support| local
(Required) Specifies the type of user, either local or support, based on the following guidelines:

support
The recipient is your product support organization, using a default value (an auto-populated email address). This setting is used with the Call Home feature. For any other use case, contact your product support organization for direction.

local
All other recipients other than your product support organization. Select the `local` usertype unless otherwise instructed by your product support organization.

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

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

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

-inventory on | off
(Optional) Specifies whether this recipient receives inventory email notifications. The default value is `off`.

Description

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

Note: Before you can set the **`usertype`** parameter to `support`, turn the **`-warning`** and **`-info`** flags `off`.

Remember: When considering e-mail addresses:

- Alphanumeric characters and additionally underscore (_), at symbol (@), and dot (.) characters are permitted.
- There must be exactly one @ character in the string, and the @ characters must not start or end the string.
- A plus (+) character is permitted before the @ character.

For syslog servers, SNMP servers, and email users with type set to `local`, the valid combinations of notification types are:

- `error`, `warning`, and `info`
- `error` and `warning`

An invocation example

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

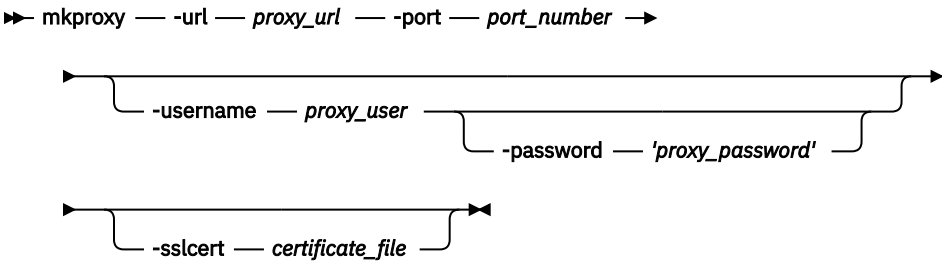
The resulting output:

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

mkproxy

Use the **mkproxy** command to create a system-wide web proxy server configuration.

Syntax



Parameters

-url *proxy_url*
(Required) Specifies the proxy server URL with a protocol prefix in fully qualified domain name format.

-port *port_number*
(Required) Specifies the proxy server port number. The value must be an integer (1-65535).

-username *proxy_user*
(Optional) Specifies the basic authentication username. The username cannot start or end with a blank. The username can consist of a string of 1 - 64 ASCII characters except for the following characters `% : " , * '`.

-password *'proxy_password'*
(Optional) Specifies the 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. Do not specify the **-password** parameter if you want to have an empty password. You should enclose the password in single quotation marks.

-sslcert *certificate_file*
(Optional) Specifies the certificate file path and file name for the proxy server.

Description

This command creates a system proxy server.

Invocation examples

```
mkproxy -url http://proxy.example.ibm.com -port 8080

mkproxy -url http://proxy.example.ibm.com -port 8080 -username albatross -password
yellowsubmarine

mkproxy -url http://proxy.example.ibm.com -port 8080 -sslcert /upgrade/proxy-cert.pem

mkproxy -url http://proxy.example.ibm.com -port 8080
CMMVC8725E The command failed because DNS servers are not configured

mkproxy -url http://123456 -port 8080
CMMVC6523E The URL that you have entered is not valid.

mkproxy -port 8080
CMMVC5707E Required parameters are missing.

mkproxy -url http://proxy.example.ibm.com -port XYZ
CMMVC5706E An invalid argument has been entered for the [port] parameter.

mkproxy -url http://proxy.example.ibm.com -port 8080 -sslcert /upgrade/no-such-file.txt
CMMVC8031E file not found.
```

```
mkproxy -url http://proxy.example.ibm.com -port 8080 -sslcert /upgrade/not-a-cert.txt
CMMVC8794E Invalid certificate file.
```

The resulting output:

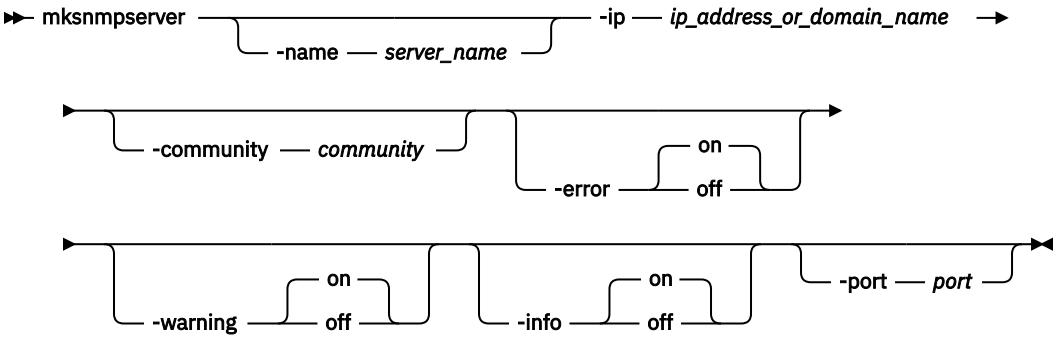
```
No feedback
```

mksnmpserver

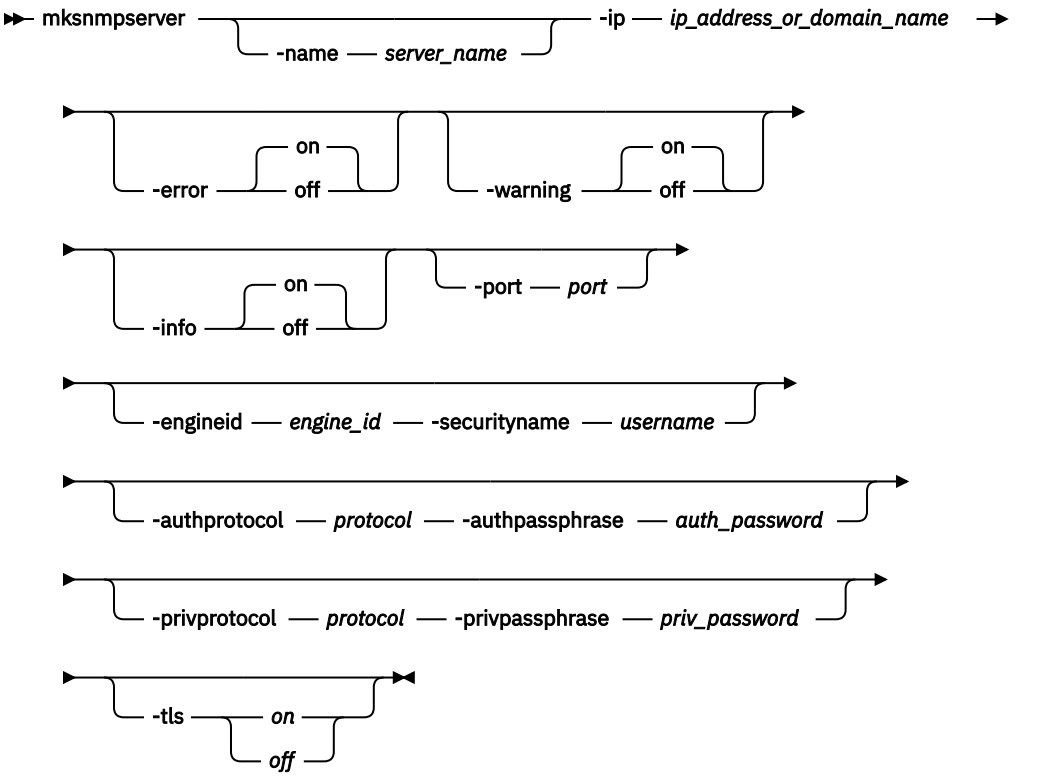
Use the **mksnmpserver** command to create a Simple Network Management Protocol (SNMP) server to receive notifications.

Syntax

This diagram is for SNMPv2.



This diagram is for SNMPv3.



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

-ip *ip_address_or_domain_name*
(Required) Specifies the IP address or domain name of the SNMP server. The IP address must be a valid IPv4 or IPv6 address, or a fully qualified domain name. If you specify a domain name, a DNS server must be configured on your system. You can use the **mkdnsserver** command to configure DNS servers.

-community *community*
(Optional) Specifies the community name for the SNMPv2 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.

-engineid *engine_id*
(Optional) Specifies the engine ID that identifies the local system to the SNMPv3 server. This value has a maximum of 32 bytes, and must meet the requirements outlined in RFC 3411. You must specify **-securityname** with this parameter.

-securityname *username*
(Optional) Specifies the security name for the SNMPv3 server. This value has a max of 32 characters. You must specify **-engineid** with this parameter.

-authprotocol *protocol*
(Optional) Specifies the authentication protocol for the SNMPv3 server. Valid values are **sha1**, **md5 , sha224**, **sha256**, **sha384** or **sha512**. You must specify **-authpassphrase** with this parameter.

-authpassphrase *auth_password*
(Optional) Specifies the authentication password for the SNMPv3 server. This value must be in the range of 8 - 64 characters. You must specify **-authprotocol** with this parameter.

-privprotocol *protocol*
(Optional) Specifies the privacy protocol for the SNMPv3 server. Valid values are **aes128**, **des, aes192**, **aes256**. You must specify **-privpassphrase** with this parameter.

-privpassphrase *priv_password*
(Optional) Specifies the privacy password for the SNMPv3 server. This value must be in the range of 8 - 64 characters. You must specify **-privprotocol** with this parameter.

-tls on | off
(Optional) Specifies whether SNMP traps sent over a TLS connection. Values can be on or off. Default is set to off.

Description

This command creates an SNMPv2 or SNMPv3 server to receive notifications.

The system supports a maximum of 6 SNMP servers.

For syslog servers, SNMP servers, and email users with type set to `local`, the valid combinations of notification types are:

- `error`, `warning`, and `info`
- `error` and `warning`

Note: The port defaults to 10162 instead of 162 if tls is set to on and no other port is specified.

An invocation example to create an SNMPv2 server

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

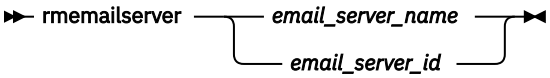
The resulting output:

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

rmemailserver

Use the **rmemailserver** command to delete the specified email server object.

Syntax



Parameters

email_server_name | email_server_id
(Required) Specifies the name or ID of the email server object to be deleted.

Description

Use this command to delete an existing email server object that describes a remote Simple Mail Transfer Protocol (SMTP) email server. You must specify either the current name or the ID of the object that is returned at creation time. Use the **lsemailserver** command to obtain this ID.

Note: Email service stops when the last email server is removed. Use the **startemail** command to reactivate the email and inventory notification function after at least one email server is configured.

An invocation example

```
rmemailserver email4
```

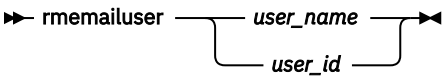
The resulting output:

```
none
```

rmemailuser

Use the **rmemailuser** command to remove a previously defined email recipient from the system.

Syntax



Parameters

user_name | user_id
(Required) Specifies the user ID or user name of the email recipient to remove.

Description

This command removes an existing email recipient from the system.

An invocation example to remove email recipient manager2008

```
rmemailuser manager2008
```

The resulting output:

```
No feedback
```

rmproxy

Use the **rmproxy** command to remove the system-wide web proxy server configuration.

Syntax



Description

This command removes the system-wide web proxy server configuration.

An invocation example

```
rmproxy
```

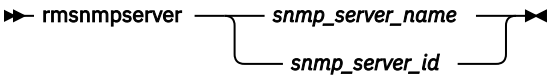
The resulting output:

```
No feedback
```

rmsnmpserver

Use the **rmsnmpserver** command to delete the specified Simple Network Management Protocol (SNMP) server.

Syntax



Parameters

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

Description

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

An invocation example

```
rmsnmpserver snmp4
```

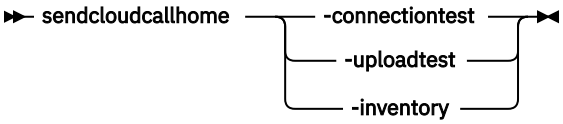
The resulting output:

```
No feedback
```

sendcloudcallhome

Use the **sendcloudcallhome** command to send Call Home information directly to a server in the Cloud.

Syntax



Parameters

-connectiontest
(Required) Initiates a connection check to a server in the Cloud. The command **svcinfo lscloudcallhome** shows the results whether a successful connection is achieved.

-uploadtest
(Required) Uploads a test message to IBM Support. A confirmation receipt is sent via email.

-inventory
(Required) Sends the current inventory to a server in the Cloud.

Description

The Call Home feature must be enabled with the **chcloudcallhome** command.

Note: The **-connectiontest**, **-uploadtest**, and **-inventory** parameters are mutually exclusive.

An invocation example

```
sendcloudcallhome -uploadtest
```

The following output is displayed:

```
No feedback
```

sendinventoryemail

Use the **sendinventoryemail** command to send an inventory email notification to all email recipients able to receive inventory email notifications. There are no parameters for this command.

Syntax



Parameters

There are no parameters for this command.

Description

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

An invocation example

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

```
sendinventoryemail
```

The resulting output:

```
No feedback
```

setemail (Discontinued)

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

startemail

Use the **startemail** command to activate the email and inventory notification function. There are no parameters for this command.

Syntax



Parameters

There are no parameters for this command.

Description

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

Note: This command fails if the **chemail** command has not been used to provide adequate configuration details. The following **chemail** parameters must be specified:

- **reply**
- **contact**
- **primary**
- **location**
- **-organization**
- **-address**
- **-city**
- **-zip**
- **-country**

An invocation example to start the email error notification service

```
startemail
```

The resulting output:

```
No feedback
```

stopemail

Use the **stopemail** command to stop the email and inventory notification function. There are no parameters for this command.

Syntax



Parameters

There are no parameters for this command.

Description

This command stops the email error notification function. No emails are sent to users until the **startemail** command is reissued.

An invocation example to stop the email and inventory notification function

```
stopemail
```

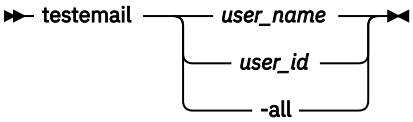
The resulting output:

```
No feedback
```

testemail

Use the **testemail** command to send an email notification to one user or all users of the email notification function to verify correct operation.

Syntax



Parameters

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

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

Description

This command sends test emails to the specified users. The email recipient expects to receive the test email within a specified service time. If the email is not received within the expected time period, the recipient must contact the administrator to ensure that the email settings for the user are correct. If a problem persists, check your product support information.

The email recipient uses the test email to check that the Simple Mail Transfer Protocol (SMTP) name, the IP address, the SMTP port, and the user address are valid.

An invocation example that sends a test email to user ID manager2008

```
testemail manager2008
```

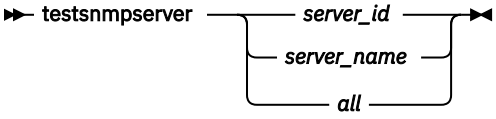
The resulting output

```
No feedback
```

testsnmpserver

Use the **testsnmpserver** command to send a test Simple Network Management Protocol (SNMP) trap. A result message will only be displayed if the SNMP server is configured to use TLS.

Syntax



Parameters

testsnmpserver server_id | server_name | all
(Required) Send a test trap to all configured SNMP servers.

Description

An invocation example to test a single SNMP server by ID

```
testsnmpserver 0
```

The resulting output:

```
No feedback.
```

Chapter 14. Enclosure commands

Use the enclosure commands to manage enclosures and their properties.

addcontrolenclosure

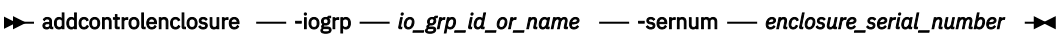
Use the **addcontrolenclosure** command to add control enclosures to the system.

Note:

The **addcontrolenclosure** command is deprecated for all IBM Storage FlashSystem systems except for IBM SAN Volume Controller. The command will fail as additional I/O groups cannot be added to any IBM Storage FlashSystem.

Additional I/O groups can be added to IBM SAN Volume Controller systems only.

Syntax



Parameters

-iogrp *io_grp_id_or_name*
The I/O group in which you want to put the control enclosure.

-sernum *enclosure_serial_number*
The serial number of the control enclosure you want to add.

Description

Use this command to add a control enclosure to the system.

Note: The **addcontrolenclosure** command is supported only if the multiple I/O groups compatibility feature is enabled while adding a second I/O group.

Note: A system with a storage partition is only allowed to have a single I/O group.

An invocation example

```
addcontrolenclosure -iogrp 0 -sernum 2361443
```

The following output is displayed:

```
Enclosure containing Node, id [x], successfully added
```

addcontrolenclosure

Use the **addcontrolenclosure** command to add control enclosures to the system.

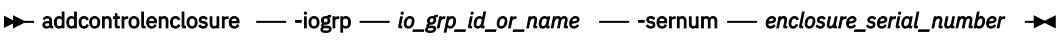
Note:

The **addcontrolenclosure** command is deprecated for all IBM Storage FlashSystem systems except for IBM SAN Volume Controller. The command will fail as additional I/O groups cannot be added to any IBM Storage FlashSystem.

Additional I/O groups can be added to IBM SAN Volume Controller systems only.

© Copyright IBM Corp. 2003, 2025

Syntax



Parameters

-iogrp *io_grp_id_or_name*
The I/O group in which you want to put the control enclosure.

-sernum *enclosure_serial_number*
The serial number of the control enclosure you want to add.

Description

Use this command to add a control enclosure to the system.

Note: The **addcontrolenclosure** command is supported only if the multiple I/O groups compatibility feature is enabled while adding a second I/O group.

Note: A system with a storage partition is only allowed to have a single I/O group.

An invocation example

```
addcontrolenclosure -iogrp 0 -sernum 2361443
```

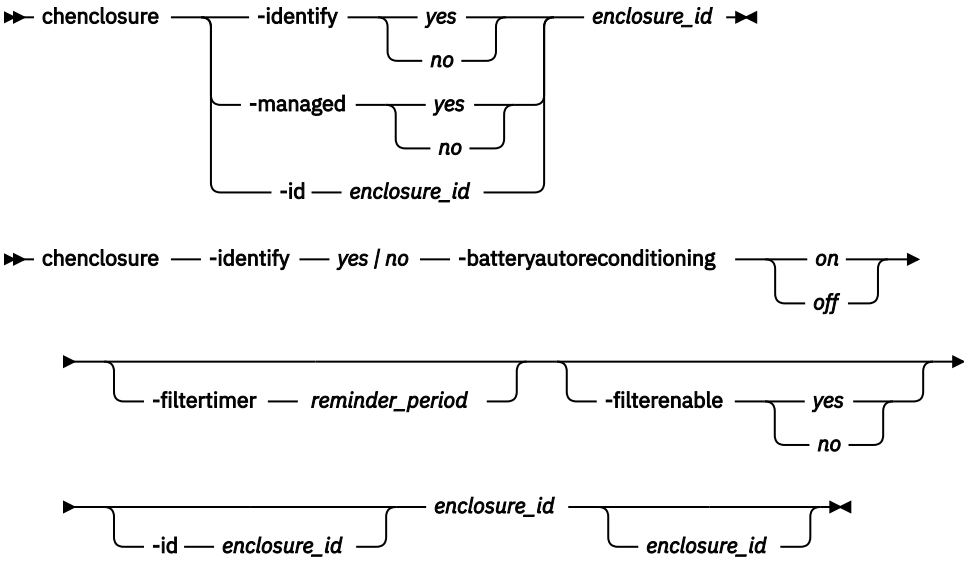
The following output is displayed:

```
Enclosure containing Node, id [x], successfully added
```

chenclosure

Use the **chenclosure** command to modify enclosure properties.

Syntax



Parameters

Note: Optional parameters are mutually exclusive. Exactly one of the optional parameters must be set.

-managed *yes|no*
(Optional) Changes the enclosure to a managed or unmanaged enclosure.

Note: If `chenclosure -managed yes` is run on an enclosure that contains compressed drives, it causes all compressed drives with a state of unused in the enclosure to perform a format. The format must complete before the drives are available for use.

-id ***enclosure_id***
(Optional) Changes the enclosure ID after you replace the enclosure, and enables you to control what is on the front panel.

enclosure_id
(Required) Specifies the enclosure that you want to modify.

-identify ***yes | no***
(Optional) Turns on the identify LED light (but causes no change in the views).

-filtertimer ***reminder_period***
(Optional) Specifies the time (in days) between change air filter reminder events (if the change air filter reminder event is enabled). The default is 90 days.

-filterenable ***yes | no***
(Optional) Enables the change air filter reminder event.

-id ***enclosure_id***
(Optional) Changes the enclosure ID after you replace the enclosure, and enables you to control what is on the front panel.

enclosure_id
(Required) Specifies the enclosure that you want to modify.

enclosure_id
(Optional) Specifies the enclosure that you want to modify.

-batteryautoreconditioning ***on | off***
(Optional) Specifies that you want to turn automated battery reconditioning on or `off`.

Description

Use this command to modify enclosure properties.

An invocation example to change the enclosure ID from 7 to 4

```
chenclosure -id 4 7
```

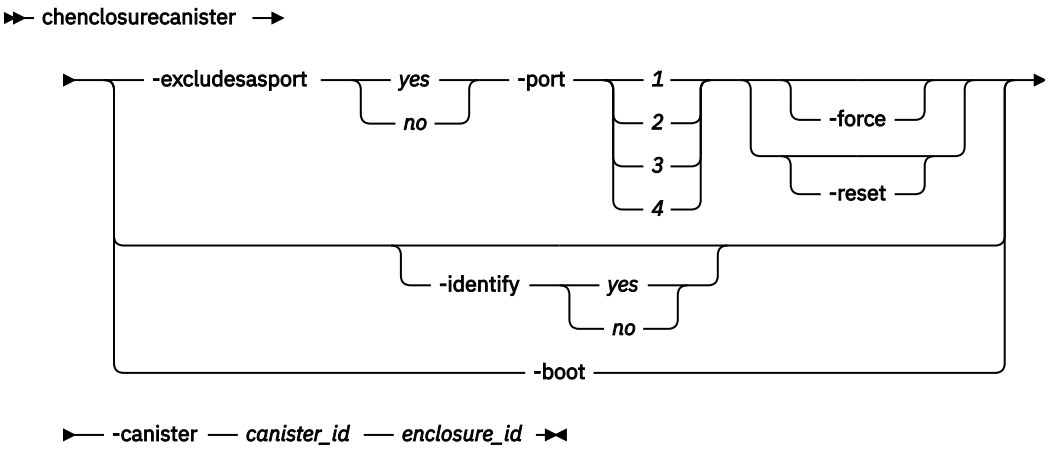
The resulting output:

```
No feedback
```

chenclosurecanister

Use the **chenclosurecanister** command to modify the properties of an enclosure canister.

Syntax



Note: Exactly one of the optional parameters must be set.

1. The **-port** and **-excludesasport** parameters must be specified together.
2. Exactly one of the optional parameters must be set.

Parameters

Note: Optional parameters are mutually exclusive.

-excludesasport *yes | no*
(Optional) Excludes or includes the specified SAS port. The **-port** and **-excludesasport** parameters must be specified together.

You can use the **-force** flag if there are dependent volumes.

Important: Using the **-force** flag might result in loss of access to your data.

-force
(Optional) Forces the enclosure on the canister to be excluded.

Important: Using the **-force** parameter might result in a loss of access to your data. Use it only under the direction of your product support group or representative.

-reset
(Optional) Resets the enclosure on the canister.

Important: Using the **-reset** parameter when the partner canister is not online can lead to a loss of access to your drives (and data). Specify `lsdependentvdisks - enclosure ID -canister ID` to determine whether access to the data for one or more volumes depends on the canister being reset. Using the **-reset** parameter might result in a loss of access to your drives (and data). Use it only under the direction of your product support group or representative.

-identify *yes | no*
(Optional) Changes the state of fault light-emitting diode (LED) either to or from *slow_flashing*.

-canister *canister_id*
Specifies the canister to which you want to apply the change.

enclosure_id
Specifies the enclosure for which the canister is a member.

-boot
Boots the canister if the canister is shut down.

Note: It is a must that the target canister's partner is online to run the boot command. The command fails if the canister partner is unreachable.

Note: The **-boot** parameter is only valid for FlashSystem 9500.

Description

This command enables you to modify the properties of an enclosure canister. You must also designate a port (by using the **-port** parameter) when this parameter is used.

An invocation example to exclude SAS port 1 on canister 2 of enclosure 1

```
chenclosurecanister -excludesasport yes -port 1 -canister 2 1
```

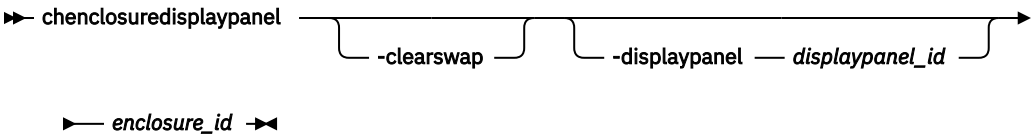
The resulting output:

```
No feedback
```

chenclosuredisplaypanel

Use the **chenclosuredisplaypanel** command to modify the properties for an enclosure display panel.

Syntax



Parameters

-clearswap
(Optional) Specifies that the enclosure display panel swap bit be cleared.

-displaypanel ***sem_id***
(Optional) Specifies the display panel ID to change. The value must be a number.

enclosure_id
(Required) Specifies the enclosure ID for the enclosure that contains the display panel. The value must be a number 1 - 99.

Description

This command modifies the properties for an enclosure display panel.

An invocation example

```
chenclosuredisplaypanel -clearswap -displaypanel 1 3
```

The resulting output:

```
No feedback
```

chenclosurefanmodule

This command is for use with products that have hot-swappable fan assemblies. Use the **chenclosurefanmodule** command to modify the properties of an enclosure fan module.

Syntax



Parameters

-fanmodule fanmodule_id
Specifies the fan module in the enclosure that is identified in the command.

-identify *yes* | *no*
Turns on a flashing light-emitting diode (LED) on the enclosure to make it easier to identify the fan assembly. The values are yes or no.

enclosure_id
Specifies the enclosure that contains the fan module.

Description

Use this command to modify the properties of an enclosure fan module.

An invocation example

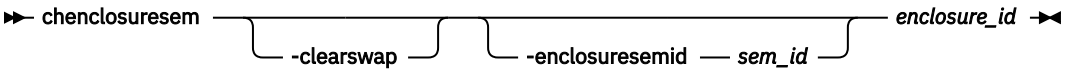
```
chenclosurefanmodule -fanmodule 2 -identify yes 1
```

No output results from using this command.

chenclosuresem

Use the **chenclosuresem** command to modify the properties of an enclosure SEM.

Syntax



Parameters

-clearswap
(Optional) Specifies that the enclosure SEM swap bit be cleared.

-enclosuresemid *sem_id*
(Optional) Specifies the enclosure SEM ID. The value must be a number 1 - 2.

enclosure_id
(Required) Specifies the enclosure ID for the enclosure that contains the SEM. The value must be a number 1 - 99.

Description

This command modifies the properties of an enclosure SEM.

An invocation example

```
chenclosuresem -clearswap -enclosuresemid 1 8
```

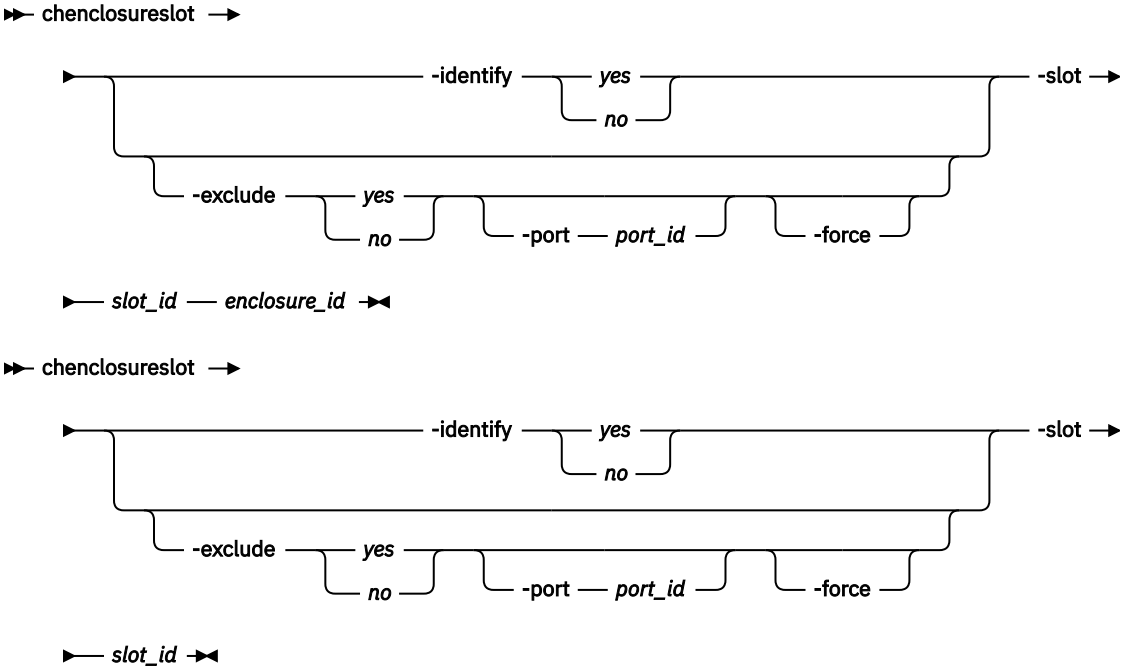
The resulting output:

```
No feedback
```

chenclosureslot

Use the **chenclosureslot** command to modify the properties of an enclosure slot.

Syntax



Note:

1. Optional parameters are mutually exclusive.
2. You can only specify **-port** or **-force** when you also specify **-exclude**.
3. Exactly one of the optional parameters must be set.
4. Using **-force** has an effect on the operation of **-exclude** *yes*.

Parameters

-identify ***yes | no***
(Optional) Change the state of fault light-emitting diode (LED) to or from *slow_flashing*.

-exclude ***yes | no***
(Optional) Ensures that an enclosure slot port is excluded. The following list gives details of the options you can use with this parameter:

- **-exclude** *yes* **-port** *port_id* **-slot** *slot_id enclosureid*: The port that you specify with *port_id* is excluded. If the current state of the port is `excluded_by_enclosure`, `excluded_by_drive`, or `excluded_by_cluster`, this command appears to have no effect. However, if the current state of the port is `online`, then that state changes to `excluded_by_cluster`. The port remains excluded until you rerun this command with *no* selected.

Attention: This command checks for dependent volumes. If issuing this command would result in losing access to data, then the command fails and an error message displayed. You can use the **-force** flag to ignore these errors, but this might result in loss of access to data.

- **-exclude** *no* **-port** *port_id* **-slot** *slot_id enclosureid*: The port is put into `online` state, provided there are no other reasons to exclude the port. If you issue this command when the port is `online`, then it has no effect. However, if you issue this command when the port is `excluded`, then the port state does one of the following items:
  - Change to `online` status immediately.
  - Change to `online` status after all other reasons for the port to be excluded have been removed.
- **-exclude** *yes* | *no* **-slot** *slot_id enclosureid*: If you issue this command without defining a port, the command simultaneously acts on both ports.

-port ***1*** **|** ***2***
(Optional) Specifies the port on the canister to be excluded. If it is not specified, **-exclude** acts on both ports.

-force
(Optional) Forces the port on the canister to be excluded.

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

-slot ***slot_id***
(Required) Specifies the slot ID. The value must be a number from 1 - 92.

The slots are numbered 1 (leftmost) to 24 (rightmost) when viewed from in front of the enclosure, in 24-slot enclosures. In 12-slot enclosures, the slots are arranged in numerical order in three rows with four slots. For example, the:

- First row contains slots 1, 2, 3, and 4 (in that order)
- Second row contains slots 5, 6, 7, and 8 (in that order)
- Third row contains slots 9, 10, 11, and 12 (in that order)

(Required) Specifies the slot ID. The slots are numbered 1 (leftmost) to 12 (rightmost) when viewed from in front of the enclosure.

Description

These commands enable you to modify the properties of an enclosure slot.

Note: For systems with NVMe drives, the physical removal of faulty drives (after the device is turned off) is necessary to prevent the accidental powering on of the drive that might cause a recurrence of the issue.

An invocation example to turn on the identify LED on slot 7 in enclosure 1

```
chenclosureslot -identify yes -slot 7 1
```

The resulting output:

```
No feedback
```

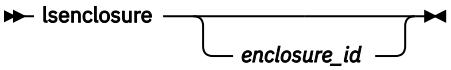
sainfo chenclosurevpd (Deprecated)

The **sainfo chenclosurevpd** command is deprecated. Use the **chvpd** command instead.

lsenclosure

Use the **lsenclosure** command to view a summary of the enclosures.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

enclosure_id
Detailed information for the enclosure that you specify.

Description

This command displays a summary of the enclosures (including status information for canisters and power and cooling units, and other enclosure attributes). This table shows the possible outputs:

Table 82. ***lsenclosure*** *output*

Attribute	Description
id	Indicates the ID of the enclosure.
status	Indicates whether an enclosure is visible to the SAS network: • online if a managed or unmanaged enclosure is visible. • offline if a managed enclosure is not visible, and other fields hold their last known values. • degraded if an enclosure is visible, but not down both strands.
type	Indicates the type of enclosure: • control • expansion
managed	Indicates whether the enclosure is managed: • yes • no
IO_group_id	Indicates the I/O group the enclosure belongs to; blank if canisters are connected to two different I/O groups.

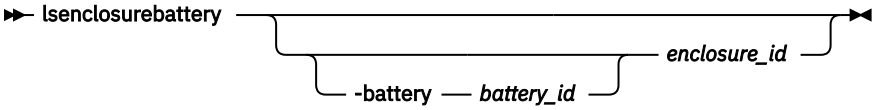
*Table 82. **lsenclosure** output (continued)*	
Attribute	**Description**
IO_group_name	Indicates the I/O group the enclosure belongs to; blank if canisters are connected to two different I/O groups.
error_sequence_number	Indicates the error log number of the highest priority error for this object. This attribute is typically blank; however, if a problem exists (for example, the status has degraded), then it contains the sequence number of that error.
product_MTM	Indicates the product machine type and model.
serial_number	Indicates the serial number of the enclosure. This serial number is the product serial number, which indicates the enclosure and its contents. The enclosure has its own serial number, which is embedded in the FRU_identity 11S data.
FRU_part_number	Indicates the FRU part number of the enclosure.
FRU_identity	Indicates the 11S serial number that combines the manufacturing part number and the serial number.
total_canisters	Indicates the maximum number of canisters for this enclosure type.
online_canisters	Indicates the number of canisters that are contained in this enclosure that are online.
total_PSUs	Indicates the number of power and cooling units in this enclosure.
online_PSUs	Indicates the number of power-supply units (PSUs) contained in this enclosure that are online.
drive_slots	Indicates the number of drive slots in the enclosure.
machine_part_number	Indicates the enclosure's machine part number.Blank.
machine_signature	Indicates a machine signature unique to the control enclosure and representing the serial number and machine part number. The format is a hyphenated string of 19 hexadecimal characters. **Remember:** Expansion enclosures do not have a machine signature.
PIB_FRU_part_number	Indicates the FRU part number for the Power® Interposer Board (PIB).
PIB_FRU_identity	Indicates the FRU ID for the PIB.
encrypted	Indicates whether the enclosure and drives are encrypted.
filter_timer	Indicates the length of time (in days) between filter events.
filter_enabled	Indicates whether the filter timer is enabled. The values are `yes` and `no`.
filter_reset_date	Indicates the date and time, in the format, in the format `YYMMDDHHMMSS`, that an air filter event is triggered.
site_id	This value is blank.
site_name	This value is blank.
total_sems	Indicates the total number of secondary expander modules (SEMs) that are in the system. The value must be a number 0 - 2.
online_sems	Indicates the total number of SEMs in the system that are online. The value must be a number 0 - 2.

Table 82. **lsenclosure** output (continued)	
Attribute	**Description**
battery_auto_reconditioning	Indicates the automated battery reconditioning state. The values are on or off.

lsenclosurebattery

Use the **lsenclosurebattery** command to display information about the batteries. The batteries are located in the node canisters.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-battery *battery_id*
(Optional) Provides a detailed view of the specified battery. Valid only when an enclosure is specified.

enclosure_id
(Optional) Lists the batteries for the specified enclosure.

Description

This command displays information about the batteries, which are located in the node canisters. The concise view displays a line for each battery slot in every enclosure whether a battery exists for that slot or not. Batteries are not shown for expansion enclosures.

This table shows possible outputs.

Table 83. lsenclosurebattery outputs	
Attribute	**Description**
enclosure_id	Identifies the identity of the enclosure that contains the battery.
battery_id	Identifies the battery in the enclosure.
status	Identifies the status of the battery: • online indicates that the battery is present and working as usual. • degraded indicates that the battery is present but not working as usual. • offline indicates that the battery cannot be detected.

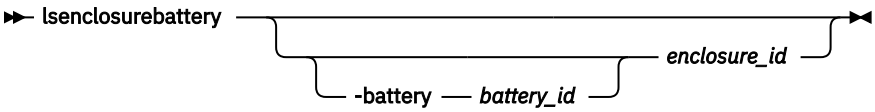
Table 83. lsenclosurebattery outputs (continued)	
Attribute	**Description**
recondition_needed	Identifies that the battery needs to be reconditioned or it must start reconditioning soon. **Remember:** If this message is permanent, there might be an error that prevents the recondition from starting.
percent_charged	Indicates the charge of battery (in a percentage).
end_of_life_warning	Identifies the battery's end of life (with a warning noise). The values are `yes` and `no`. **Important:** Replace the battery.
FRU_part_number	Identifies the FRU part number of the battery.
FRU_identity	Identifies the `11S` number, combining the manufacturing part number and the serial number.
firmware_level	The version of the microcode image that is installed on the battery.
error_sequence_number	Indicates the error log (or event log) number of the highest priority error for this object. This output field is typically blank. However, if there is a problem (for example, the status is `degraded`), then it contains the sequence number of that error event.
remaining_charge_capacity_mAh	Identifies the battery's remaining charge capacity in milliampere-hour (mAh).
full_charge_capacity_mAh	Identifies the fully charged capacity of the battery in mAh (this value diminishes as the battery ages).
compatibility_level	Identifies the battery driver software must support this level to operate with this battery - this value comes from the battery vital product data (VPD).
last_recondition_timestamp	Identifies a system timestamp, in the format `YYMMDDHHMMSS`, for when the last successful calibration of the gas gauge occurred, where: • `Y` is year • (The first) `M` is month • `D` is day • `H` is hour • (The second) `M` is minute • `S` is second
powered_on_hours	Identifies the number of hours the battery is in a powered node (not necessarily the same node).
cycle_count	Identifies the number of charge or discharge cycles that are performed on the battery.
fault_led	Indicates whether the fault or identify light-emitting diode (LED) is on, off, or flashing. The values are: • on, which indicates that the fault LED is on • `off`, which indicates that both the fault and identify LED are off • `slow_flashing`, which indicates that the identify LED is blinking

Table 83. lsenclosurebattery outputs (continued)	
Attribute	**Description**
redundant	Indicates whether you can remove this battery. The values are yes and no.
redundant	Indicates whether you can remove this battery. The values are yes and no.
swapped	Indicates that the battery is replaced or reseated. The values are yes and no.

lsenclosurebattery

Use the **lsenclosurebattery** command to display information about the batteries. The batteries are located in the node canisters.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-battery ***battery_id***
(Optional) Provides a detailed view of the specified battery. Valid only when an enclosure is specified.

enclosure_id
(Optional) Lists the batteries for the specified enclosure.

Description

This command displays information about the batteries, which are located in the node canisters. The concise view displays a line for each battery slot in every enclosure whether a battery exists for that slot or not. Batteries are not shown for expansion enclosures.

This table shows possible outputs.

Table 84. lsenclosurebattery outputs	
Attribute	**Description**
enclosure_id	Identifies the identity of the enclosure that contains the battery.
battery_id	Identifies the battery in the enclosure.

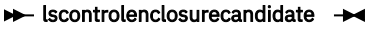
Table 84. lsenclosurebattery outputs (continued)	
Attribute	**Description**
status	Identifies the status of the battery: • `online` indicates that the battery is present and working as usual. • `degraded` indicates that the battery is present but not working as usual. • `offline` indicates that the battery cannot be detected.
recondition_needed	Identifies that the battery needs to be reconditioned or it must start reconditioning soon. **Remember:** If this message is permanent, there might be an error that prevents the recondition from starting.
percent_charged	Indicates the charge of battery (in a percentage).
end_of_life_warning	Identifies the battery's end of life (with a warning noise). The values are `yes` and `no`. **Important:** Replace the battery.
FRU_part_number	Identifies the FRU part number of the battery.
FRU_identity	Identifies the 11S number, combining the manufacturing part number and the serial number.
firmware_level	The version of the microcode image that is installed on the battery.
error_sequence_number	Indicates the error log (or event log) number of the highest priority error for this object. This output field is typically blank. However, if there is a problem (for example, the status is `degraded`), then it contains the sequence number of that error event.
remaining_charge_capacity_mAh	Identifies the battery's remaining charge capacity in milliampere-hour (mAh).
full_charge_capacity_mAh	Identifies the fully charged capacity of the battery in mAh (this value diminishes as the battery ages).
compatibility_level	Identifies the battery driver software must support this level to operate with this battery - this value comes from the battery vital product data (VPD).
last_recondition_timestamp	Identifies a system timestamp, in the format `YYMMDDHHMMSS`, for when the last successful calibration of the gas gauge occurred, where: • `Y` is year • (The first) `M` is month • `D` is day • `H` is hour • (The second) `M` is minute • `S` is second
powered_on_hours	Identifies the number of hours the battery is in a powered node (not necessarily the same node).
cycle_count	Identifies the number of charge or discharge cycles that are performed on the battery.

Table 84. lsenclosurebattery outputs (continued)	
Attribute	**Description**
fault_led	Indicates whether the fault or identify light-emitting diode (LED) is on, off, or flashing. The values are: • on, which indicates that the fault LED is on • off, which indicates that both the fault and identify LED are off • slow_flashing, which indicates that the identify LED is blinking
redundant	Indicates whether you can remove this battery. The values are yes and no.
redundant	Indicates whether you can remove this battery. The values are yes and no.
swapped	Indicates that the battery is replaced or reseated. The values are yes and no.

lscontrolenclosurecandidate

Use the **lscontrolenclosurecandidate** command to display a list of all control enclosures you can add to the current system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

Description

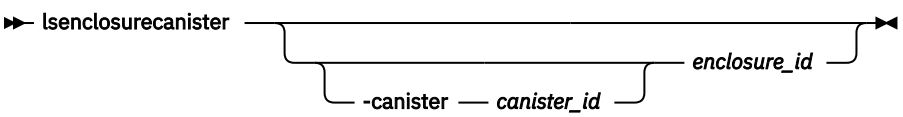
Table 85 on page 325 provides the possible values that are applicable to the attributes that are displayed as data in the output views.

Table 85. **lscontrolenclosurecandidate** attribute values	
Attribute	**Value**
serial_number	Indicates the serial number for the enclosure.
product_MTM	Indicates the MTM for the enclosure.

lsenclosurecanister

Use the **lsenclosurecanister** command to view a detailed status for each canister in an enclosure.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

enclosure_id
(Optional) Lists the canisters for the specified enclosure.

Description

This command displays a detailed status for each canister in an enclosure. This table shows the possible outputs:

Table 86. ***lsenclosurecanister*** *output*

Attribute	Description
enclosure_id	Indicates the enclosure that contains the canister.
canister_id	Indicates which of the canisters in the enclosure it is.
status	Indicates the status of the canister. The values are: • online indicates that the canister is present and working normally. • degraded indicates that the canister is present but not working normally • offline indicates that the canister could not be detected.
type	Indicates the type of canister. The values are node or expansion.
node_id	Indicates the node that corresponds to this canister. The value is blank if the canister is not a node, or if the node is offline, or if the node is not part of the clustered system.
node_name	Indicates the node that corresponds to this canister. The value is blank if the canister is not a node, or if the node is offline, or if the node is not part of the clustered system.
temperature	Indicates the temperature of the canister in degrees centigrade.
internal_fault_LED_off	Indicates the status of the internal_fault_LED_off. The values are on, off, or slow-flashing.
controller_fault_LED_off	Indicates the status of the controller_fault_LED_off. The values are on, off, or slow-flashing.
fault_LED	Indicates the status of the fault_LED. The values are on, off, or slow-flashing.
SES_status	Indicates the Small Computer System Interface (SCSI) status of the connection between the device and the canister. The values are online and offline.
FRU_part_number	Indicates the field-replaceable unit (FRU) part number of the canister.
FRU_identity	Indicates the 11S number that combines the manufacturing part number and the serial number.

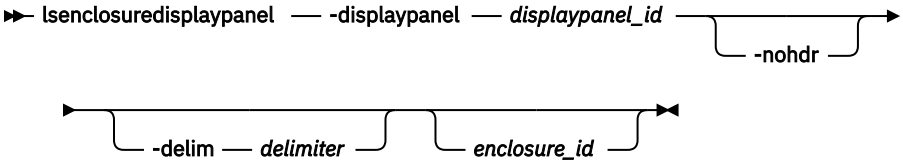
*Table 86. **lsenclosurecanister** output (continued)*	
Attribute	**Description**
WWNN	Indicates the Fibre Channel (FC) worldwide node name (WWNN) of the canister (node canisters only).
temperature	Indicates the temperature of the canister (in degrees Celsius). If the temperature goes below `0`, `s0` is displayed. The value must be in the range `0` - `245`. **Remember:** The temperature value is not an ambient temperature value. It is an internal temperature sensor value.
error_sequence_number	Indicates the error log number of the highest priority error for this object. This value is typically blank; however, if there is a problem (for example, the status is degraded), then the value contains the sequence number of that error.
SAS_port_1_status	Indicates whether there is damage to the cable between SAS ports: • `online` • `offline` • `excluded` (meaning logged in but cannot communicate with the canister) • `degraded` (meaning the SAS cable is not fully functional) • Blank (which can appear on control canisters; see **lsportsas**)
SAS_port_2_status	Indicates whether there is damage to the cable between SAS ports: • `online` • `offline` • `excluded` (meaning logged in but cannot communicate with the canister) • `degraded` (meaning the SAS cable is not fully functional) • Blank (which can appear on control canisters; see **lsportsas**)
firmware_level	Indicates the firmware level of the Small Computer System Interface (SCSI) Enclosure Services (SES) code, or canister firmware version, that is running on the canister.
firmware_level_2	Indicates the version of the first other microcode image (canister bootloader version) that is installed on the canister.
firmware_level_3	Indicates the version of the second other microcode (canister PSoC version) image that is installed on the canister.
firmware_level_4	Indicates the version of the third other microcode image (canister CPLD version) that is installed on the canister.
firmware_level_6	Indicates the canister CPLD version firmware level that is installed on the canister.
fan_1_status	Indicates the status of fan 1: • `online` indicates that the fan is present and working • `degraded` indicates that the fan is present, but is not working • `offline` indicates that the fan is not present or cannot be detected

*Table 86. **lsenclosurecanister** output (continued)*	
Attribute	**Description**
fan_2_status	Indicates the status of fan 2: • online indicates that the fan is present and working • degraded indicates that the fan is present, but is not working • offline indicates that the fan is not present or cannot be detected
fan_1_swapped	Indicates the swapped bit is or is not available for fan 1: • yes indicates that the fan is available • no indicates that the fan is not available
fan_2_swapped	Indicates the swapped bit is or is not available for fan 2: • yes indicates that the fan is available • no indicates that the fan is not available
fan_1_fault_LED	Indicates the state of the combined fault and identify light-emitting diodes (LEDs) of fan 1: • off indicates that there is no fault mode • slow_flashing indicates identify mode • on indicates fault mode **Remember:** When the LED is in identify mode, it conceals whether there is a fault present because it always flashes. When you remove it from identity mode, the LED becomes on or off.
fan_2_fault_LED	The state of the combined fault and identify LEDs of fan 2: • off indicates that there is no fault mode • slow_flashing indicates identify mode • on indicates fault mode **Remember:** When the LED is in identify mode, it conceals whether there is a fault present because it always flashes. When you remove it from identity mode, the LED becomes on or off.
fan_1_fru_identity	Indicates the 11S number for fan 1 that is a combination of the fan: • Manufacturing part number • Serial number (for fan 1)
fan_2_fru_identity	Indicates the 11S number for fan 2 that is a combination of the fan: • Manufacturing part number • Serial number (for fan 2)
fan_1_fru_part_number	Indicates the field-replaceable unit (FRU) part number for fan 1.
fan_1_fru_part_number	Indicates the field-replaceable unit (FRU) part number for fan 2.

lsenclosuredisplaypanel

Use the **lsenclosuredisplaypanel** command to display information about the display panel in an enclosure.

Syntax



Parameters

-displaypanel *displaypanel_id*
(Required) Specifies the display panel ID for the display panel that is being displayed. The value must be a number.

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

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

-delim *delimiter*
(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a 1-byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

enclosure_id
(Required) Specifies the enclosure ID for the enclosure data that is being displayed. The value must be a number 1 - 99.

Description

This command displays information about the display panel in an enclosure.

This table provides the attribute values that can be displayed as output view data.

Table 87. ***lsenclosuredisplaypanel*** *output*

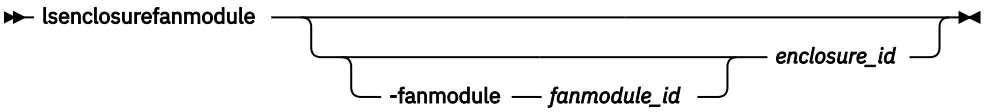
Attribute	Description
enclosure_id	Indicates the enclosure ID for the enclosure that contains the display panel. The value must be a number 1 - 99.
displaypanel_id	Indicates the ID of the display panel that is in the enclosure. The value must be a number.
status	Indicates the display panel status for the display panel that is in the enclosure. The values are: • `online` • `degraded` • `offline`

Table 87. **lsenclosuredisplaypanel** output (continued)	
Attribute	**Description**
error_sequence_number	Indicates the event log sequence number for the current event that is logged against the secondary expander module (SEM). The value is blank if there is no event to log.
FRU_part_number	Indicates the FRU part number of the display panel. The value must be a 7-character numeric string.
FRU_identity	Indicates the FRU identity of the display panel. The value must be a 22-character alphanumeric string.

lsenclosurefanmodule

Use the **lsenclosurefanmodule** command to report the status of each fan module and the contained fans in an enclosure.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-fanmodule *fanmodule_id*
(Optional) Specifies the ID of the fan module for which data is displayed. The possible values are 1 or 2, and any other value returns no output.

enclosure_id
(Optional) Specifies the ID of the enclosure for which data is displayed.

Description

The command reports status for fan modules and the contained fans in an enclosure.

This table provides the possible values that are applicable to the attributes that are displayed as data in the output views.

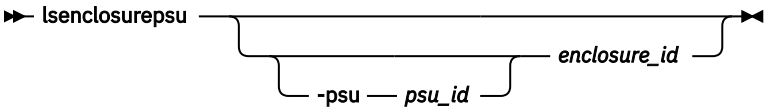
Table 88. **lsenclosurefanmodule** attribute values	
Attribute	**Value**
enclosure_id	Indicates the enclosure ID of the enclosure that contains the fan module.
fan_module_id	Indicates the fan module ID of the fan module that is in the enclosure. The possible values are 1 or 2.
status	Indicates the combined status of the fan module and any contained fans. The values are: • online • offline • degraded

Table 88. **lsenclosurefanmodule** attribute values (continued)	
Attribute	**Value**
error_sequence_number	Indicates the event log sequence number of the current event that is logged against the fan module. It is blank if there is no current event.
FRU_part_number	Indicates the part number of the fan module.
FRU_identity	Indicates the FRU identity of the fan module.
fault_LED	Indicates the status of the fault light-emitting diode (LED) on the fan module: • on, which indicates that the LED is on • off, which indicates that the LED is off • unknown, which indicates that the LED status is unknown

lsenclosurepsu

Use the **lsenclosurepsu** command to view information about each power-supply unit (PSU) in the enclosure.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-psu *psu_id*
(Optional) Valid only when the **enclosure_id** is specified. Provides a detailed view of the PSU for the specified enclosure.

enclosure_id
(Optional) Lists the PSUs for the specified enclosure.

Description

This command displays information about each power-supply unit (PSU) in the enclosure. This table shows the possible outputs:

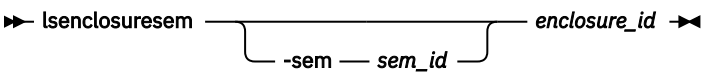
Table 89. **lsenclosurepsu** output	
Attribute	**Description**
enclosure_id	Indicates the ID of the enclosure that contains the PSU.
psu_id	Indicates the ID of the PSU in the enclosure.

*Table 89. **lsenclosurepsu** output (continued)*	
Attribute	**Description**
status	Indicates the status of the power and cooling unit in the enclosure: • online indicates that a PSU is present and working normally. • offline indicates that a PSU cannot be detected. • degraded indicates that a PSU is present but not working normally.
input_failed	• on indicates that no usable input power is detected from the power distribution unit. • off indicates that the input power is OK.
output_failed	• on indicates that no usable output power is detected from the power supply unit. • off indicates that the output power is OK.
input_power	• ac indicates that the power supply requires AC power input. • dc indicates that the power supply requires DC power input. • unknown indicates that the power supply is not known or cannot be determined.
fan_failed	• on indicates that if the AC, DC, and fan LEDs are all on, there is a PSU fault. If only the fan LED is on, then there is a fan failure. • off indicates that the fans in this PSU are OK.
redundant	Indicates (yes or no) whether you can disconnect the power supply. • If the value is yes, there are sufficient PSUs to power the enclosure, so the PSU can be disconnected and removed from the enclosure. • If the value is no, do not remove the PSU from the enclosure and do not disconnect it from its power source. This will make the enclosure to lose power and shut down.
error_sequence_number	Indicates the error log (or event log) number of the highest priority error for this object. This value is typically blank. However, if there is a problem (for example, the status is degraded), then it contains the sequence number of that error event.
FRU_part_number	Indicates the FRU part number of the PSU.
FRU_identity	Indicates the 11S number, combining the manufacturing part number and the serial number.

lsenclosuresem

Use the **lsenclosuresem** command to display the status (or any pertinent data) about secondary expander modules (SEM) in a 5U92 system.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-sem ***sem_id***
(Optional) Specifies the SEM ID for the SEM data that is being displayed.

enclosure_id
(Required) Specifies the enclosure ID for the enclosure data that is being displayed. The value must be an integer 1 - 99.

Description

This command displays the status (or any pertinent data) about SEM in a 5U92 system.

A 5U92 system is a 5U enclosure that can contain up to 92 3.5-inch drives (but must be used in an expansion enclosure).

This table provides the attribute values that can be displayed as output view data.

Table 90. ***lsenclosuresem*** *output*

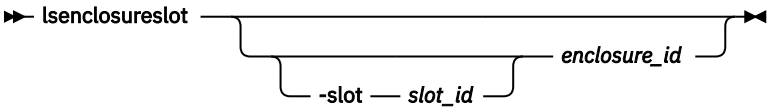
Attribute	Description
id	Indicates the enclosure ID of the enclosure that contains the disk drawer. The value must be a number 1 - 99.
sem_id	Indicates the SEM ID for the SEM that is in the enclosure. The value must be a number, 1 or 2.
status	Indicates the SEM status for the SEM that is in the enclosure. The values are: • `online, which indicates that the SEM is online` • `degraded, which indicates that the SEM is degraded` • `offline, which indicates that the SEM is offline`
expander_1_status	Indicates the status of the first or lowest order expander index. The values are: • `online, which indicates that the SEM is online` • `degraded, which indicates that the SEM is degraded` • `offline, which indicates that the SEM is offline`
expander_2_status	Indicates the status of the second expander index (or *e+1*, where *e* is the lowest-order or first expander index). The values are: • `online, which indicates that the SEM is online` • `degraded, which indicates that the SEM is degraded` • `offline, which indicates that the SEM is offline`
error_sequence_number	Indicates the event log sequence number of the current event that is logged against the SEM. The value is blank if there is no error to log.
FRU_part_number	Indicates the field-replaceable unit (FRU) part number for the SEM. The value must be a 7-character numeric string.
FRU_identity	Indicates the FRU ID for the SEM. The value must be a 22-character alphanumeric string.
firmware_level_1	Indicates the SCSI Enclosure Services (SES) firmware level of the lowest order expander index. The value must be a 22-character alphanumeric string.

Table 90. **lsenclosuresem** output (continued)	
Attribute	**Description**
firmware_level_2	Indicates the bootloader firmware level of the lowest order expander index. The value must be a 22-character alphanumeric string.
firmware_level_3	Indicates the SES firmware level of the second lowest order expander index. (or *s+1*, where *s* is the lowest-order or first expander index). The value must be a 22-character alphanumeric string.
firmware_level_4	Indicates the bootloader firmware level of the second lowest order expander index. (or *b+1*, where *b* is the lowest-order or first expander index). The value must be a 22-character alphanumeric string.
firmware_level_5	Indicates the complex programmable logic device (CPLD) firmware level of the lowest order expander index. The value must be a 22-character alphanumeric string.
firmware_level_6	Indicates the CPLD firmware level of the second lowest order expander index (or *c+1*, where *c* is the lowest-order or first expander index). The value must be a 22-character alphanumeric string.

lsenclosureslot

Use the **lsenclosureslot** command to view information about each drive slot in the enclosure.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-slot***slot_id***
(Optional) Specifies the slot to display information for (it provides a detailed view for that enclosure slot). This parameter is valid only when an enclosure is specified. The value must be a number in the range 1 - 92.

Note:

If slot information is requested for a slot that does not exist on the specified enclosure, value that is displayed is blank.

enclosure_id
(Optional) Lists slots for that enclosure. Must be specified whether **-slot** is used.

Description

This command displays information about each drive slot in the enclosure, such as whether a drive is present, and the port status for that drive. This table shows the possible outputs:

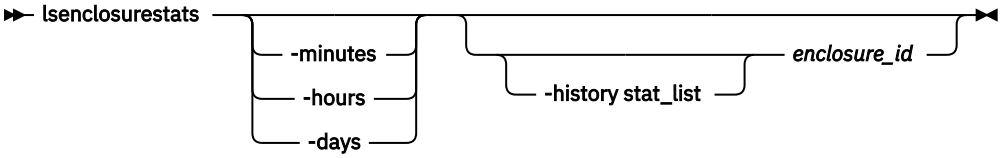
*Table 91. **`lsenclosureslot`** output*

Attribute	Description
enclosure_id	The identity of the enclosure that contains the drive slot.
slot_id	Identifies which of the drive slots in the enclosure it is.
port_1_status	The status of enclosure slot port 1. If the port is bypassed for multiple reasons, only one is shown. In order of priority, they are: • `online` indicates that the enclosure slot port 1 is online • `excluded_by_drive` indicates that the drive excluded the port • `excluded_by_enclosure` indicates that the enclosure excluded the port • `excluded_by_system` indicates that the system excludes the port This field is blank for enclosures with NVMe drives.
port_2_status	The status of enclosure slot port 2. If the port is bypassed for multiple reasons, only one is shown. In order of priority, they are: • `online` indicates that the enclosure slot port 2 is online. • `excluded_by_drive` indicates that the drive excluded the port. • `excluded_by_enclosure` indicates that the enclosure excluded the port. • `excluded_by_system` indicates that the system excludes the port. This field is blank for enclosures with NVMe drives.
fault_LED	The state of the combined fault and identify light-emitting diodes (LEDs): • `off` indicates no fault • `slow_flashing` indicates the `identify` mode **Note:** When the LED is in identify mode, it conceals whether there is a fault present (it always flashes). When you remove it from identity mode, the LED is turned on or `off`. • on indicates fault
powered	Indicates whether the slot is powered on. The values are `yes` or no.
drive_present	Indicates if a drive is in the slot. The drive can be working, dead, or powered off. The values are `yes` (present) or no (empty).
drive_id	Indicates the ID of the drive in the slot; blank if there is no drive present, or if there is a drive present but it is offline and unmanaged.
error_sequence_number	Indicates the error log number of the highest priority error for this object. This value is typically blank. However, if there is a problem (for example, the status is degraded), then it contains the sequence number of that error.

lsenclosurestats

Use the **lsenclosurestats** command to display the most recent values (averaged) of all enclosure statistics. It can also display a history of those values for any subset of the available statistics.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-history stat_list
(Optional) Produces a history of values for enclosure statistics.

-minutes
(Optional) Displays stats at minutes granularity.

-hours
(Optional) Displays stats at hours granularity.

-days
(Optional) Show stats at days granularity.

Note:

- Flags are mutually exclusive. In absence of any the live (5 sec) samples are displayed and used with or without **-history** parameter.
- Each granularity will show up to 60 samples per stat when used with **-history** parameter.

enclosure_id
(Optional) Indicates the unique enclosure identifier (a number in the range 1 - 99).

Description

Remember: This command cannot be used for products that do not support environmental statistics.

If you specify **-history stat_list** you must also specify *enclosure_id*. Filtering is supported for the concise view but not the detailed view.

Multiple statistical histories can be requested. The limit is the current maximum number of different statistical names published in the concise view. The concise view defines the output order.

For the detailed view, enclosure power is averaged over thirty seconds to provide immediate power.

Note: Averaging applies only to populated samples

Enclosure power is not averaged in the output if **-history** is specified.

The following is an invocation example for products that do not support environmental statistics - a message is displayed:

```
lsenclosurestats
```

This is the resulting output:

```
CMMVC6051E An unsupported action was selected.
```

This table displays information about chassis-specific enclosure properties and shows possible outputs for products that support environmental statistics.

Table 92. lsenclosurestats outputs

Attribute	Description
enclosure_id	Indicates the enclosure identifier; it can be a numeric character in the range 1 - 264.
sample_time	Indicates the time during which the sample occurred.
stat_name	Indicates the name of the statistical field.
stat_current	Indicates the current value of the statistical field.
stat_peak	Indicates the peak value of the statistic field. The last 5 minutes are used for samples.
stat_peak_time	Indicates the time that the peak occurred.
stat_value	Indicates the value of the statistic.

Remember: Filtering is supported on the `enclosure_id` and `stat_name` fields by using the concise view.

This table provides the possible values that are applicable to the values that are displayed for the **`stat_name`** attribute.

Table 93. Stat_name field values

Value	Description
power_w	Displays the power that is consumed in watts.
temp_c	Displays the ambient temperature in Celsius.
temp_f	Displays the ambient temperature in Fahrenheit.

lssasfabric

Use the **`lssasfabric`** command to see which canisters are visible to a node, and the order of these canisters.

Syntax



Parameters

`-delim` ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

`-nohdr`
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

Description

Use this command to see which canisters are visible to a node, and the order of these canisters. This table describes possible outputs:

Table 94. ***lssasfabric*** *output*

Attribute	Description
enclosure_id	Indicates the identity of the enclosure the strand goes to.
canister_id	Indicates the canister in the enclosure that the strand goes to.
canister_port_id	Indicates the canister port that the strand goes to.
control_enclosure_id	Indicates the identity of the enclosure that the strand comes from.
node_canister_id	Indicates the identity of the canister the strand comes from.
node_canister_port_id	Indicates the node canister port that the strand is from. This must be the same as the chain ID.
position	Indicates the position in the strand or chain.
IO_group_id	Indicates the I/O group that the strand belongs to. This must be the same as the enclosure IO group.
IO_group_name	Indicates the I/O group the strand belongs to. This must be the same as the enclosure IO group.
node_id	Indicates the identity of the node that the strand is from. This is the same physical object as the node_canister
node_name	The name of the node that the strand is from. This is the same physical object as the node_canister.

resetleds

Use the **resetleds** command to simultaneously switch off all light-emitting diodes (LEDs) in the system, including node battery LEDs.

Syntax



Parameters

None.

Description

The **resetleds** command simultaneously switches off all LEDs in the system, including node battery LEDs. This ensures that any identity LED switched on is the only one in the system that is switched on. This command works only on LEDs that are on systems that can communicate, which means they are online or of a supported type. This command fails if an object is offline, or if the enclosure is an unsupported type. This command does not affect LEDs:

- On independently-controlled objects

- On offline objects
- With hardware-only controls

An invocation example

```
resetleds
```

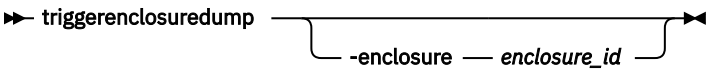
The resulting output:

```
No feedback
```

triggerenclosuredump

Use the **triggerenclosuredump** command to force the specified enclosure or enclosures to dump data.

Syntax



Parameters

-enclosure *enclosure_id*
(Optional) The ID of the enclosure you want to dump.

Description

This command requests the canisters in the enclosure or enclosures that are specified to dump data. Then, the dumped data is collected and moved to `/dumps/enclosure` on the nodes that are connected to the enclosure. There is one file for each canister successfully dumped. The files might be on different nodes. Dumps provide information about your system to your product support personnel, who have tools to interpret the dump data. Use the **cpdumps** command to copy the files from the system. This command does not disrupt access to the enclosures. The system limits the number of enclosure statesaves in the directory to 20 per node.

To trigger enclosure dumps from `enclosure 5`

```
triggerenclosuredump -enclosure 5
```

The resulting output:

```
The data is dumped to the /dumps/enclosure directory if command is successful.
```

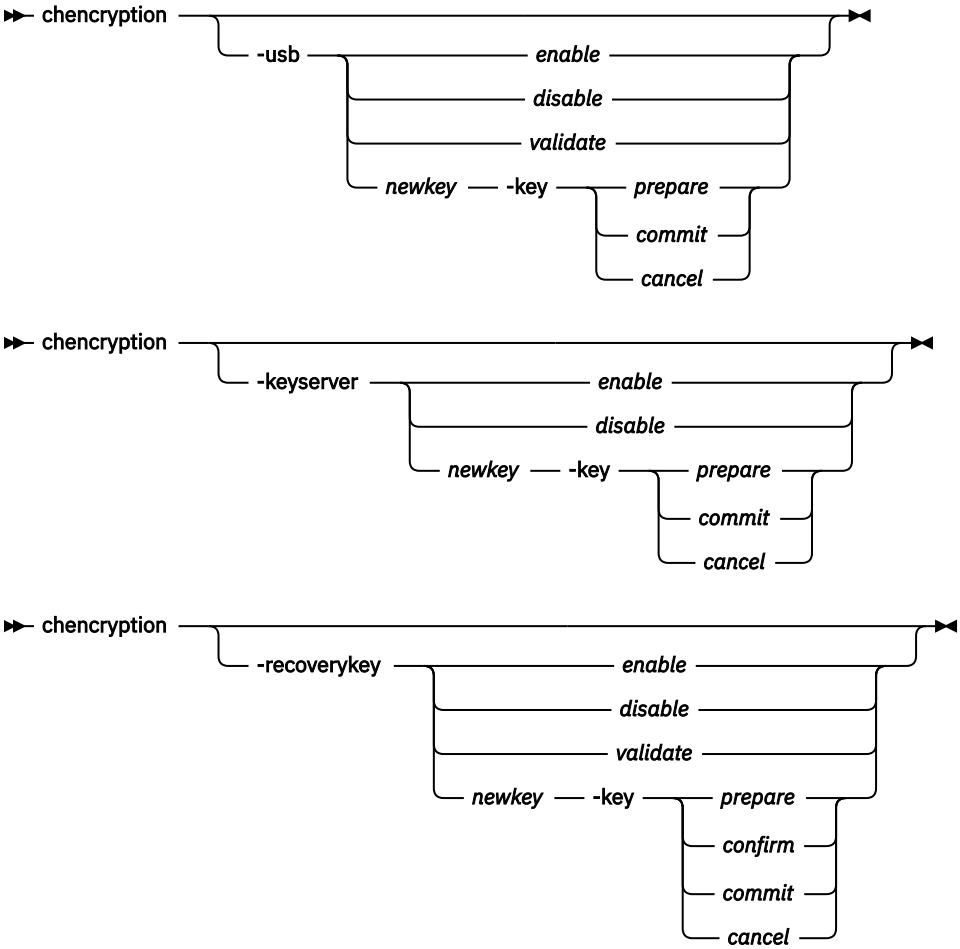
Chapter 15. Encryption commands

Use the encryption and security commands are used to create, change, or list system encryption feature details.

chencryption

Use the **chencryption** command to manage the encryption state of the system.

Syntax



Parameters

-usb ***enable | disable | validate | newkey***

(Required if you do not specify **-keyserver**, **-recoverykey**) Specifies whether encryption with USB flash drives is enabled (or disabled) or the encryption keys are validated. You can also create new encryption keys that are stored on Universal Serial Bus (USB) flash drives.

-usb ***enable***

Enables encryption capability on the system with USB flash drives. Then specify `-usb` *newkey* to create new keys. Use this command when the system has encryption hardware and encryption licenses (for example, the **lsencryption** value for `status` is set to `licensed`).

© Copyright IBM Corp. 2003, 2025

-usb *disable*
Disables the encryption capability of the system with USB flash drives. If no encryption key is prepared this operation is complete and no further action is needed. Do not use this command if an encryption key is prepared or encrypted objects exist.

Remember: This removes all knowledge of the encryption keys from the system, but does not remove any encryption key files from USB flash drives.

-usb *validate*
Verifies that encryption keys are present on the USB flash drive and makes sure that the keys match the system encryption keys. Use this command when encryption is enabled and encryption keys exist (for example, **lsencryption** value for `usb_rekey` is set to no).

-usb *newkey*
Generates a new encryption key on USB flash drives that are attached to the system. Use this command only if the minimum number of USB flash drives that can be used as key material stores are attached to the system (as reported by **lsportusb**). When you specify this parameter, the **-key** option must also be supplied.

-keyserver *enable | disable | newkey*
(Required if you do not specify **-usb**, **-recoverykey**) Specifies whether encryption with key servers is enabled (or disabled). You can also create new encryption keys that are stored on key servers.

-keyserver *enable*
Enables encryption capability on the system with key servers. Use this command when the system has encryption hardware and encryption licenses (for example, the **lsencryption** value for `keyserver_status` is set to `licensed`).

-keyserver *disable*
Disables the encryption capability of the system with key servers. If no encryption key is prepared this operation is complete and no further action is needed. Do not use this command if an encryption key is prepared or encrypted objects exist.

-keyserver *newkey*
Generates a new encryption key on the primary key server that is attached the system. You must also specify **-key** when you specify this parameter.

-recoverykey *enable | disable | validate | newkey*
(Required if you do not specify **-usb**, **-recoverykey**) Specifies the encryption task to be performed for the encryption recovery key. You can enable, disable, or validate the encryption recovery key. You can also create a new encryption recovery key for the system.

Note: The **-recoverykey** parameter can only be entered when at least one of USB flash drives or key servers have been configured already.

-key *prepare | confirm | commit | cancel*
Manages the creation of a new key for the system. When you specify the value confirm, the **-recoverykey** option must also be supplied.

Description

Use this command to manage the encryption state of the system. You must specify either **-usb**, **-keyserver**, **-internal**, or **-recoverykey**.

You can use this command to enable or disable encryption using USB flash drives or key servers (but you cannot disable encryption if there are any encrypted objects). You can also enable or disable the encryption recovery key for the system. There are four possible arguments:

- `enable`, which enables encryption
- `disable`, which disables encryption
- `validate`, which validates encryption key files on all of the USB flash drives currently installed in the system. The **lsportusb** command should be used to check the results of this command when **-usb** is specified. Refer to "lsportusb" on page 178 for more information. When **-recoverykey** is specified, this validates the recovery key entered matches the system's recovery key.

Note: The `validate` option does not apply to key server encryption.

- `newkey`, which specifies a new key for encryption

You can also perform a rekey of the external USB key or key server key material, which is divided into three stages:

- `prepare`, which generates new keys and sets up the system to change encryption keys during apply
- confirm, which confirms that the supplied recovery key matches the correct key for the system (only specify this with the **`-recoverykey`** option)
- `commit`, which includes applying new keys (and copying key material)
- `cancel`, which rolls back the key setup that is performed during the `prepare` and cancels the rekey request

You can use both USB flash drive and key server encryption in parallel on the same system. However, you must configure and administer these encryption methods independently.

An invocation example

```
chencryption -usb enable
```

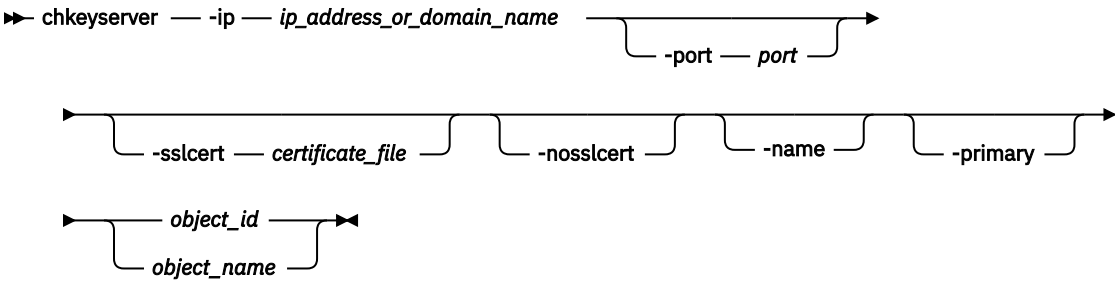
The resulting output:

```
No feedback
```

chkeyserver

Use the **`chkeyserver`** command to change the attributes for a key server object.

Syntax



Parameters

`-ip` ***ip_address_or_domain_name***
(Required) Specifies the key server's IP address or domain name. This must be a valid IPv4 or IPv6 address, or a fully qualified domain name. IPv6 addresses can be zero compressed. If you specify a fully qualified domain name, a DNS server must be configured on your system. You can use the **`mkdnsserver`** command to configure DNS servers.

`-port` ***port***
(Optional) Specifies the key server's TCP/IP port. The value must be a number 1 - 65535. The default value is the same as the default port used for key servers of the currently enabled type.

`-sslcert` ***certificate_file***
(Optional) Specifies the key server's internally-signed certificate. The value must be a file path string.

`-nosslcert`
(Optional) Specifies the removal of the key server internally-signed certificate.

`-name`
(Optional) Specifies the key server object name. The value must be an alphanumeric string.

-primary
(Optional) Specifies the primary key server.

object_id | object_name
(Required) Specifies the object name or ID that you want to modify.

Description

This command changes the attributes for a key server object.

When a primary key server is configured, that key server must be defined before a rekey operation occurs. A primary object (such as a server) can be configured at any time when a defined primary server is present. A rekey operation without a defined primary key server fails.

An invocation example

```
chkeyserver -primary vardy2
```

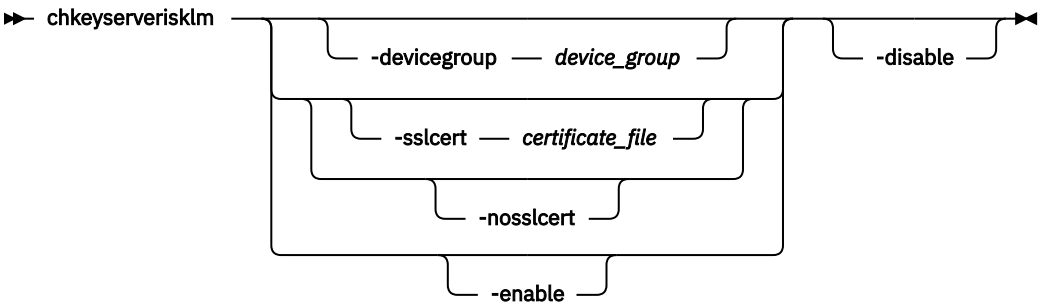
The resulting output:

```
No feedback
```

chkeyserverisklm

Use the **chkeyserverisklm** command to change the system-wide IBM Security Key Lifecycle Manager key server configuration.

Syntax



Parameters

-devicegroup ***device_group***
(Optional) Specifies a specific device group that the system uses with a key server. The value must be an alphanumeric string no more than 16 characters long. The default device group name is FLASHSYSTEM

Note: The specified device name must begin with a letter (not a number) and cannot contain an underscore.

-sslcert ***certificate_file***
(Optional) Specifies the certificate authority (CA) certificate for the key server. This parameter cannot be used with **-nosslcert**. The value must specified in base64-encoded PEM format.

-nosslcert
(Optional) Specifies that the CA certificate on the key server is deleted. This parameter cannot be used with **-sslcert**.

-enable
(Optional) Enables the specified key server type.

-disable
(Optional) Disables the specified key server type.

Important: Do not specify **-disable** with other parameters.

Description

This command changes the system-wide IBM Security Key Lifecycle Manager key server configuration.

An invocation example

```
chkeyserverisklm -devicegroup JVAR_IBRA -sslcert /dumps/CA_certificate.pem -enable
```

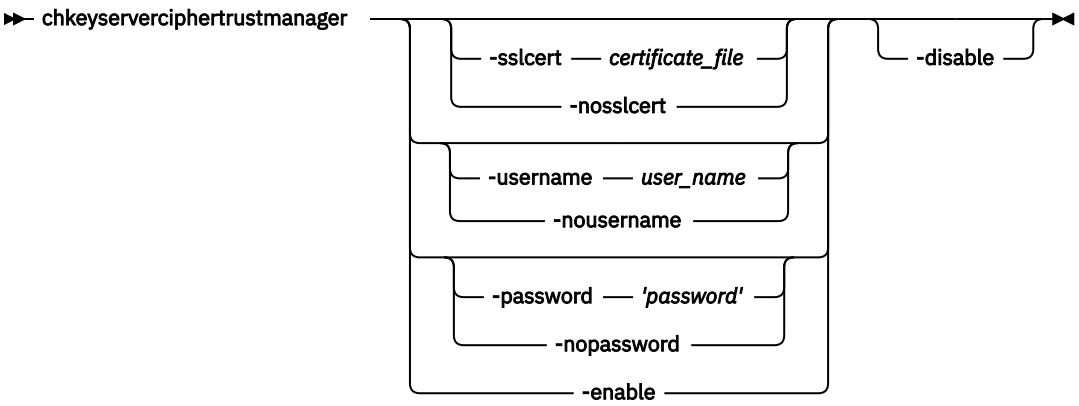
The resulting output:

```
No feedback
```

chkeyserverciphertrustmanager

Use the **chkeyserverciphertrustmanager** command to change the system-wide Thales CipherTrust Manager or Gemalto SafeNet KeySecure key server configuration.

Syntax



Parameters

-sslcert *certificate_file*
(Optional) Specifies the CA certificate that was used to sign the key server certificate.

-nosslcert
(Optional) Removes the existing CA certificate.

-username *user_name*
(Optional) Specifies the username that is used for authentication with the key servers. The value must be an alphanumeric string with a maximum of 64 characters.

-nousername
(Optional) Clears the username that is used for authentication with the key servers.

-password *'password'*
(Optional) Specifies the password that is used for authentication with the key servers. The value must be an alphanumeric string with a maximum of 64 characters. You must enclose the password in single quotation marks.

-nopassword
(Optional) Clears the password that is used for authentication with the key servers.

-enable
(Optional) Enables the Thales CipherTrust Manager key server type.

-disable
(Optional) Disables the Thales CipherTrust Manager key server type.

Description

This command changes the Thales CipherTrust Manager key server configuration.

Keep the following items in mind when you use this command:

- The parameters **-sslcert** and **-nosslcert** are mutually exclusive.
- The parameters **-username** and **-nousername** are mutually exclusive.
- The parameters **-password** and **-nopassword** are mutually exclusive.
- The parameters **-sslcert**, **-username**, and **-password** can be set while the status is disabled (or enabled).
- The parameter **-disable** is mutually exclusive with all other parameters.

An invocation example

```
chkeyserverciphertrustmanager -sslcert /tmp/keysecureCA.pem -enable
```

The resulting output:

```
No feedback
```

chkeyserverkeysecure (Deprecated)

The **chkeyserverkeysecure** command has been deprecated. Use the **chkeyserverciphertrustmanager** command instead.

lsencryption

Use the **lsencryption** command to display system encryption information.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

Description

Use this command to display output that is related to the system encryption state.

This table describes possible outputs.

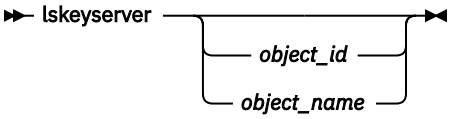
*Table 95. **lsencryption** output*	
Attribute	**Value**
status	Indicates the system USB encryption status. • not_supported, which indicates that the system has no supported encryption function. • not_licensed, which indicates that the system supports USB encryption but not all licenses are installed. • licensed, which indicates that the system has licenses that are installed for all encryption-capable hardware. • enabled, which indicates that system encryption by using USB flash drives is working and ready to create encrypted storage.
error_sequence_number	Indicates the event log sequence number of any problem that affects encryption. If there is no problem, it is blank.
usb_rekey	Indicates the state of the Universal Serial Bus (USB) rekey process. • no, which indicates that there is no rekey process ongoing, but keys exist. • no_key, which indicates that there is no rekey process and keys do not exist. • prepared, which indicates that a rekey process is active and the system prepares a new key that is waiting for this command to be issued: **chencryption -usb newkey -key commit**. • committing, which indicates that a commit is in progress.
usb_key_copies	Indicates the number of USB devices that prepared keys are written to.
usb_key_filename	Indicates the name of the file that contains the current encryption key.
usb_rekey_filename	Indicates the name of the file that contains the current prepared encryption key.
keyserver_status	Indicates the encryption status for key server encryption. The values are: • not_supported, which indicates that the system has no supported encryption function. • not_licensed, which indicates that the system supports key server encryption but not all licenses are installed. • licensed, which indicates that the system has licenses that are installed for all encryption-capable hardware. • enabled, which indicates that system encryption by using key servers is working and ready to create encrypted storage.
keyserver_rekey	Indicates the state of the key server rekey process. The values are: • no, which indicates that there is no rekey process ongoing, but keys exist. • no_key, which indicates that there is no rekey process and keys do not exist. • prepared, which indicates that a rekey process is active and the system prepares a new key that is waiting for this command to be issued: **chencryption -keyserver newkey -key commit**. • committing, which indicates that a commit is in progress.
keyserver_pmk_uid	Indicates the UID for the key server.
keyserver_rekey_pmk_uid	Indicates the UID (after a rekey process) for the key server.

Table 95. **lsencryption** output (continued)	
Attribute	**Value**
recovery_key_status	Indicates the status for the encryption recovery key. The possible values can be: • not_supported • not_licensed • licensed • enabled
recovery_key_rekey	Indicates the state of the encryption recovery key re-key process. The possible values can be: • no • no_key • prepared • confirmed • committing
recovery_key_name	Indicates the name of the committed recovery key. The string can be in the *recoverykey_<cluster_id_alias>_<key_uid>_<cluster_name>* format.
recovery_key_rekey_name	Indicates the name of the prepared recovery key. The string can be in the `recoverykey_<cluster_id_alias>_<key_uid>_<cluster_name>` format.

lskeyserver

Use the **lskeyserver** command to display the key servers that are available to the system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

object_id | object_name
(Optional) Specifies the object name or ID that you want to display details for.

Description

This command displays all key servers that are available to the system.

This table provides the attribute values that can be displayed as output view data.

Table 96. lskeyserver output	
Attribute	**Description**
id	Indicates the key server ID.

Table 96. lskeyserver output (continued)	
Attribute	**Description**
name	Indicates the key server name.
status	Indicates the key server status type. The values are: • `online` • `degraded` • `offline`
err_seq_num	Indicates the event log sequence number of the highest priority problem that affects the key server.
IP_address	Indicates the IP address or domain name of the key server.
port	Indicates the key server TCP/IP port.
type	Indicates the key server type.
primary	Indicates whether the server is a primary server.
cert_set	Indicates whether a certificate exists for this key server object.
certificate	Indicates a human-readable description of the SSL certificate. The value reads `0 fields` if there is no certificate.

lskeyserverisklm

Use the **lskeyserverisklm** command to display the system-wide IBM Security Key Lifecycle Manager key server configuration.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

Description

This command displays the system-wide IBM Security Key Lifecycle Manager key server configuration.

This table provides the attribute values that can be displayed as output view data.

Table 97. lskeyserverisklm output	
Attribute	**Description**
status	Indicates the key server status type. The values are: • `disabled` • `enabled_inactive` • `prepared` • `enabled_active`

Table 97. lskeyserverisklm output (continued)	
Attribute	**Description**
device_group	Indicates the device group. The value is a 16-character alphanumeric string.
certificate	Indicates a human readable description of the server SSL certificate that is generated by the system.

lskeyserverciphertrustmanager

Use the **lskeyserverciphertrustmanager** command to display the system-wide Thales CipherTrust Manager or Gemalto SafeNet KeySecure key server configuration.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

Description

This table shows the possible output:

Table 98. ***lskeyserverciphertrustmanager*** *Output*	
Attribute	**Description**
status	Displays the status of this key server type. The possible values are: • `disabled` • `enabled_inactive` • `prepared` • `enabled_active`
username_set	Displays whether a username is set for this key server type. The values are `yes` or no.
password_set	Displays whether a password is set for this key server type. The values are `yes` or no.
certificate	Displays a human readable description of the whole server SSL certificate (as generated by `openssl x509 -text`). This value is blank if no certificate exists.

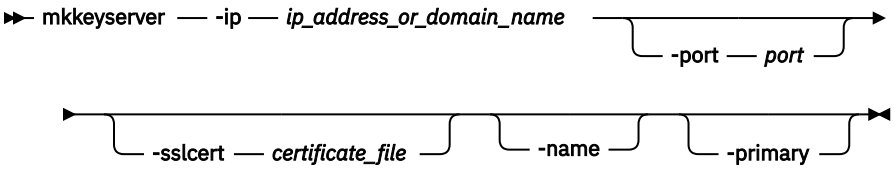
lskeyserverkeysecure (Deprecated)

The **lskeyserverkeysecure** command has been deprecated. Use the **lskeyserverciphertrustmanager** command instead.

mkkeyserver

Use the **mkkeyserver** command to create a key server object.

Syntax



Parameters

-ip ***ip_address_or_domain_name***
(Required) Specifies the key server's IP address or domain name. This must be a valid IPv4 or IPv6 address or a fully qualified domain name. IPv6 addresses can be zero compressed. If you specify a fully qualified domain name, a DNS server must be configured on your system. You can use the **mkdnsserver** command to configure DNS servers.

-port ***port***
(Optional) Specifies the key server's TCP/IP port. The value must be a number 1 - 65535. The default value is the same as the default port used for key servers of the currently enabled type.

-sslcert ***certificate_file***
(Optional) Specifies the key server's internally-signed certificate. The value must be a file path string.

-name
(Optional) Specifies the key server object name. The value must be an alphanumeric string.

-primary
(Optional) Specifies the primary key server.

Description

This command creates a key server object.

The primary key server object is created by specifying **-primary**. If key management is enabled, you must use the primary key server object to create keys.

Note: When a primary key server is configured, that key server must be defined before a rekey operation occurs. A primary object (such as a server) can be configured at any time when a defined primary server is present. When you create keys, the system uses the key server that is configured as the primary key server. For multi-primary key server configurations, any key server can be selected as the primary. A rekey operation without a defined primary key server fails.

When a key server object is created, it is automatically validated. If the validation is not successful, the command fails and an error message is displayed.

An invocation example

```
mkkeyserver -ip servername.example.com -sslcert /tmp/isklm_public_server_cert.pem -primary
```

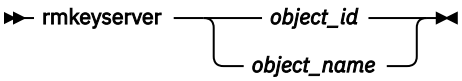
The resulting output:

```
Key Server, id [0], successfully created
```

rmkeyserver

Use the **rmkeyserver** command to remove a key server object.

Syntax



Parameters

object_id | object_name
(Required) Specifies the object name or ID to remove.

Description

This command removes a key server object.

An invocation example

```
rmkeyserver 1
```

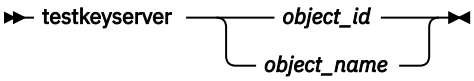
The resulting output:

```
No feedback
```

testkeyserver

Use the **testkeyserver** command to test key server objects.

Syntax



Parameters

object_id | object_name
(Required) Specifies the object name or ID to be validated.

Description

The **testkeyserver** command validates that a configured keyserver is operating correctly, and is available to provide the correct encryption key if needed. It will complete successfully if:

- All nodes can connect to the key server, and successfully authenticate.
- The key server responds to Key Management Interoperability Protocol (KMIP) requests within a reasonable time.
- The key server returns the correct key for this system, which would enable access to encrypted data.

If any of these steps fails, the command will produce an error message.

This command ensures that the key server can connect to both canisters, which means that it makes sure that the:

- IP path is clear

- Certificates are exchanged correctly

This command also ensures that the key server has its configured keys available.

An invocation example

```
testkeyserver 0
```

The resulting output:

```
The key server task completed successfully.
```

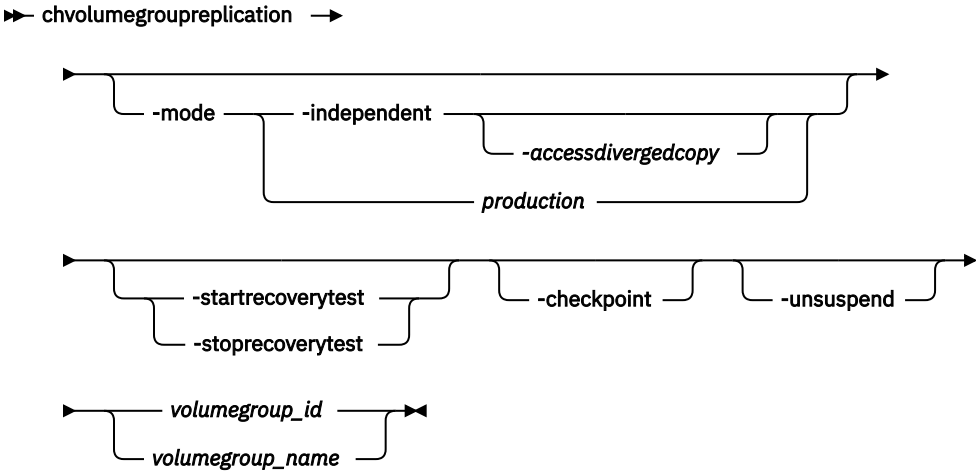
Chapter 16. Policy-based replication commands

Policy-based replication uses volume group and replication policies to automatically deploy and manage the replication. Policy-based replication significantly simplifies configuring, managing, and monitoring replication between two systems. Use the following set of commands to change, create or delete the replication policies.

chvolumegroupreplication

Use the **chvolumegroupreplication** command to change the replication properties of a volume group.

Syntax



Parameters

-mode ***independent | production***
(Optional) Specifies the replication mode for the local volume group. Value can be *independent* or *production.*

-accessdivergedcopy
(Optional) Allow independent access to be enabled to a recovery copy of a volume group when the image that exists may be one that has diverged from the production copy of the volume group. This parameter is required to enable independent access before a synchronisation has completed following a recovery test or restarting replication from the independent state.

Note:

- Changing the replication mode to independent while a recovery test is active will promote the volumes in-use for the recovery test to independent copies.
- Starting a recovery test is prohibited if the recovery copy does not have a valid recovery point.
- By default, independent access is prohibited if the recovery copy does not have a valid recovery point, but this can be bypassed using the **-accessdivergedcopy** option.

-checkpoint
(Optional) Specifies that the system should start a replication checkpoint. The **checkpoint_achieved** field in the **lsvolumegroupreplication** view indicates when the recovery copy of the volume group has a recovery point no earlier than the time this command completes.

© Copyright IBM Corp. 2003, 2025

This command can be issued on the production or recovery copy of a volume group.

This parameter cannot be specified with any other parameters and cannot be used for volume groups in a HA partition.

-startrecoverytest
(Optional) Starts a disaster recovery test on the local system, enabling host access to the volumes in the specified volume group. This operation can only be used on the system that has the recovery copy of the volume group.

-stoprecoverytest
(Optional) Stops an active disaster recovery test on the local system.

-unsuspend

(Optional) Specifies that replication is allowed to continue after a system failure, such as a dead domain or T3 recovery.

Note: This operation can only be used after a system failure, when the location status of the local system is reported as suspended in the **lsvolumegroupreplication** view. To recover from this state, you have the following options for recovery:

- Before you fail over to the recovery system, you can validate data on the production copy. To failover to the recovery system, issue chvolumegroupreplication -mode independent on the recovery system. This action is recommended when the production system is reporting the suspended state.
- On the system that is reporting the suspended state, issue the **chvolumegroupreplication -unsuspend** command to restart replication in the current direction. Use this command when the recovery copy of the data is reporting the suspended state.

volumegroup_id | volumegroup_name
(Mandatory) Specifies the ID or name of the local volume group to be modified.

A concise invocation example

```
chvolumegroupreplication -mode independent volumegroup4
```

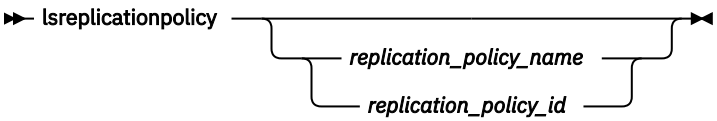
The concise resulting output:

```
No feedback.
```

lsreplicationpolicy

Use the **lsreplicationpolicy** command to display a list of all replication policies on the system.

Syntax



Parameters

replication_policy_name | replication_policy_id
(Optional) Specifies the name or ID of the replication policy.

Note: When this parameter is used, a detailed view of the specified replication policy is displayed. If this parameter is not specified then a concise view of all replication policies is displayed.

Description

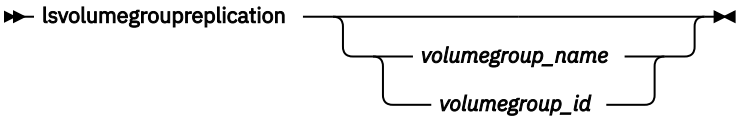
This table describes possible outputs.

*Table 99. **lsreplicationpolicy** output*	
Attribute	**Value**
id	Indicates the ID of the replication policy.
name	Indicates the name of the replication policy.
topology	Indicates the topology of the replication policy.
rpo_alert	Indicates the RPO alert threshold time (in seconds) for the replication policy.
volume_group_count	Indicates the number of volume groups that are associated with this replication policy.
location1_system_id	Indicates the system ID in location 1 of the replication policy.
location1_system_name	Indicates the system name in location 1 of the replication policy.
location1_iogrp_id	Indicates the ID of the I/O group in location 1 of the replication policy.
location2_system_id	Indicates the system ID in location 2 of the replication policy.
location2_system_name	Indicates the system name in location 2 of the replication policy.
location2_iogrp_id	Indicates the ID of the I/O group in location 2 of the replication policy.
partition_count	Indicates the number of partition associated with the replication policy. **Note:** A replication policy assigned to one or more partitions will display a count to indicate the number of partitions that use the replication policy and the volume group count will be zero.
partition_name	Indicates the name of the partition associated with the replication policy.
partition_id	Indicates the id of the partition associated with the replication policy.

lsvolumegroupreplication

Use the **lsvolumegroupreplication** command to display the replication information for the volume group.

Syntax



Parameters

volumegroup_name | volumegroup_id
(Optional) Specifies the name or ID of the volume group.

Note: When this parameter is used, a detailed view of the replication information for the specified volume group is displayed. If this parameter is omitted, then a concise view of the replication information for all volume groups using policy-based replication is displayed.

Description

This table describes possible outputs.

*Table 100. **lsvolumegroupreplication** output*	
Attribute	**Value**
id	Indicates the system ID of the local volume group.
name	Indicates the name of the volume group.
replication_policy_id	Indicates the ID of the replication policy that is associated with the volume group.
replication_policy_name	Indicates the name of the replication policy that is associated with the volume group.
local_location	Indicates the location number (1 or 2) for the local system in the replication policy. This identifier can be used to determine which of the location1 or location2 fields correspond to the local system.
local_partition_location	Indicates the position (upper or lower) of the local partition in the replication topology diagram. This identifier can be used to determine which of the upper/lower replication mode fields correspond to the local system.
location1_system_id	Indicates the ID of the system in the first position in the topology diagram.
location1_system_name	Indicates the name of the system in the first position in the topology diagram.
location2_system_id	Indicates the ID of the system located in the second position in the topology diagram.
location3_system_id	Indicates the ID of the system located in the third position in the topology diagram.
location2_system_name	Indicates the name of the system in the second position in the topology diagram.
location3_system_name	Indicates the name of the system in the third position in the topology diagram.
location1_replication_mode	Indicates the replication mode of the volume group on the system in location 1 of the replication policy. **Note:** • This field is populated only if the volume group is associated with a replication policy that specifies systems in specific locations. • This field is deprecated and may be removed in a future release. • Use **upper_partition_replication_mode** or **lower_partition_replication_mode** instead.
location2_replication_mode	Indicates the replication mode of the volume group on the system in location 2 of the replication policy. **Note:** • This field is populated only if the volume group is associated with a replication policy that specifies systems in specific locations. • This field is deprecated and may be removed in a future release. • Use **upper_partition_replication_mode** or **lower_partition_replication_mode** instead.
upper_partition_replication_mode	Indicates the replication mode of the volume group on the systems in the upper partition in the topology diagram.
lower_partition_replication_mode	Indicates the replication mode of the volume group on the systems in the lower partition in the topology diagram.

*Table 100. **lsvolumegroupreplication** output (continued)*	
Attribute	**Value**
location1_volumegroup_id	Indicates the ID of the volume group on the system in location 1 of the topology diagram.
location2_volumegroup_id	Indicates the ID of the volume group on the system in location 2 of the topology diagram.
location3_volumegroup_id	Indicates the ID of the volume group on the system in location 3 of the topology diagram.
location1_status	Indicates the health status of the system in location 1 of the topology diagram. The following statuses can be displayed: • healthy • unhealthy • disconnected • suspended • blank A status of *unhealthy* indicates that there is an error that affects replication. A status of *suspended* indicates that user intervention is required.
location2_status	Indicates the health status of the system in location 2 of the topology diagram. The following statuses can be displayed: • healthy • unhealthy • disconnected • suspended • blank A status of *unhealthy* indicates that there is an error that affects replication. A status of *suspended* indicates that user intervention is required.
location1_running_recovery_point	Indicates the total time (in second) the volume group in location 1 of the replication policy is behind the production copy of the volume group. This field is populated when replication is running.
location2_running_recovery_point	Indicates the total time (in second) the volume group in location 2 of the replication policy is behind the production copy of the volume group. This field is populated when replication is running.
location3_running_recovery_point	Indicates how many seconds the volume group in this location is behind the production copy of the volume group. This field is populated when replication is running.
location1_fixed_recovery_point	Indicates the time of the recovery point for the volume group in this location. It is displayed as blank or in YYMMDDHHMMSS format in UTC.
location2_fixed_recovery_point	Indicates the time of the recovery point for the volume group in this location. It is displayed as blank or in YYMMDDHHMMSS format in UTC.
location3_fixed_recovery_point	Indicates the time of the recovery point for the volume group in this location. It is displayed as blank or in YYMMDDHHMMSS format in UTC.

*Table 100. **lsvolumegroupreplication** output (continued)*	
Attribute	**Value**
location1_within_rpo	Indicates whether the current recovery point for the volume group in this location is within the recovery point objective (RPO) of the replication policy associated with the production copy of the volume group. The following statuses can be displayed: • yes • no • blank **Note:** If the volume group exceeds the RPO of the replication policy, an alert is logged on the production system. This field is only populated when the current replication mode of the volume group is **recovery**.
location2_within_rpo	Indicates whether the current recovery point for the volume group in this location is within the recovery point objective (RPO) of the replication policy associated with the production copy of the volume group. The following statuses can be displayed: • yes • no • blank **Note:** If the volume group exceeds the RPO of the replication policy, an alert is logged on the production system. This field is only populated when the current replication mode of the volume group is **recovery**.
location3_within_rpo	Indicates whether the current recovery point for the volume group in this location is within the recovery point objective (RPO) of the replication policy associated with the production copy of the volume group. The following statuses can be displayed: • yes • no • blank **Note:** If the volume group exceeds the RPO of the replication policy, an alert will be logged on the production system. This field is only populated when the current replication mode of the volume group is **recovery**.
location1_sync_required	Indicates the total amount of data that must be copied to update the recovery point for the volume group in this location.
location2_sync_required	Indicates the total amount of data that must be copied to update the recovery point for the volume group in this location.
location3_sync_required	Indicates the total amount of data that must be copied to update the recovery point for the volume group in this location.
location1_sync_remaining	Indicates the remaining amount of data that must be copied to update the recovery point for the volume group in this location. This value can be subtracted from the `location1_sync_required` value to obtain the amount of data that has already been copied.
location2_sync_remaining	Indicates the remaining amount of data that must be copied to update the recovery point for the volume group in this location. This value can be subtracted from the `location2_sync_required` value to obtain the amount of data that has already been copied.

*Table 100. **lsvolumegroupreplication** output (continued)*	
Attribute	**Value**
location3_sync_remaining	Indicates the remaining amount of data that must be copied to update the recovery point for the volume group in this location. This value can be subtracted from the `location3_sync_required` value to obtain the amount of data that has already been copied.
link1_status	Indicates the status of replication between location 1 and location 2 in the replication policy. The following can be displayed: • running • synchronizing • synchronized • suspended • independent • disconnected • blank **Note:** • This field is populated only if the volume group is associated with a replication policy that specifies systems in specific locations. • This field is deprecated and may be removed in a future release. – For DR replication, use **dr_link_status** – For HA replication, use **upper_partition_ha_link_status**
upper_partition_ha_link_status	Indicates the status of HA replication between the systems in the upper partition of the topology diagram. The following can be displayed: • synchronizing • synchronized • suspended • disconnected • unknown • blank

*Table 100. **lsvolumegroupreplication** output (continued)*	
Attribute	**Value**
dr_link_status	Indicates the status of DR replication for the volume group. The following can be displayed: • running • running_missing_connectivity • suspended_missing_connectivity • suspended • independent • disconnected • unknown • blank For a 2-site-async-dr topology, this is the status of replication between location 1 and location 2 in the replication policy. The following can be displayed: • running • suspended • disconnected • unknown For an async-dr topology, this is the status of replication between the systems in the upper and lower partitions of the topology diagram. The following can be displayed: • running_missing_connectivity • suspended_missing_connectivity • suspended • independent • unknown
location1_previous_replication_mode	Indicates the previous replication mode of the volume group on the system in location 1 of the replication policy. The following values can be displayed: • production • recovery • blank **Note:** • This field is populated only if the volume group is associated with a replication policy that specifies systems in specific locations. • This field is deprecated and may be removed in a future release. • Use **upper_partition_previous_replication_mode** or **lower_partition_previous_replication_mode** instead.

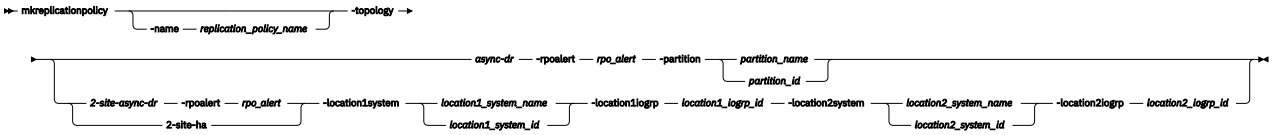
*Table 100. **lsvolumegroupreplication** output (continued)*	
Attribute	**Value**
location2_previous_replication_mode	Indicates the previous replication mode of the volume group on the system in location 2 of the replication policy. • production • recovery • blank **Note:** • This field is populated only if the volume group is associated with a replication policy that specifies systems in specific locations. • This field is deprecated and may be removed in a future release. • Use **upper_partition_previous_replication_mode** or **lower_partition_previous_replication_mode** instead.
upper_partition_previous_replication_mode	Indicates the previous replication mode of the volume group on the systems in the upper partition in the topology diagram. **Note:** • This field is only displayed when the current replication mode of the volume group is **independent**. • This field applies only to async-dr replication between partitions.
lower_partition_previous_replication_mode	Indicates the previous replication mode of the volume group on the systems in the lower partition in the topology diagram. **Note:** • This field is only displayed when the current replication mode of the volume group is **independent**. • This field applies only to async-dr replication between partitions.
location3_status	Indicates the health status of the system in location 3 of the topology diagram. The following can be displayed: • healthy • unhealthy • disconnected • suspended • blank
location1_last_write_time	Indicates the time that volumes in the volume group at this location last received a write I/O from a host. It is displayed as blank or in YYMMDDHHMMSS formation UTC. **Note:** This field only displays when the current replication mode of the volume group is **independent** and the system in this location is accessible. **Note:** If a location is configured for both high availability and disaster recovery, the value is only reported on the location with the active management system.

*Table 100. **lsvolumegroupreplication** output (continued)*	
Attribute	**Value**
location2_last_write_time	Indicates the time that volumes in the volume group at this location last received a write I/O from a host. It is displayed as blank or in YYMMDDHHMMSS format in UTC. **Note:** This field only displays when the current replication mode of the volume group is **independent** and the system in this location is accessible. **Note:** If a location is configured for both high availability and disaster recovery, the value is only reported on the location with the active management system.
location3_last_write_time	Indicates the time that volumes in the volume group at this location last received a write I/O from a host. It is displayed as blank or in YYMMDDHHMMSS format in UTC. **Note:** This field is only populated when the replication mode of the volume group is **independent** and the system in this location is accessible.
partition_id	Indicates the ID of the partition associated with the volume group.
partition_name	Indicates the name of the partition associated with the volume group.
checkpoint_achieved	Indicates whether the recovery point on the recovery copy is no earlier than the time the last checkpoint was started for this volume group by using the **chvolumegroupreplication -checkpoint** command on this system. This field is only populated if a replication policy with a 2-site-async-dr topology is assigned to the volume group and the volume group is a production or recovery copy. Otherwise, the field is blank.
recovery_test_active	Indicates whether a disaster recovery test is currently active on the recovery copy of this volume group. The following values can be displayed: • yes • no • blank
ha_replication_policy_id	Indicates the ID of the HA replication policy assigned to the storage partition containing this volume group.
ha_replication_policy_name	Indicates the name of the HA replication policy assigned to the storage partition containing this volume group.

mkreplicationpolicy

Use the **mkreplicationpolicy** command to create a replication policy.

Syntax



A new replication policy can only be created when all the systems that are specified by the policy are online and connected.

Parameters

-name *replication_policy_name*
(Optional) Specifies the name of the policy.

Note: If a replication policy with the same name exists on the remote system, the command fails. Specify a unique name and rerun the command.

-topology *2-site-async-dr | 2-site-ha | async-dr*
(Mandatory) Specifies the topology of the policy.

`2-site-async-dr:` 2-site asynchronous disaster recovery

- This type of policy is used with volume groups that are not in a storage partition.
- The **-location** parameters are used to specify the systems and the I/O groups in those systems that are used for replication.
- The **-rpoalert** parameter must be specified for this type of policy.
- This type of policy can be used for 2-site replication only.

`async-dr:` asynchronous disaster recovery

- This type of policy is used with volume groups that are in a storage partition.
- The storage partitions are linked to define the systems that are used for disaster recovery replication. The **-location** parameters are not used.
- The partition parameter must be used to specify the partition in which the replication policy is created.
- The **-rpoalert** parameter must be specified for this type of policy.
- This type of policy can be used for 2-site disaster recovery and 3-site (high availability with disaster recovery) configurations.

`2-site-ha:` *High availability*

- This type of policy is used with storage partitions.
- The **-location** parameters are used to specify the systems and the I/O groups in those systems that are used for replication.
- This type of policy can be used for 2-site high availability and 3-site (high availability with disaster recovery) configurations.

-rpoalert *rpo_alert*
Specifies the recovery point objective (RPO) alert threshold in seconds. An alert will be logged if the recovery point exceeds this value.

The minimum value is 60 seconds (one minute) and the maximum value is 86400 seconds (1 day). The value must be a multiple of 60 seconds.

-partition *partition_name | partition_id*
Specifies the partition in which this replication policy will be created. This parameter can only be specified for an async-dr topology.

-location1system *location1_system_name | location1_system_id*
Specifies the name or ID of the system in location 1 of the topology.

-location1iogrp *location1_iogrp_id*
Specifies the ID of the I/O group of the system in location 1 of the topology.

Note: The I/O group parameters accept only numeric ID.

-location2system *location2_system_name | location2_system_id*
Specifies the name or ID of the system in location 2 of the topology.

-location2iogrp *location2_iogrp_id*
Specifies the ID of the I/O group of the system in location 2 of the topology.

Note: The I/O group parameters accept only numeric ID.

An invocation example

The following example creates a replication policy for asynchronous replication:

```
mkreplicationpolicy -name bath_bristol_dr -topology 2-site-async-dr
                -location1system bath -location1iogrp 0 -location2system bristol
-location2iogrp 0
                -rpoalert 60
```

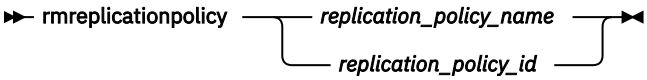
The resulting output:

```
Replication policy, id [2], successfully created
```

rmreplicationpolicy

Use the **rmreplicationpolicy** command to delete a replication policy.

Syntax



Parameters

replication_policy_name | replication_policy_id
(Mandatory) Specifies the name or ID of the replication policy to be removed.

Note:

- A replication policy can only be removed when it is not associated with any volume group.
- If the remote system is accessible, the command removes the policy from the remote system.
- If the remote system is not connected, the replication policy is removed from the remote system when the systems are reconnected.
- A replication policy cannot be removed if it is assigned to a partition.

An invocation example to remove a replication policy

```
rmreplicationpolicy replicationpolicy5
```

The resulting output:

```
No feedback.
```

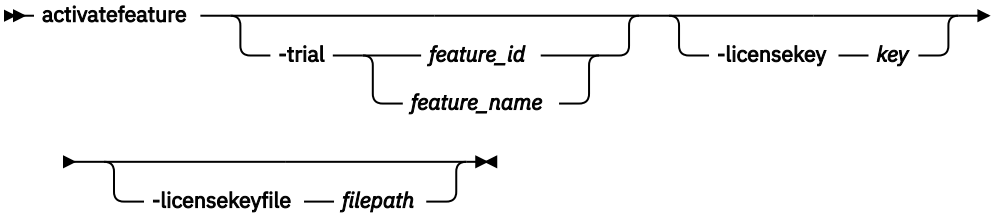
Chapter 17. Licensing and featurization commands

Use the licensing and featurization commands to work with licensed system functions.

activatefeature

Use the **activatefeature** command to activate a feature (using a license key or keyfile) or feature trial period.

Syntax



Parameters

-trial ***feature_id | feature_name***
(Optional) Activates the trial period for the feature of the specified ID that uses an unsigned 16-bit integer:

- Valid integer values are `0`, `1`, and `3`.
- Valid names are `turbo_performance`, `easy_tier`, and `remote_mirroring`.

-licensekey ***key***
(Optional) Provides the license key to activate a feature that contains 16 hexadecimal characters organized in four groups of four numbers with each group separated by a hyphen (such as `0123-4567-89AB-CDEF`).

-licensekeyfile ***filepath***
(Optional) Provides the full path-to-file containing all required license information by using an alphanumeric string that contains 1 - 256 characters.

Description

All parameters are mutually exclusive.

A license key file can contain one or more license keys. If you specify a key file, every key in the file is applied to the system. The license key is checked against the node or control enclosure serial number, machine type, and model. If no valid keys exist in the file, the command cannot complete successfully on the system. If you cannot apply a key successfully to the system, the command adds any remaining keys.

You must have one key for each node or control enclosure. Specify `activatefeature -licensekeyfile` with an `.xml` file that contains all node or control enclosure keys. Or, specify `activatefeature -licensekey` one time per node or control enclosure.

If a feature is already activated and you activate a feature again by using a key, the command completes successfully.

Remember:

- You cannot complete a trial when a feature is activated.
- You can activate a feature while a trial is in progress.

© Copyright IBM Corp. 2003, 2025

An invocation example

```
activatefeature -trial 1
```

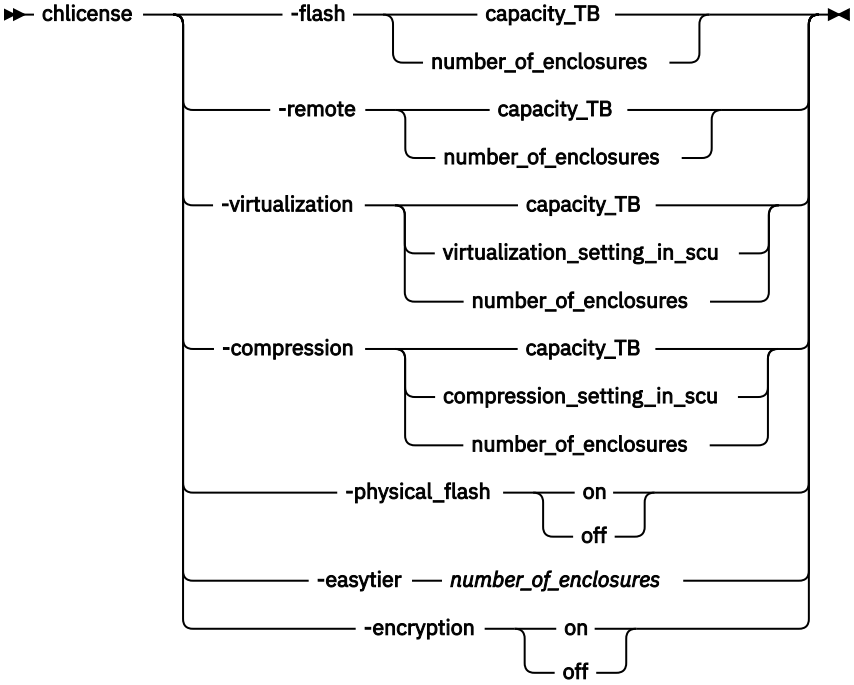
The resulting output:

```
Activation of a trial is a one time operation.  Are you sure you wish to continue? Yes
```

chlicense

Use the **chlicense** command to change license settings for system functions.

Syntax



Parameters

-flash ***capacity_TB | number_of_enclosures***
(Required) Changes system licensing for the FlashCopy function. Depending on your system, specify a capacity in terabytes (TB) or specify the total number of internal and external enclosures for the FlashCopy function.

-remote ***capacity_TB | number_of_enclosures***
(Required) Changes system licensing for policy-based replication and policy-based high availability functions. Depending on your system, specify a capacity in terabytes (TB) or specify the total number of internal and external enclosures that you have licensed on your system. You must have an enclosure license for all enclosures.

-virtualization ***capacity_TB | virtualization_setting_in_scu | number_of_enclosures***
(Required) Changes system licensing for the Virtualization function. Depending on your system, specify a capacity in terabytes (TB) or specify the total number of storage capacity units (SCUs) that you are licensed to virtualize across tiers of storage on your system or specify the number of enclosures of external storage that you are authorized to use.

-compression ***capacity_TB | virtualization_setting_in_scu | number_of_enclosures***
(Required) Changes system licensing for the compression function. Depending on your system, specify a capacity value in terabytes (TB) or specify the total number of storage capacity units (SCUs) that you are licensed to virtualize across tiers of storage on your system or specify the total number of internal

and external enclosures that you have licensed on your system. Only models that have installed extra memory module (16-GB DIMM) support compression.

-physical_flash *on | off*
(Required) For physical disk licensing, this parameter enables or disables the FlashCopy function. The default value is **off**.

-easytier *number_of_enclosures*
(Required) Specifies the number of enclosures on which you can run Easy Tier.

-encryption *on | off*
(Required) Specifies whether the encryption license function is enabled or disabled. The values are on or off.

Description

The **chlicense** command changes license settings for the system. Any change that is made is logged as an event in the license setting log.

Any error that is placed in the license settings log results in a generic error being placed in the system error log. The command-line tool return code also notifies you that you are using an unlicensed function.

deactivatefeature

Use the **deactivatefeature** command to deactivate a feature (such as encryption) or suspend a feature trial period.

Syntax



Parameters

feature_id
(Required) Deactivates the feature (or feature trial). This ID is the unique ID as displayed when you use the **lsfeature** command, and is an incremental number (in the range 0 - 320).

Description

An invocation example

```
deactivatefeature 1
```

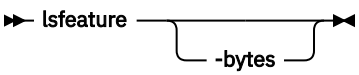
The following output is displayed:

```
You are removing the ability to use a feature of this system. Are you sure you wish to
continue?  Y
```

lsfeature

Use the **lsfeature** command to list the features that are available for the current system code release. You can also list trial or entitlement information and license keys.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-bytes
(Optional) Displays all capacities as bytes.

Description

This command lists the features that are available for the current system code release. You can also list trial or entitlement information and license keys.

This table provides the attribute values that can be displayed as output view data.

Table 101. ***lsfeature*** *outputs*

Attribute	Possible Values
id	Indicates a number that refers to each feature of the system.
name	Indicates the feature name with a 16-character alphanumeric string: • easy_tier • remote_mirroring • flashcopy_upgrade • turbo_performance • encryption
state	Indicates the current state of the feature: • active • inactive • trial_available • trial_active • trial_expired
license_key	Indicates the key that is used in feature activation with a string that consists of 16 hexadecimal characters that are organized in four groups of four numbers with each group that is separated by a hyphen (such as 0123-4567-89AB-CDEF).
trial_expiration_date	Indicates the trial expiration date as long as the state is trial_available or trial_active. This value is displayed in the format YYYYMMDD.

Note: A license key's association with an enclosure can be determined if:

- The enclosure that is associated with the key contains at least one node that is added to the cluster and the node is either online of offline at the time the view is queried.
- The enclosure that is associated with the key contains at least one node that is a candidate for clustered system membership. The node must be online at the time the view is being queried.

lslicense

Use the **lslicense** command to display current license settings for system features.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

Description

The **lslicense** command displays license settings for system features, including virtualization settings.

Use the **chlicense** command to change the feature license settings. Because the feature license settings are entered when the system is first created, you must update the settings if you change your license.

This table provides the possible values that are applicable to the attributes that are displayed as data in the output views.

Table 102. ***lslicense*** *output*

Attribute	Possible Values
used_flash	Indicates the amount of capacity used with FlashCopy.
used_remote	Indicates the amount of capacity used with replication features such as policy-based replication.
used_virtualization	Indicates the amount of capacity used for virtualization.
license_flash	Indicates the amount of capacity licensed for FlashCopy features.
license_remote	Indicates the amount of capacity licensed for replication features.
license_virtualization	Indicates the amount of capacity licensed for virtualization settings.
license_physical_disks	Indicates the amount of physical disk space available for the license.
license_physical_flash	Indicates whether the license's physical flash is on or off.
license_easy_tier	Indicates which enclosures Easy Tier can be run on.
license_virtualization	Indicates the external virtualization settings. The value must be an integer that indicates the number of enclosures.

*Table 102. **lslicense** output (continued)*	
Attribute	**Possible Values**
used_encryption	Indicates how much encryption is being used. The values are on if it is enabled and `off` if it is not enabled.
license_encryption	Indicates whether encryption is allowed or not. The values are on if it is allowed and `off` if it is not allowed.
scu_ratio_scm	Indicates the storage capacity unit (SCU) ratio for SCM tier storage. The value must be a number with two decimal places.
used_virtualization_tier_scm	Indicates the amount of capacity used for virtualization on **`tier_scm`** storage. The value must be zero on platforms where virtualization is not licensed as per SCU.
used_virtualization_tier_flash	Indicates the amount of capacity used for virtualization on **`tier0_flash`** and **`tier1_flash`** storage. The value must be zero on platforms where virtualization is not licensed as per SCU.
used_virtualization_tier_enterprise	Indicates the amount of capacity used for virtualization on **`tier_enterprise`** storage. The value must be zero on platforms where virtualization is not licensed as per SCU.
used_virtualization_tier_nearline	Indicates the amount of capacity used for virtualization on **`tier_nearline`** storage. The value must be zero on platforms where virtualization is not licensed as per SCU.

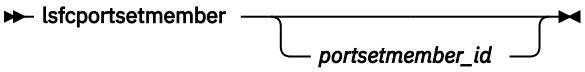
Chapter 18. Fibre Channel portset

Use the following set of commands to associate a Fibre Channel I/O port ID with portset, list the mapping between Fibre Channel I/O port ID and portset objects and to remove Fibre Channel I/O port ID association from the portset.

lsfcportsetmember

Use the **lsfcportsetmember** command to list the mapping between Fibre Channel I/O port ID and portset objects.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

portsetmember_id
(Optional) Specifies the ID of the portset membership. The ID value is in decimal of 0 - 1727.

Description

This table provides the attribute values that can be displayed as output view data.

Table 103. lsfcportsetmember output

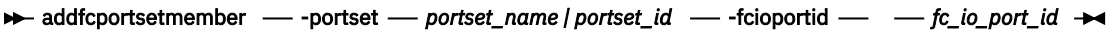
Attribute	Description
id	Indicates the ID of the map. The value can be a decimal number in the range 0 - 1727.
fc_io_port_id	Indicates the Fibre Channel I/O port ID of the port. The value can be a decimal number 1 to maximum Fibre Channel I/O ports.
portset_id	Indicates the portset id to which Fibre Channel I/O port ID is added. The value can be a decimal number 0 to 71.
portset_name	Indicates the portset name to which Fibre Channel I/O port ID is added. The value can be alphanumeric string.
owner_id	Indicates the ID of the ownership group.
owner_name	Indicates the name of the ownership group.

© Copyright IBM Corp. 2003, 2025

addfcportsetmember

Use the **addfcportsetmember** command to associate a Fibre Channel IO port ID with portset. A Fibre Channel IO port ID can be associated to multiple portsets at the same time and this association applies to all nodes in a cluster.

Syntax



Parameters

-portset *portset_name | portset_id*
(Required) Specifies the ID or name of the portset. The specified value must match one of the portsets listed by **lsportset**. The portset must be of host type with port type fc.

-fcioportid *fc_io_port_id*
(Required) Specifies the Fibre Channel IO port ID to be associated with portset. The **fc_io_port_id** is as listed in output of **lsportfc** command. The value can be a decimal number in the range 1 to maximum Fibre Channel IO port count.

Note: When the **addfcportsetmember** command runs, it changes I/O paths for the hosts in the same portset. Admin can run device discovery from the FC hosts added in the same portset to allow host to discover the devices from the valid I/O paths.

An invocation example (Add FC portid 1 on each node in cluster to portset5)

```
addfcportsetmember -portset portset5 -fcioportid 1
```

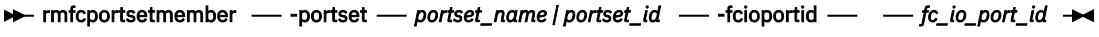
The resulting output:

```
No feedback
```

rmfcportsetmember

Use the **rmfcportsetmember** command to remove Fibre Channel IO port ID association from the portset.

Syntax



Note: When the **rmfcportsetmember** command runs, it changes I/O paths for the hosts in the same portset. Admin can run device discovery from the FC hosts added in the same portset to allow host to discover the devices from the valid I/O paths.

Parameters

-portset *portset_name | portset_id*
(Required) Specifies the ID or name of the portset. The specified value must match one of the portsets listed by **lsportset**. The Fibre Channel IO port must be associated with specified portset.

-fcioportid *fc_io_port_id*
(Required) Specifies the Fibre Channel IO port ID from output of the **lsfcportsetmember** command. The value can be a decimal number in the range 1 to maximum Fibre Channel IO port count.

Note: When the **addfcportsetmember** command runs, it changes I/O paths for the hosts in the same portset. Admin can run device discovery from the FC hosts added in the same portset to allow host to discover the devices from the valid I/O paths.

An invocation example (Removing association between portset name portset1 and port id 1 for every node in cluster)

```
rmfcportsetmember -portset portset1 -fcioportid 1
```

The resulting output:

```
No feedback.
```

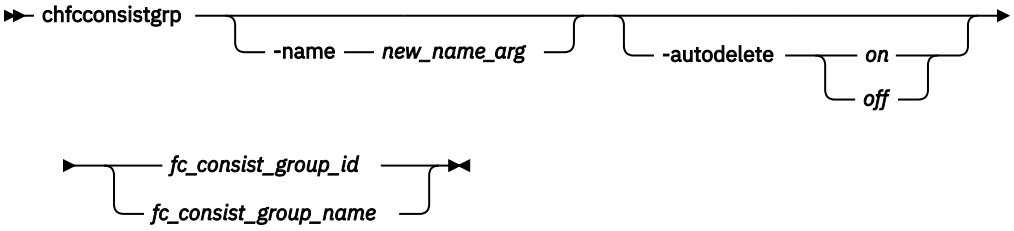
Chapter 19. FlashCopy commands

Use the FlashCopy commands to work with FlashCopy system methods and functions.

chfcconsistgrp

Use the **chfcconsistgrp** command to change the name of a consistency group or mark the group for auto-deletion.

Syntax



Parameters

-name *new_name_arg*
(Optional) Specifies the new name to assign to the consistency group.

-autodelete *on | off*
(Optional) Deletes the consistency group when the last mapping that it contains is deleted or removed from the consistency group.

fc_consist_group_id | fc_consist_group_name
(Required) Specifies the ID or existing name of the consistency group that you want to modify.

Description

The **chfcconsistgrp** command changes the name of a consistency group, marks the group for auto-deletion, or both.

Note: Maps that are `rc_controlled` are not shown in the view when this command is specified.

An invocation example

```
chfcconsistgrp -name testgrp1 fcconsistgrp1
```

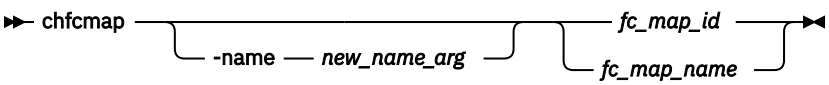
The resulting output:

```
No feedback
```

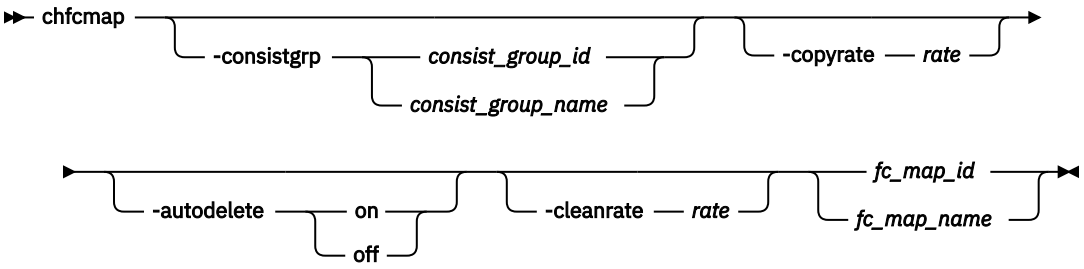
chfcmap

Use the **chfcmap** command to modify attributes of an existing mapping.

Syntax



© Copyright IBM Corp. 2003, 2025



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.

-consistgrp *consist_group_id | consist_group_name*
(Optional) Specifies the consistency group for which you want to modify the mapping. To modify a map such that it is stand-alone (not in a consistency group), specify `-consistgrp 0`.

Restriction: 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 - 150. The default value is 50. A value of 0 indicates no background copy process. For the supported **-copyrate** values and their corresponding rates, see the table in the description.

-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 delete a mapping after the background copy is completed. If the background copy is completed, the mapping is deleted immediately.

-cleanrate *rate*
(Optional) Sets the cleaning rate for the mapping. The *rate* value can be 0 - 150. The default value is 50.

fc_map_id | fc_map_name
(Required) Specifies the ID or name of the mapping to modify. Enter the ID or name last on the command line.

Description

The **chfcmap** command modifies attributes of an existing mapping.

Important: You must enter the *fc_map_id | fc_map_name* last on the command line.

If you have created several FlashCopy mappings for a group of volumes that contain elements of data for the same application, you can assign these mappings to a single FlashCopy consistency group. You can then issue a single prepare command and a single start command for the whole group, for example, so that all of the files for a particular database are copied at the same time.

The **-copyrate** parameter specifies the copy rate. If 0 is specified, background copy operations are disabled. The **-cleanrate** parameter specifies the rate for cleaning the target volume. The cleaning process is active only 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 can be completed.

The following table describes 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 104. Relationship between the rate, data rate, and grains per second values			
User-specified rate attribute value	**Data copied/sec**	**256 KB grains/sec**	**64 KB grains/sec**
1 - 10	128 KB	0.5	2
11 - 20	256 KB	1	4
21 - 30	512 KB	2	8
31 - 40	1 MB	4	16
41 - 50	2 MB	8	32
51 - 60	4 MB	16	64
61 - 70	8 MB	32	128
71 - 80	16 MB	64	256
81 - 90	32 MB	128	512
91 - 100	64 MB	256	1024
101 - 110	128 MB	512	2048
111 - 120	256 MB	1024	4096
121 - 130	512 MB	2048	8192
131 - 140	1 GB	4096	16384
141 - 150	2 GB	8192	32768

Restriction: Maps that are `rc_controlled` are not shown in the view after this command is processed.

Note: For FlashCopy mappings that are used for policy-based replication, the **chfcmap** command can fail. These mappings are system-owned and are created automatically as part of the policy-based replication. However, you can change the **-copyrate** and **-cleanrate** parameters of a reverse mapping and also change the **-cleanrate** parameter on a forward mapping.

An invocation example

```
chfcmap -name testmap 1
```

The resulting output:

```
No feedback
```

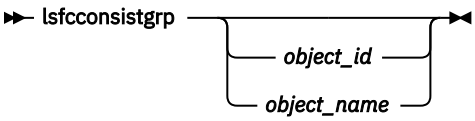
lsfcconsistgrp

Use the **lsfcconsistgrp** command to display a concise list or a detailed view of FlashCopy consistency groups that are visible to the system. This information is useful for tracking FlashCopy consistency groups.

The list report style can be used to obtain two styles of report:

- A list that contains concise information about all of the FlashCopy consistency groups on a system. (Each entry in the list corresponds to a single FlashCopy consistency group.)
- The detailed information about a single FlashCopy consistency group.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

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 that match the filtering requirements that are specified by the **-filtervalue** parameter are displayed.

Description

This command returns a concise list or a detailed view of FlashCopy consistency groups that are visible to the system.

The following list provides values of the `status` attribute that are displayed as data in the output views:

status
Indicates the status. The values are:

- `idle_or_copied`
- `preparing`
- `prepared`
- `copying`
- `stopped`
- `suspended`
- `stopping`
- Empty

id
Indicates the mapping ID.

name
Indicates the mapping name.

start_time
Indicates the time that the group was started in YYMMDDHHMMSS format (or blank).

autodelete
Indicates whether auto deletion is on or `off`.

FC_mapping_id
Indicates the FlashCopy mapping ID.

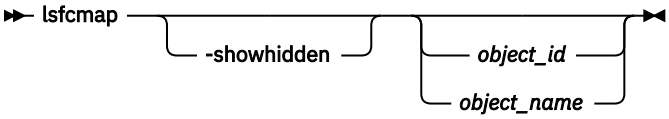
FC_mapping_name
Indicates the FlashCopy mapping name.

size_mismatch
Indicate one or more of the FlashCopy mappings in the group that has size mismatching. The value is set to yes or no.

lsfcmap

Use the **lsfcmap** command to generate a list that contains concise information about all of the FlashCopy mappings that are visible to the system, or detailed information for a single FlashCopy mapping.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-showhidden
(Optional) Volume group snapshot mappings (snapshot, clone, thinclone and restore) and policy-based replication mappings are system-owned FlashCopy mappings. To display system-owned mappings in the concise view, use the **-showhidden** parameter.

object_id | object_name
(Optional) Specifies the name or ID of an object.

Description

This command returns a concise list or a detailed view of FlashCopy mappings that are visible to the system.

The following attributes are displayed:

id
Displays the mapping ID.

name
Displays the mapping name.

source_vdisk_id
Displays the source volume ID.

source_vdisk_name
Displays the source volume name.

target_vdisk_id
Displays the target volume ID.

target_vdisk_name
Displays the target volume name.

group_id
Displays the group ID.

group_name
Displays the group name.

status
Displays the status:

- idle_or_copied
- preparing
- prepared
- copying
- stopped
- suspended
- stopping

progress
Displays the progress. The percentage of the target volume that has its own copy of the data.

copy_rate
Displays the copy rate 0-150, where 50 is 2MB/s, 60 is 4MB/s etc.

start_time
Displays the time that the copy was last started. It is in the format YYMMDDHHMMSS. If a copy is not started, a blank line is displayed.

dependent_mappings
Displays any dependent mappings.

autodelete
Specifies if autodelete is on or off. If on, then the map will automatically delete its self once the progress is 100%, the mappings state is idle_copied and the target volume has no dependent mappings.

clean_progress
Indicates the clean progress. The percentage of the target volume that is not required by a dependent target volume.

clean_rate
Indicates the clean rate, 0-150, where 50 is 2MB/s, 60 is 4MB/s etc.

incremental
Indicates whether incremental is on or off.

difference
Indicates the difference for an incremental mapping by showing the percentage of the source volume that has changed since the mapping was last initiated.

grain_size
Indicates the grain size used by the FlashCopy mappings metadata, can be 64KiB or 256KiB.

IO_group_id
Displays the I/O group ID.

IO_group_name
Displays the I/O group name.

partner_FC_id
Displays the partner FlashCopy ID. The partner is a mapping whose source volume matches this maps target volume and target volume matches this maps source volume.

partner_FC_name
Displays the partner FlashCopy name. The partner is a mapping whose source volume matches this maps target volume and target volume matches this maps source volume.

restoring
Displays the restoring status. The values are yes or no.

rc_controlled (Deprecated)

keep_target
Displays the target and source volume availability. The values are:

- yes indicates that source volume will be offline if the target volume is offline and the mappings is copying
- no indicates that if target volume is offline and the mapping is copying, then mapping will be force stopped if copy-on-write operations are prevented by the offline target volume.

restore_progress

Displays the percentage of the source volume that is restored from the target.

fc_controlled

Indicates that the FlashCopy mapping is either a cloud tiering mapping or a vVols restore FlashCopy mapping that cannot be modified.

owner_id

Indicates the ID of the ownership group.

owner_name

Indicates the name of the ownership group.

size_mismatch

Indicates that the source and target VDisks have different virtual capacity. The value is set to yes or no.

size_capacity

Indicates that size of the source and target at the time the FlashCopy mapping was last started.

is_snapshot

Indicates the FlashCopy mapping is used to maintain a snapshot.

snapshot_id

Indicates the ID of the snapshot associated with the FlashCopy mapping.

type

Indicates the type of the FlashCopy mapping.

- **generic**
  Indicates the standard mapping, includes policy-based replication change volume mappings.

  snapshot
  If is_snapshot is yes then volume group snapshot mapping, otherwise a vVols snapshot or thinclone.

  thinclone

  To display a thinclone mapping, requires the **-showhidden** to view.

  clone
  If source volume is the target volume group snapshot mappings then its a VGS clone mapping otherwise it is a vVols clone or restore mapping.

  restore
  Indicates the volume group snapshot restore mapping.

  cloud_backup
  Indicates the cloud tiering backup mapping.

cloud_restore
Indicates the cloud tiering restore mapping.

accesspoint
Indicates the policy-based replication 3-site accesspoint forward mapping.

accesspoint_restore
Indicates the policy-based replication 3-site accesspoint reverse mapping.

reverse_change
Indicates the policy-based replication change volume reverse mapping.

forward_change
Indicates the policy-based replication change volume forward mapping.

lsfcmapcandidate

Use the **lsfcmapcandidate** command to list all of the volumes that are associated with fewer than 256 FlashCopy mappings.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

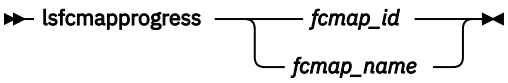
Description

This command returns a list of volumes that are associated with fewer than 256 FlashCopy mappings.

lsfcmapprogress

Use the **lsfcmapprogress** command to display the progress of the background copy of a FlashCopy mapping. This information is displayed as a percentage-completed value.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

fcmap_id | fcmap_name
(Required) Specifies that you want the report to display the progress of the background copy for the designated FlashCopy mapping.

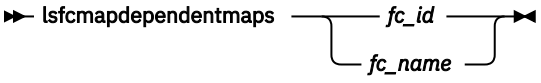
Description

This command reports a percentage for the progress of the background copy being copied on the specified FlashCopy mapping.

lsfcmapdependentmaps

Use the **lsfcmapdependentmaps** command to display the FlashCopy mappings that depend on the user specified mapping.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

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 by using the **-force** parameter.

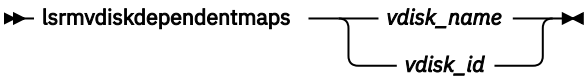
Use the dependent_mapping_count field in the FlashCopy map detailed view (displayed when you process the **lsfcmap** command) as an indicator of dependent mappings in progress. A count of zero means no dependent copies.

Note: Time lapses from processing the **lsfcmap** command to processing the **lsfcmapdependentmaps** command can cause differences between the number of dependent mappings that are being processed and the number that the **lsfcmap** command shows.

lsrmvdiskdependentmaps

Use the **lsrmvdiskdependentmaps** command to display all FlashCopy mappings that must be stopped for the specified volume to be deleted.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

vdisk_name | vdisk_id
(Required) Specifies the name or ID of the volume for which the FlashCopy mappings are displayed.

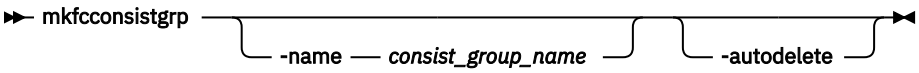
Description

This command returns a list of the FlashCopy mappings that must be stopped before the specified volume can be deleted. Any mappings that are returned in the list for the volume are automatically stopped when the volume is deleted with the **force** option.

mkfcconsistgrp

Use the **mkfcconsistgrp** command to create a new FlashCopy consistency group and identification name.

Syntax



Parameters

-name ***consist_group_name***
(Optional) Specifies a name for the consistency group. If you do not specify a consistency group name, a name is automatically assigned to the consistency group. For example, if the next available consistency group ID is `id=2`, the consistency group name is `fccstgrp2`.

Note: Consistency group names must be an alphanumeric string of up to 15 characters.

-autodelete
(Optional) Deletes the consistency group when the last mapping that it contains is deleted or removed from the consistency group.

Description

This command creates a new consistency group and identification name. The ID of the new group is displayed when the command process completes.

If you have created several FlashCopy mappings for a group of volumes that contain elements of data for the same application, you might find it convenient to assign these mappings to a single FlashCopy consistency group. You can then issue a single prepare command and a single start command for the whole group, for example, so that all of the files for a particular database are copied at the same time.

An invocation example

```
mkfcconsistgrp
```

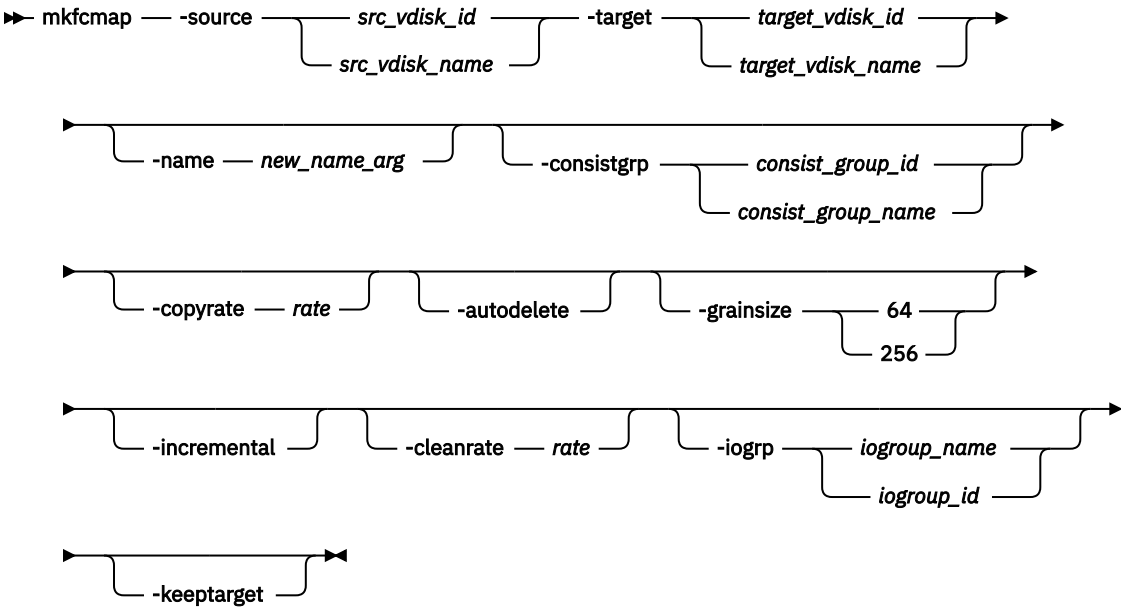
The resulting output:

```
FlashCopy Consistency Group, id [1], successfully created
```

mkfcmap

Use the **mkfcmap** command to create a new FlashCopy mapping, which maps a source volume to a target volume for subsequent copying.

Syntax



Parameters

-source *src_vdisk_id | src_vdisk_name*
(Required) Specifies the ID or name of the source volume.

-target *target_vdisk_id | target_vdisk_name*
(Required) Specifies the ID or name of the target volume.

-name *new_name_arg*
(Optional) Specifies the name to assign to the new mapping.

-consistgrp *consist_group_id | consist_group_name*
(Optional) Specifies the consistency group to add the new mapping to. If you do not specify a consistency group, the mapping is treated as a stand-alone mapping.

-copyrate *rate*
(Optional) Specifies the copy rate. The *rate* value can be 0 - 150. The default value is 50. A value of 0 indicates no background copy process. For the supported **-copyrate** values and their corresponding rates, see the table in the command description.

-autodelete
(Optional) Specifies that a mapping is to 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. After set, this value cannot be changed.

Remember: If either the source or target disk contains compressed copies, the default value is 64 (unless source or target disk is part of a mapping with grain size 256 KB).

-incremental
(Optional) Marks the FlashCopy mapping as an incremental copy. The default is nonincremental. After set, this value cannot be changed.

-cleanrate *rate*
(Optional) Sets the cleaning rate for the mapping. The *rate* value can be 0 - 150. The default value is 50.

-iogrp *iogroup_name | iogroup_id*
(Optional) Specifies the I/O group for the FlashCopy bitmap. After set, this value cannot be changed. The default I/O group is either the source volume, if a single target map, or the I/O group of the other FlashCopy mapping to which either the source or target volumes belong.

Note: If not enough bitmap space is available to complete this command, more space is automatically allocated in the bitmap memory (unless the maximum bitmap memory is reached).

-keeptarget
(Optional) Specifies that the target volume and source volume availability should be kept the same. If the target becomes unavailable, the source is also made unavailable (instead of stopping the FlashCopy mapping).

Description

This command creates a new FlashCopy mapping. This mapping persists until it is manually deleted, or until it is automatically deleted when the background copy completes and the **autodelete** parameter is set to on. The source and target volumes must be specified on the **mkfcmap** command. The **mkfcmap** command fails if the source and target volumes are not identical in size. Issue the **lsvdisk -bytes** command to find the exact size of the source volume for which you want to create a target disk of the same size. The target volume that you specify cannot be a target volume in an existing FlashCopy mapping. A mapping cannot be created if the resulting set of connected mappings exceeds 256 connected mappings.

The mapping can optionally be given a name and assigned to a consistency group, which is a group of mappings that can be started with a single command. These are groups of mappings that can be processed at the same time. It enables multiple volumes to be copied at the same time, which creates a consistent copy of multiple disks. This consistent copy of multiple disks is required by some database products in which the database and log files reside on different disks.

If the specified source and target volumes are the target and source volumes, respectively, of an existing mapping, then the mapping that is 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.

Note: You cannot use this command if a volume is part of a mapping and cloud snapshot is enabled on the volume.

The **copyrate** parameter specifies the copy rate. If `0` is specified, background copy is disabled. The **cleanrate** parameter specifies the rate for cleaning the target volume. The cleaning process is only active if the mapping is in the copying state and the background copy is completed, the mapping is in the copying state and the background copy is disabled, or the mapping is in the stopping state. You can disable cleaning when the mapping is in the copying state by setting the **cleanrate** parameter to `0`. If the **cleanrate** is set to `0`, the cleaning process runs at the default rate of 50 when the mapping is in the stopping state to ensure that the stop operation completes.

This table provides the relationship of the copy rate and cleaning rate values to the attempted number of grains to be split per second. A grain is the unit of data that is represented by a single bit.

Note: The default grain size of 64 KB for compressed volumes applies only to compressed volumes in regular pools if:

- Either the source or target volumes are compressed within a regular pool.
- Either the source or target volumes are themselves sources or targets of compressed volumes in regular pools (i.e. in a cascade where other volumes in the cascade are compressed volumes in a regular pool).

Otherwise, the default grain size is 256 KB for data reduction compressed volumes, provided the volumes are not involved in cascades with compressed volumes in regular pools.

Table 105. Relationship between the rate, data rate, and grains per second values			
User-specified rate attribute value	**Data copied/sec**	**256 KB grains/sec**	**64 KB grains/sec**
1 - 10	128 KB	0.5	2
11 - 20	256 KB	1	4
21 - 30	512 KB	2	8
31 - 40	1 MB	4	16
41 - 50	2 MB	8	32
51 - 60	4 MB	16	64
61 - 70	8 MB	32	128
71 - 80	16 MB	64	256
81 - 90	32 MB	128	512
91 - 100	64 MB	256	1024
101 - 110	128 MB	512	2048
111 - 120	256 MB	1024	4096
121 - 130	512 MB	2048	8192
131 - 140	1 GB	4096	16384
141 - 150	2 GB	8192	32768

Remember: This command cannot be used on a volume that is owned by a file system.

Note: Maps that are **rc_controlled** are not shown in the view when this command is specified.

An invocation example

```
mkfcmap -source 0 -target 2 -name mapone
```

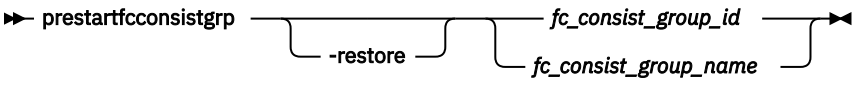
The resulting output:

```
FlashCopy Mapping, id [1], successfully created
```

prestartfcconsistgrp

Use the **prestartfcconsistgrp** command to prepare a consistency group (a group of FlashCopy mappings) so that the consistency group can be started. This command flushes the cache of any data that is destined for the source volume and forces the cache into the write-through mode until the consistency group is started.

Syntax



Parameters

-restore
(Optional) Specifies the restore flag. This forces the consistency group to be prepared even if the target volume of one of the mappings in the consistency group is being used as a source volume of another active mapping. An active mapping is in the copying, suspended, or stopping state.

fc_consist_group_id | fc_consist_group_name
(Required) Specifies the name or ID of the consistency group that you want to prepare.

Description

This command prepares a consistency group (a group of FlashCopy mappings) to subsequently start. The preparation step ensures that any data that resides in the cache for the source volume is first flushed to disk. This step ensures that the FlashCopy target volume is identical to what has been acknowledged to the host operating system as having been written successfully to the source volume.

You can use the **restore** parameter to force the consistency group to be prepared even if the target volume of one or more mappings in the consistency group is being used as a source volume of another active mapping. In this case the mapping restores as shown in the **lsfcmap** view. If the **restore** parameter is specified when preparing a consistency group where none of the target volumes are the source volume of another active mapping, then the parameter is ignored.

You must issue the **prestartfcconsistgrp** command to prepare the FlashCopy consistency group before the copy process can be started. When you have assigned several mappings to a FlashCopy consistency group, you must issue a single prepare command for the whole group to prepare all of the mappings at once.

The consistency group must be in the idle_or_copied or stopped state before it can be prepared. When you enter the **prestartfcconsistgrp** command, the group enters the preparing state. After the preparation is complete, the consistency group status changes to prepared. At this point, you can start the group.

If FlashCopy mappings are assigned to a consistency group, the preparing and the subsequent starting of the mappings in the group must be performed on the consistency group rather than on an individual FlashCopy mapping that is assigned to the group. Only stand-alone mappings, which are mappings that are not assigned to a consistency group, can be prepared and started on their own. A FlashCopy consistency group must be prepared before it can be started.

An invocation example

```
prestartfcconsistgrp 1
```

The resulting output:

```
No feedback
```

prestartfcmap

Use the **prestartfcmap** command to prepare a FlashCopy mapping so that it can be started. This command flushes the cache of any data that is destined for the source volume and forces the cache into the write-through mode until the mapping is started.

Syntax



Parameters

-restore
(Optional) Specifies the restore flag. This forces the mapping to be prepared even if the target volume is being used as a source volume in another active mapping. An active mapping is in the copying, suspended, or stopping state.

fc_map_id | fc_map_name
(Required) Specifies the name or ID of the mapping to prepare.

Description

This command prepares a single mapping for subsequent starting. The preparation step ensures that any data that resides in the cache for the source volume is first transferred to disk. This step ensures that the copy that is made is consistent with what the operating system expects on the disk.

The **restore** parameter can be used to force the mapping to be prepared even if the target volume is being used as a source volume of another active mapping. In this case, the mapping is restoring as shown in the **lsfcmap** view. If the **restore** parameter is specified when preparing a mapping where the target volume is not the source volume of another active mapping, then the parameter is ignored.

Note: To prepare a FlashCopy mapping that is part of a consistency group, you must use the **prestartfcconsistgrp** command.

The mapping must be in the `idle_or_copied` or `stopped` state before it can be prepared. When the **prestartfcmap** command is processed, the mapping enters the preparing state. After the preparation is complete, it changes to the prepared state. At this point, the mapping is ready to start.

Attention: This command can take a considerable amount of time to complete. For example, while a volume is in the `prepared` state, response times might increase.

An invocation example

```
prestartfcmap 1
```

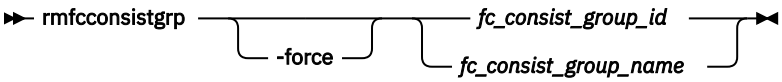
The resulting output:

```
No feedback
```

rmfcconsistgrp

Use the **rmfcconsistgrp** command to delete a FlashCopy consistency group.

Syntax



Parameters

-force
(Optional) Specifies that all of the mappings that are associated with a consistency group that you want to delete are removed from the group and changed to stand-alone mappings. This parameter is only required if the consistency group that you want to delete contains mappings.

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

fc_consist_group_id | fc_consist_group_name
(Required) Specifies the ID or name of the consistency group that you want to delete.

Description

This command deletes the specified FlashCopy consistency group. If there are mappings that are members of the consistency group, the command fails unless you specify the **-force** parameter. When you specify the **-force** parameter, all of the mappings that are associated with the consistency group are removed from the group and changed to stand-alone mappings.

To delete a single mapping in the consistency group, you must use the **rmfcmap** command.

Note: Maps that are `rc_controlled` are not shown in the view when this command is specified.

An invocation example

```
rmfcconsistgrp fcconsistgrp1
```

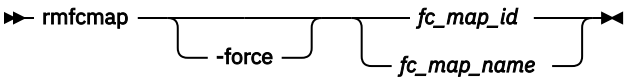
The resulting output:

```
No feedback
```

rmfcmap

Use the **rmfcmap** command to delete an existing mapping.

Syntax



Parameters

-force
(Optional) Specifies that the target volume is brought online. This parameter is required if the FlashCopy mapping is in the stopped state.

fc_map_id | fc_map_name
(Required) Specifies the ID or name of the FlashCopy mapping to delete. Enter the ID or name last on the command line.

Description

The **rmfcmap** command deletes the specified mapping if the mapping is in the `idle_or_copied` or `stopped` state. If it is in the stopped state, the **-force** parameter is required. If the mapping is in any other state, you must stop the mapping before you can delete it.

Deleting a mapping only deletes the logical relationship between the two volumes; it does not affect the volumes themselves. However, if you force the deletion, the target volume (which might contain inconsistent data) is brought back online.

An invocation example

```
rmfcmap testmap
```

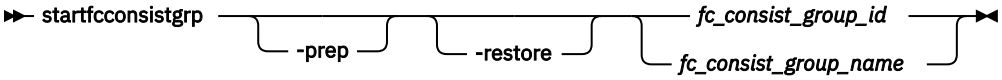
The resulting output:

```
No feedback
```

startfcconsistgrp

Use the **startfcconsistgrp** command to start a FlashCopy consistency group of mappings. This command makes a point-in-time copy of the source volumes at the moment that the command is started.

Syntax



Parameters

-prep
(Optional) Specifies that the designated FlashCopy consistency group be prepared prior to starting the FlashCopy consistency group. A FlashCopy consistency group must be prepared before it can be started. When you use this parameter, the system automatically issues the **prestartfcconsistgrp** command for the group that you specify.

-restore
(Optional) Specifies the restore flag. When combined with the **prep** option, this forces the consistency group to be prepared even if the target volume of one of the mappings in the consistency group is being used as a source volume in another active mapping. An active mapping is in the copying, suspended, or stopping state.

fc_consist_group_id | fc_consist_group_name
(Required) Specifies the ID or name of the consistency group mapping to start.

Description

This command starts a consistency group, which results in a point-in-time copy of the source volumes of all mappings in the consistency group. You can combine the **restore** parameter with the **prep** parameter to force the consistency group to be prepared prior to starting, even if the target volume of one or more mappings in the consistency group is being used as a source volume of another active mapping. In this case, the mapping is restoring as shown in the **lsfcmap** view. If the **restore** parameter is specified when starting a consistency group where none of the target volumes are the source volume of another active mapping, the parameter is ignored.

If a consistency group is started and the target volume of the mapping being started has up to four other incremental FlashCopy mappings using the target, the incremental recording is left on. If there are more than four other incremental FlashCopy mappings using the target volume, the incremental recording for all of these mappings is turned off until they are restarted.

Note: The **startfcconsistgrp** command can take some time to process particularly if you have specified the **prep** parameter. If you use the **prep** parameter, you give additional processing control to the system because the system must prepare the mapping before the mapping is started. If the prepare process takes too long, the system completes the prepare but does not start the consistency group. In this case, error message CMMVC6209E displays. To control the processing times of the **prestartfcconsistgrp** and **startfcconsistgrp** commands independently of each other, do not use the **prep** parameter. Instead, first issue the **prestartfcconsistgrp** command, and then issue the **startfcconsistgrp** command to start the copy.

An invocation example

```
startfcconsistgrp -prep 2
```

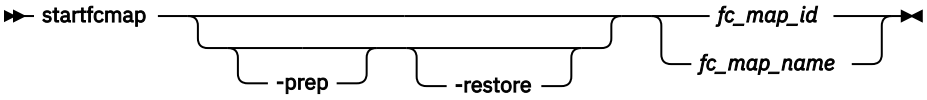
The resulting output:

```
No feedback
```

startfcmap

Use the **startfcmap** command to start a FlashCopy mapping. This command makes a point-in-time copy of the source volume at the moment that the command is started.

Syntax



Parameters

-prep
(Optional) Specifies that the designated mapping be prepared prior to starting the mapping. A mapping must be prepared before it can be started. When you use this parameter, the system automatically issues the **prestartfcmap** command for the group that you specify.

Note: If you have already used the **prestartfcmap** command, you cannot use the **-prep** parameter on the **startfcmap** command; the command fails. However, if the FlashCopy has successfully prepared before, the **startfcmap** command succeeds.

-restore
(Optional) Specifies the restore flag. When combined with the **prep** option, this forces the mapping to be prepared even if the target volume is being used as a source volume in another active mapping. An active mapping is in the copying, suspended, or stopping state.

fc_map_id | fc_map_name
(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 volume. You can combine the **restore** parameter with the **prep** parameter to force the mapping to be prepared prior to starting, even if the target volume is being used as a source volume of another active mapping. In this case, the mapping is restoring as shown in the **lsfcmap** view. If the **restore** parameter is specified when starting a mapping where the target volume is not the source volume of another active mapping, the parameter is ignored and the mapping is not restoring as shown in the **lsfcmap** view.

If a mapping is started and the target volume of the mapping being started has up to four other incremental FlashCopy mappings using the target, the incremental recording is left on. If there are more than four other incremental FlashCopy mappings using the target volume, the incremental recording for all of these mappings is turned off until they are restarted.

Note: The **startfcmap** command can take some time to start, particularly if you use the **prep** parameter. If you use the **prep** parameter, you give additional starting control to the system. The system must prepare the mapping before the mapping is started. To keep control when the mapping starts, you must issue the **prestartfcmap** command before you issue the **startfcmap** command.

An invocation example

```
startfcmap -prep 2
```

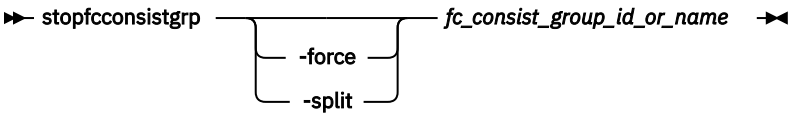
The resulting output:

```
No feedback
```

stopfcconsistgrp

Use the **stopfcconsistgrp** command to stop all processing that is associated with a FlashCopy consistency group that is in one of the following processing states: prepared, copying, stopping,or suspended.

Syntax



Parameters

-force
(Optional) Specifies that all processing that is associated with the mappings of the designated consistency group be stopped immediately.

Note: When you use this parameter, all FlashCopy mappings that depend on the mappings in this group (as listed by the **lsfcmapdependentmaps** command) are also stopped.

If the **-force** parameter is not specified, the command is rejected if the target volume of the FlashCopy consistency group is the primary in a relationship that is mirroring I/O:

- consistent_synchronized
- consistent_copying
- inconsistent_copying

-split
(Optional) Breaks the dependency on the source volumes of any mappings that are also dependent on the target volume. This parameter can only be specified when stopping a consistency group where all maps in the group have progress of 100 as shown by the **lsfcmap** command.

fc_consist_group_id_or_name
(Required) Specifies the name or ID of the consistency group that you want to stop.

Description

This command stops a group of mappings in a consistency group. If the copy process is stopped, the target disks become unusable unless they already contain complete images of the source. Disks that contain complete images of the source have a progress of 100, as indicated in the **lsfcmap** command output. The target volume is reported as offline if it does not contain a complete image. Before you can access this volume, the group of mappings must be prepared and restarted.

If the consistency group is in the **idle_or_copied** state, the **stopfcconsistgrp** command has no effect and the consistency group stays in the **idle_or_copied** state.

The **split** option can be used when all of the maps in the group have progress of 100. It removes the dependency of any other maps on the source volumes. It might be used prior to starting another FlashCopy consistency group whose target disks are the source disks of the mappings being stopped. Once the consistency group has been stopped with the **split** option, the other consistency group could then be started without the **restore** option.

Note: Maps that are **rc_controlled** are not shown in the view when this command is specified.

An invocation example

```
stopfcconsistgrp testmapone
```

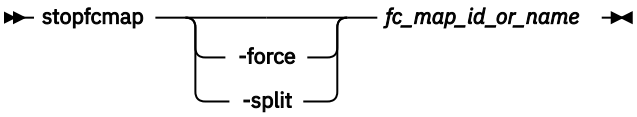
The resulting output

```
No feedback
```

stopfcmap

Use the **stopfcmap** command to stop all processing that is associated with a FlashCopy mapping that is in one of the following processing states: prepared, copying, stopping, or suspended.

Syntax



Parameters

-force
(Optional) Specifies that all processing that is associated with the designated mapping stops immediately.

Note: When you use this parameter, all FlashCopy mappings that depend on this mapping (as listed by the **lsfcmapdependentmaps** command) are also stopped.

If the **-force** parameter is not specified, the command is rejected if the target volume of the FlashCopy mapping is the primary in a relationship that is mirroring I/O:

- consistent_synchronized
- consistent_copying
- inconsistent_copying

-split
(Optional) Breaks the dependency on the source volume of any mappings that are also dependent on the target disk. Specify this parameter only when you are stopping a map with a progress of 100 as shown by the **lsfcmap** command.

fc_map_id_or_name
(Required) Specifies the name or ID of the mapping to stop.

Description

This command stops a single mapping. If the copy process is stopped, the target disk becomes unusable unless it already contained a complete image of the source. A complete image is indicated by a map with a progress of 100, as shown by the **lsfcmap** command. For incomplete mappings, prepare and reprocess the mapping before you use the target disk.

Only stand-alone mappings can be stopped with the **stopfcmap** command. Stop mappings that belong to a consistency group with 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.

The **split** option can be used when the mapping has progress of 100. It removes the dependency of any other mappings on the source volume. It might be used before you start another FlashCopy mapping whose target disk is the source disk of the mapping that you are stopping. After the mapping stops by using the **split** option, you can start the other mapping without the **restore** option.

An invocation example

```
stopfcmap testmapone
```

The resulting output

```
No feedback
```

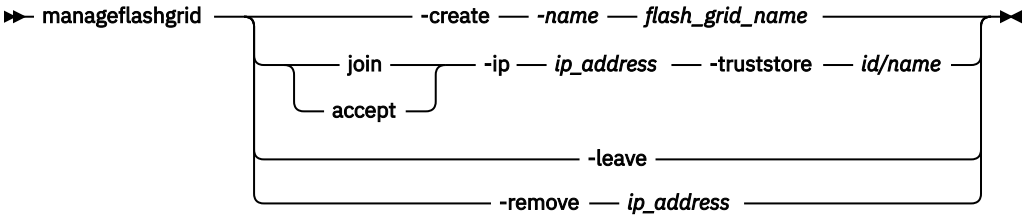
Chapter 20. FlashSystem grid commands

Use the FlashSystem grid commands to manage and display the systems in a FlashSystem grid.

manageflashgrid

Use the **manageflashgrid** command to manage a FlashSystem grid.

Syntax



Parameters

-create *flag*
(Required) Creates a new FlashSystem grid and adds the current system.

-name *string*
(Required) Specifies a name to use for the FlashSystem grid.

-join *flag*
(Optional) Requests the current system to join the specified FlashSystem grid.

-accept *flag*
(Optional) Accepts a request from another system to join the FlashSystem grid where the current system is added.

-ip *IPv4_IPv6_address_or_FQDN*
(Optional) Specifies the IP address (IPv4, IPv6, or FQDN) of the system to which the request is being made.

-truststore *integer_or_string*
(Optional) Specifies the ID or name of the entry in the truststore for the system that is identified by the IP address.

-leave *flag*
(Optional) Requests the current system to leave its FlashSystem grid.

-remove *IPv4_IPv6_address_or_FQDN*
(Optional) Removes the specified system from the FlashSystem grid.

Description

This command can be used to create a new FlashSystem grid and manage the members of the FlashSystem grid. The system that creates the FlashSystem grid is referred to as the owner of the FlashSystem grid, and it assumes responsibility for adding and removing other systems. After creating a FlashSystem grid, the following steps must be performed in order to add a system:

1. Copy the certificate of the FlashSystem grid coordinator-system to the new system, and add it to the truststore by using the **mktruststore** CLI.
2. Request to join the FlashSystem grid by using the **join** option. Provide the IP address or the FQDN of the coordinator system and the entry, which the coordinator system has in the truststore.

© Copyright IBM Corp. 2003, 2025

3. Copy the certificate from the new system to the FlashSystem grid coordinator-system, and add it to the truststore.
4. On the coordinator system, accept the request to join the FlashSystem grid by using the **accept** option. Provide the IP address or the FQDN of the new system and the entry, which the new system has in the truststore.

To remove a system from a FlashSystem grid, use the **remove** option on the coordinator system. The system that is removed from the FlashSystem grid also needs to run the **leave** option.

To create or join a FlashSystem grid, the following requirements must be met:

- The system topology must be in the replication layer.

An invocation example

The following creates a new FlashSystem grid and makes this system the coordinator system.

```
manageflashgrid -create -name MyFlashGrid
```

The resulting output:

```
No feedback
```

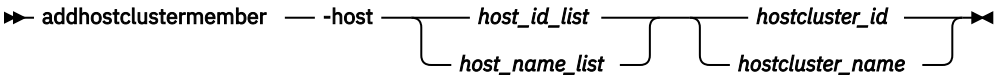
Chapter 21. Host commands

Use the host commands to work with host objects on your system.

addhostclustermember

Use the **addhostclustermember** command to add a host object to a host cluster.

Syntax



Parameters

-host ***host_id_list | host_name_list***
(Required) Specifies the host (by ID or name) to add to the host cluster.

hostcluster_id | hostcluster_name
(Required) Specifies (by ID or name) the host cluster that the host object is added to. The value for the ID must be a number and the value for the name must be an alphanumeric string.

Description

This command adds a host object to a host cluster.

When you add a host object to a host cluster, shared mappings are created. For example, if any host mappings match a host cluster mapping that is part of the same volume on the same Small Computer System Interface (SCSI) logical unit number (LUN) - with the same I/O groups - the host cluster assumes control of the mapping (which makes it a shared mapping).

Note: A host cannot be added to a host cluster if both have their individual throttling specifications defined. However, if either the host or host cluster throttling specification is present, the command succeeds.

Any mappings that do not match the shared host cluster mappings are managed by the host as private mappings.

Note: New mappings must not conflict with a shared mapping on a host system. The command fails when there are shared mappings that conflict with the host's private mappings. This includes either:

- A volume that is being mapped - but with different SCSI LUNs
- The host that has a different volume mapped but with the same SCSI LUN as a shared mapping of the host cluster

An invocation example that adds host 0 to host cluster 4

```
addhostclustermember -host 0 4
```

The resulting output:

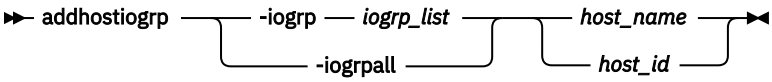
```
No feedback
```

© Copyright IBM Corp. 2003, 2025

addhostiogrp

Use the **addhostiogrp** command to map I/O groups to an existing host object.

Syntax



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.

Note: Hosts associated with a partition are by default in all I/O groups and cannot have I/O groups added or removed. The maximum number of I/O groups is hardware type dependent.

Description

This command allows you to map the list of I/O groups to the specified host object.

An invocation example

```
addhostiogrp -iogrpall testhost
```

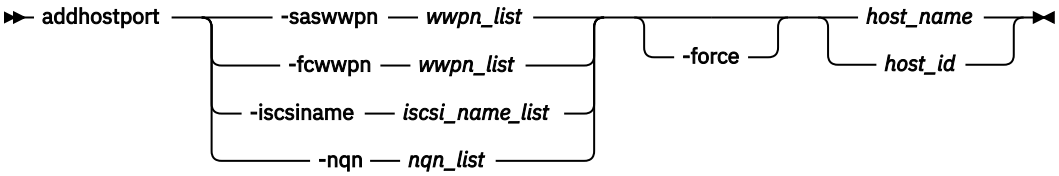
The resulting output:

```
No feedback
```

addhostport

Use the **addhostport** command to add worldwide port names (WWPNs), NVMe Qualified Names (NQNs), or InfiniBand port names to an existing host object.

Syntax



Parameters

-saswwpn *wwpn_list*
(Required if you do not use **-iscsiname**, **-fcwwpn**, or **-nqn**.) Specifies a list of Serial Attached SCSI (SAS) WWPNs with a 16-character hexadecimal string.

-fcwwpn *wwpn_list*
(Required if you do not use **-iscsiname**, **-saswwpn**, or **-nqn**.) Specifies a list of Fibre Channel (FC) WWPNs with a 16-character hexadecimal string.

-iscsiname *iscsi_name_list*
(Required if you do not use **-fcwwpn**, **-saswwpn**, or **-nqn**.) Specifies the comma-separated list of iSCSI names to add to the host. At least one WWPN or iSCSI name must be specified. You cannot use this parameter with the **-fcwwpn** or **-saswwpn** parameter.

-nqn *nqn_list*
(Required if you do not use **-fcwwpn**, **-saswwpn**, or **-iscsiname**.) Specifies the comma-separated list of NVMe Qualified Names (NQNs) to add to the host. The host protocol must be **nvme** to use this parameter.

-force
(Optional) Specifies that the list of ports be added to the host without the validation of any WWPNs or iSCSI names.

host_id | host_name
(Required) Specifies the host object to add ports to, either by ID or by name.

Description

This command adds a list of host bus adapter (HBA) WWPNs, iSCSI names, or NQNs to the specified host object. Any volumes that are mapped to this host object automatically map to the new ports.

Note: The host port can only be added if it is compatible with the protocol of the host (SCSI can have all port types except NQNs, NVMe can have only NQNs).

This command adds a list of host bus adapter (HBA) WWPNs, NVMe Qualified Names (NQNs), or InfiniBand IDs to the specified host object. Any volumes that are mapped to this host object automatically map to the new ports.

Only WWPNs that are logged-in unconfigured can be added. For a list of candidate WWPNs, use the **lssasportcandidate** or **lsfcportcandidate** command.

Only port names that are logged-in unconfigured can be added. For a list of candidate port names, use the **lsibportcandidate** or **lsfcportcandidate** command.

Some HBA device drivers do not log in to the fabric until they can recognize target logical unit numbers (LUNs). Because they do not log in, their WWPNs are not recognized as candidate ports. You can specify the **force** parameter with the **addhostport** command to stop the validation of the WWPN list.You can specify the **force** parameter with the **addhostport** command to stop the validation of the port name list.

Note:

When all I/O groups are removed from an iSCSI host, you cannot add a port to the iSCSI host until you map the iSCSI host to at least one I/O group. After you map the iSCSI host to at least one I/O group, resubmit the **addhostport** command. After you add the port to the host, you must create a host authentication entry using the **chhost** command.

The **addhostport** command fails if the:

- Host is mapped to a volume with more than one I/O group in the access set and the host port you add is an Internet Small Computer System Interface (iSCSI) name
- Port being added is from a host system that does not support volumes mapped from multiple I/O groups.

A host can have a single port type from the following list:

- sas
- fcscsi
- iscsi

- fcnvme or rdmanvme

An invocation example

```
addhostport -saswwpn 210100E08B251DD4 host1
```

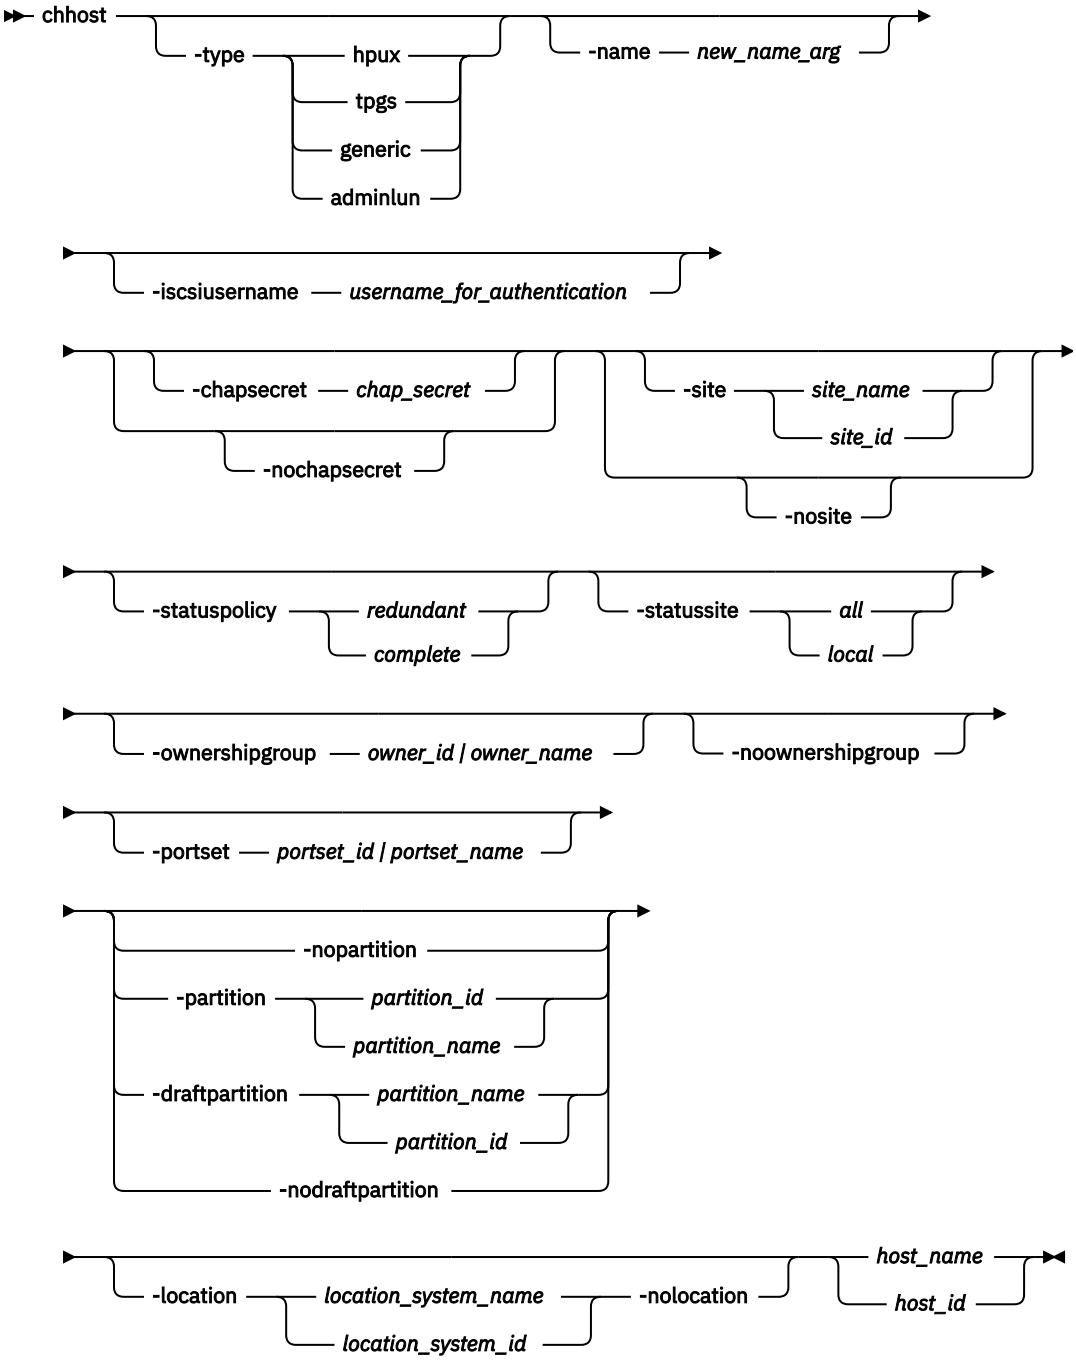
The resulting output:

```
No feedback
```

chhost

Use the **chhost** command to change the name or type of a host object. This command does not affect any existing host mappings.

Syntax



Note: Mapping a new portset to a host object can cause host application outage unless the newly mapped portset contains at least one port that is overlapping with the old portset. If volume protection is turned on, the administrator cannot change to a portset that doesn't have at least one overlapping port. To override that restriction and enforce a change, the global admin must run the **chsystem -vdiskprotectionenabled no** command to turn off volume protection for the system. Also, if a host is **offline**, the volume protection check is bypassed and **chhost** will allow a portset with non-overlapping ports to be used.

Parameters

-type hpux | tpgs | generic | openvms | adminlun | hide_secondary
(Optional) Specifies the type of host. The following values are the available host types:

- `generic` indicates the default.
- `tpgs` indicates when target port information changes (deprecated).
- `openvms` indicates OpenVMS (deprecated).
- `adminlun` indicates virtual volumes, which are enabled on the host.
- `hpux` indicates HP-UX firmware.
- `hide_secondary` (deprecated).

The host type for NVMe hosts cannot be changed. For more information about hosts that require the **type** parameter, see the host attachment documentation.

-name ***new_name_arg***
(Optional) Specifies the new name that you want to assign to the host object.

-iscsiusername ***username_for_authentication***
(Optional) Specifies the user name for a host object for one-way authentication for Internet Small Computer System Interface (iSCSI) host attachment login. If this parameter is specified, the value is taken as the "username" for one-way authentication to log in to the iSCSI host attachment. If you do not specify the **iscsiusername** parameter, the IQN of the host object is used as the user name by default. If no **iscsiusername** parameter is provided and multiple IQNs exist in the same host object, the user name of each IQN is the IQN itself. If the **iscsiusername** parameter is provided for a multiple IQN host, then log in to all hosts by using the provided IQN. If you use the **iscsiusername** parameter, you must also specify the **chapsecret** parameter.

-chapsecret ***chap_secret***
(Optional) Sets the Challenge Handshake Authentication Protocol (CHAP) secret that is 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 by using the **lsiscsiauth** command.

-nochapsecret
(Optional) Clears any previously set CHAP secret for this host. The **nochapsecret** parameter cannot be specified whether **chapsecret** is specified.

-site ***site_name | site_id***
(Optional) Specifies the numeric site value or site name of the host. The site name must be an alphanumeric value. The site ID must be 1 or 2. The host site only applies when using the stretched topology. If you are using highly available storage partitions then a location can be set instead of a site.

-nosite
(Optional) Resets the site value.

-statuspolicy ***redundant | complete***
(Optional) Specifies if *complete* or *redundant* policy is used for online status calculations.

- Complete connectivity is where online host status indicates complete connectivity, with every host port that is logged in to every node through which that host can access volumes, and degraded host status indicates one or more of these nodes does not have a login.
- Redundant connectivity calculation is where online host status indicates redundant connectivity, where enough host ports are logged in to enough nodes such that the removal of a single node or a single host port allows that host to access all of its volumes. In addition, there must be the same total login count from the host to each node within an I/O group. Degraded host status indicates non-redundant connectivity, in which a single point of failure prevents a host from accessing at least some of its volumes, or there is no single point of failure but the total login count from the host to each node within an I/O group is not the same.

The default value is *redundant*.

-statussite ***all | local***
(Optional) Specifies whether *all* sites or only the *local* host site connectivity is used to determine online status. *local* is only applicable in a dual site topology for a host with a valid site. The default value is *all*.

-ownershipgroup ***owner_id | owner_name***
(Optional) The name or ID of the ownership group to which the object is being added.

-noownershipgroup
(Optional) If specified, the object is removed from the ownership group to which it belongs.

-portset
(Optional) The name or ID of the portset to associate with the host object.

Note: Admin can run device discovery from the host after running this command as changing portset of a host changes I/O paths for the host. Running discovery allows host to discover the devices from valid I/O paths.

-partition ***partition_id|partition_name***
(Optional) The name or ID of the storage partition to associate with the host object.

Note: If the host is associated with a partition, only the following host parameters are allowed:

- **-statuspolicy**
- **-statussite**
- **-portset**
- **-nopartition**.

-nopartition
(Optional) If specified, the host is removed from the partition to which it is associated.

-draftpartition
Attempt to associate this host with the specified draft partition.

-nodraftpartition
Attempt to remove this host from the draft partition.

-location ***location_system_name | location_system_id***
(Optional) Specifies the system ID or system name that is co-located with this host. The parameter is only valid if the host is in a highly available storage partition. Setting a host location will result in the I/O submitted by this host being submitted to the system in the same location while high availability is established as this can reduce the latency of I/O operations. If a host location is not set then the host will normally submit I/O operations to the storage partition's preferred system.

-nolocation
(Optional) Remove the currently set location from 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 host mappings.

Note: When all I/O groups are removed from an iSCSI host, the **lsiscsiauth** command does not display the authentication entry for that host. Use the **addhostiogrp** command to map the iSCSI host to at least one I/O group, and then use the **addhostport** command to add the iSCSI port into it. You must also add authentication for that host by using the **chhost** command with either the **chapsecret** or **nochapsecret** parameter.

- The protocol of the portset must match the protocol of the host

Note: The 100 Gbps ports are mainly used for NVMe over RDMA and NVMe over TCP connections to host attachments. Any attempt to map an iSCSI host using **chhost** to a 100 Gbps portset will fail. This iSCSI

mapping using **chhost** cannot be completed through any plugins or GUI. This fails because a 100 Gbps portset will not match the performance expectation for iSCSI host attach connectivity. However, user can use **-force** flag CLI command to map an iSCSI host to a 100 Gbps port for less demanding application or boot from SAN applications.

- To change the portset by using the **chhost** command user has to check portset **port_type** and it must match the type of the host (FC, NVMe, or iSCSI) created.

An invocation example

```
chhost -portset myportset myhost
```

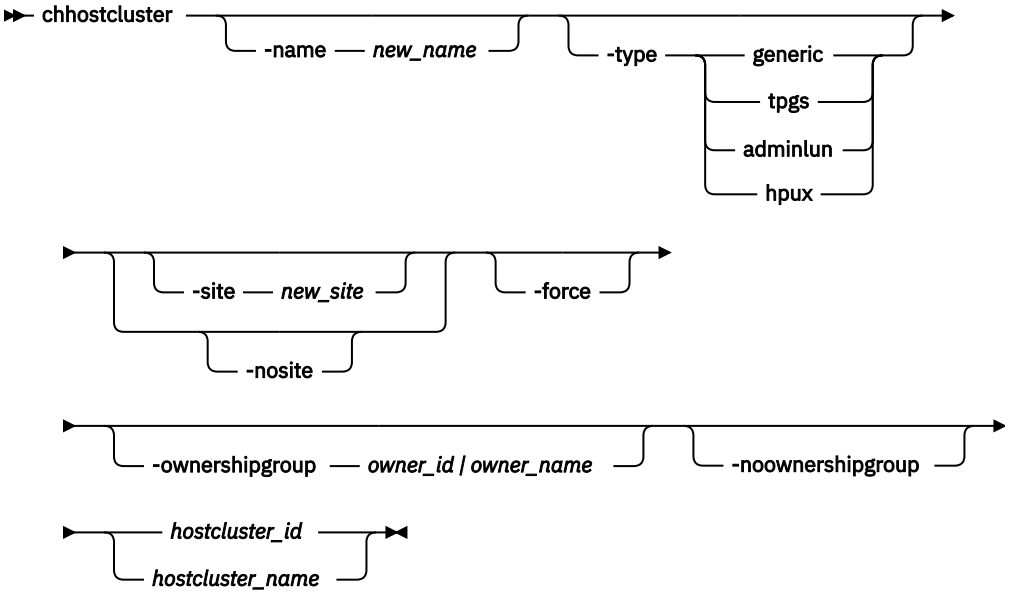
The following output is displayed:

```
No feedback
```

chhostcluster

Use the **chhostcluster** to change the name, type, or site of a host cluster object that is part of a host cluster.

Syntax



Parameters

-name *new_name*
(Optional) Specifies a name for a host cluster object. The value must be an alphanumeric string.

-type *new_type*
(Optional) Changes the type of the all hosts in the host cluster. The following values are available host types:

- `generic` indicates the default.
- `tpgs` indicates when target port information changes (deprecated).
- `openvms` indicates OpenVMS (deprecated).
- `adminlun` indicates virtual volumes, which are enabled on the host.
- `hpux` indicates HP-UX firmware.

- `hide_secondary` (deprecated).

The host type for NVMe hosts cannot be changed. For more information about hosts that require the **type** parameter, see the host attachment documentation.

-site ***new_site***
(Optional) Changes the site for all hosts in the host cluster. The value must be an alphanumeric string and the default is `site0` (meaning there is no site).

-nosite
(Optional) Resets the site value.

-force
(Optional) Specify this parameter to change the host cluster's site and also change a site for at least one host in that host cluster. You do not have to specify this parameter if you are changing site `0`.

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

-ownershipgroup ***owner_id | owner_name***
(Optional) The name or ID of the ownership group to which the object is being added.

-noownershipgroup
(Optional) If specified, the object is removed from the ownership group to which it belongs.

hostcluster_id | hostcluster_name
(Required) Specifies (by ID or name) the host cluster object to change. The value for the ID must be a number and the value for the name must be an alphanumeric string.

Description

This command changes the name, type, or a site on a host cluster object.

If you assign a host to a site and it is not specified as `site0`, you must use **-force** to change it.

Important: Do not specify `-site site0`, instead use **-nosite**.

If you specify the **-site** and **-type** parameters, the site and type properties that are changed are not attributes of the host cluster object. They are properties of the individual hosts that are members of the host cluster. These properties can be modified by using the **chhostcluster** command to change the attribute values on every host in the host cluster simultaneously (rather than modifying each host separately).

An invocation example that changes settings of host cluster hostcluster0

```
chhostcluster -name myhostcluster hostcluster0
```

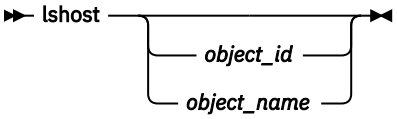
The following detailed output is displayed:

```
No feedback
```

lshost

Use the **lshost** command to generate a list with concise information about all the hosts visible to the system and detailed information about a single host.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

object_id | object_name
(Optional) Specifies the name or ID of an object. When you use this parameter, the detailed view of the specific object is returned and any value that is specified by the **-filtervalue** parameter is ignored. If you do not specify the *object_id* | *object_name* parameter, the concise view of all objects that match the filtering requirements that are specified by the **-filtervalue** parameter are displayed.

Note: In case any existing system is upgraded then all existing host objects will associate with the default portset. Similarly if host objects are created without specifying a portset they will associate with the default portset.

Description

This command returns a concise list or a detailed view of hosts visible to the system.

For Fibre Channel (FC) ports, the `node_logged_in_count` field provides the number of nodes that the host port is logged in to. For Internet Small Computer System Interface (iSCSI) ports, the `node_logged_in_count` field provides the number of iSCSI sessions from the host iSCSI qualified name (IQN).

You can map an iSCSI host to volumes that are accessible through multiple I/O groups. iSCSI hosts can access volumes that are accessible through multiple I/O groups (and single I/O groups). An iSCSI host that is mapped to a volume accessible through multiple I/O groups is `online` if it has at least one active iSCSI session with each I/O group of the access set. If volumes are not mapped to an iSCSI host, it is `degraded`.

The following list provides the different states for a fabric attach FC host port:

active
The host port is active if all nodes with volume mappings have a login for the specified worldwide port name (WWPN) and at least one node received SCSI commands from the WWPN within the last 5 minutes.

degraded
The host port is degraded if one or more nodes with volume mappings do not have a login for the specified WWPN.

inactive
The host port is inactive if all the nodes with volume mappings have a login for the specified WWPN but no nodes see any Small Computer System Interface (SCSI) commands from the WWPN within the last 5 minutes.

offline
The host port is offline if the input/output (I/O) group with volume mappings does not have a login for the specified WWPN.

The following list provides the different states for a direct attach FC host port:

active
The host port is active if a node has a login for the specified WWPN and the node receives SCSI commands from the WWPN within the last 5 minutes.

inactive
The host port is inactive if all the nodes with volume mappings have a login for the specified WWPN but no nodes see any Small Computer System Interface (SCSI) commands from the WWPN within the last 5 minutes.

offline
The host port is offline if no login exists for the specified WWPN.

If a host does not have any volume mappings, it is reported as offline or inactive.

The following list provides the different states for fabric attached NVMe over FC, NVMe over RDMA, and NVMe over TCP host ports:

active
The host port is active if all nodes with volume mappings have a login for the specified NVMe Qualified Name (NQN) and at least one node has received NVMe commands from the NQN within the last 5 minutes.

degraded
The host port is degraded if one or more nodes with volume mappings do not have a login for the specified NQN.

inactive
The host port is inactive if all the nodes with volume mappings have a login for the specified NQN but no nodes have seen any NVMe commands from the NQN within the last 5 minutes.

offline
The host port is offline if the I/O group with volume mappings does not have a login for the specified NQN.

Note: The **lshost** command presents a list of host HBA ports that are logged in to nodes. However, situations exist when the information presented can include host HBA ports that are no longer logged in or even part of the SAN fabric. For example, a host HBA port is unplugged from a switch, but **lshost** still shows the WWPN logged in to all nodes. If this action occurs, the incorrect entry is removed when another device is plugged in to the same switch port that previously contained the removed host HBA port.

The following list provides the different states for a specified `iscsiname`:

active
The `iscsiname` is active if the I/O group with volume mappings has at least one associated iSCSI session for the specified `iscsiname`.

inactive
The `iscsiname` is inactive if the host has no volume mappings but at least one iSCSI session for the specified `iscsiname` is present.

offline
The `iscsiname` is offline if volumes with access to the I/O group is mapped to a specified **iscsiname** that has no associated iSCSI sessions.

The following list provides the different states for host_status:

online
The host has full connectivity. A host that uses just one style of connectivity is online if it uses one of these types:

Fabric Attach Fibre Channel
Every port is active or inactive, and is logged in to every online node in each I/O group in which the host has volume mappings.

Direct Attach Fibre Channel
The host has an active or inactive login to every node in each I/O group in which the host has volume mappings.

Internet Small Computer System Interface (iSCSI)
The host has an iSCSI session with each I/O group with which the host has volume mappings.

NVMe over FC (FC-NVMe), NVMe over RDMA (RDMA-NVMe), and NVMe over TCP (TCP-NVMe)
All host NQNs have NVMe associations with every online node in each I/O group in which the host has volume mappings.

offline
The host has no connectivity. The reason might be because the host is powered off and is not on.

Remember: If an iSCSI host is only logged in to the I/O group for which it is not configured, the associated host object status is `offline`.

degraded
The host is not fully connected, which might be introduced by a configuration error or a hardware failure. It can cause a loss of access during any planned maintenance activity and must be corrected as soon as possible.

Remember: An iSCSI host that has no mapped volumes is `degraded` if it is logged in to some, but not all, of the I/O groups to which it belongs.

This table shows the possible outputs:

*Table 106. **`lshost`** output*

Attribute	Description
id	Indicates the unique host ID. The value is an alphanumeric value.
name	Indicates the unique host name. The value is an alphanumeric string.
port_count	Indicates the number of ports.
type	Indicates the host type.
iogrp_count	Indicates the number of I/O groups.
status	Indicates whether the host is `online` or `offline`.
protocol	Indicates whether the host uses scsi, fcnvme, tcpnvme, or rdmanvme to access the storage.
WWPN	Indicates the worldwide port name (WWPN) with a 16-character hexadecimal string.
SAS_WWPN	Indicates the serial-attached SCSI (SAS) WWPN with a 16-character hexadecimal string.
nqn	Indicates the NVMe Qualified Name.
node_logged_in_count	Indicates the number of nodes the WWPN or NQN is logged in to.
state	Indicates the state of the SAS WWPN login. The values are: • `offline` • `inactive` • `active`
sas_wwpn_count	Indicates the number of configured SAS WWPNs.
nqn_count	Indicates the number of NVMe NQN ports assigned to the host.
host_cluster_id	Indicates the unique ID for a host cluster.
host_cluster_name	Indicates the unique name for a host cluster.
protocol	Displays the host connectivity protocol.
portset_id	Displays the ID of the portset.
portset_name	Displays the name of the portset.

*Table 106. **lshost** output (continued)*	
Attribute	**Description**
iSCSI_name	Displays the iSCSI name.
partition_id	Displays the ID of the storage partition in which the host resides.
partition_name	Displays the name of the storage partition in which the host resides.
location1_status	Displays the status of the host at the location 1 system that is associated with the partition.
location2_status	Displays the status of the host at the location 2 system that is associated with the partition.
draft_partition_id	Displays the ID for draft partition associated with this host.
draft_partition_name	Displays the name of draft partition associated with this host.
ungrouped_volume_mapping	Indicates that this host has a mapping to a volume that is not in a volume group.
location_system_name	Indicates the name of the system for the host location.

lshostports

Use the **lshostports** command to list the host port information for the IBM Storage Insights.

Syntax



Description

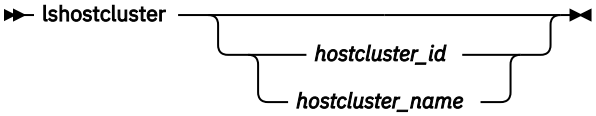
This table provides the attribute values that can be displayed as output view data.

*Table 107. **lshostports** attribute values*	
Attribute	**Value**
host_id	Indicates the ID of the host.
host_name	Indicates the name of the host.
port_id	Indicates the ID of the port.
WWPN	Indicates the worldwide port name (WWPN) of the port fibre-channel.
SAS_WWPN	Indicates the WWPN of the SAS port.
iscsi_name	Indicates the name of the ISCSI port.
nqn	Indicates the NVMe Qualified Name.
node_logged_in_count	Indicates the number of nodes logged into the fiber channel.
state	Indicates the state of the host port. State can be: • active • degraded • inactive • offline

lshostcluster

Use the **lshostcluster** command to generate a list with concise information about all the host clusters that are defined on the system or detailed information about a single host cluster.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

hostcluster_id | hostcluster_name
(Optional) Specifies the ID or name of the host cluster to display information about. If you do not enter a host cluster ID or name, the command displays a list of all host clusters.

Description

This command lists concise information about all the host clusters that are defined on the system or detailed information about a single host cluster.

This table provides the attribute values that can be displayed as output view data.

Table 108. lshostcluster output

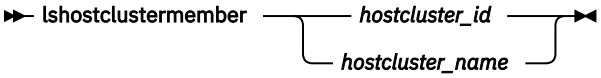
Attribute	Description
id	Indicates the host cluster ID.
name	Indicates host cluster name. The value must be an alphanumeric string of no more than 64 characters.
status	Indicates the status of the host cluster. The values are: • `online` indicates that all hosts or members are online. • `host_degraded` indicates that no hosts are offline but at least one host is degraded. • `host_cluster_degraded` indicates that one or more hosts are offline and at least one host is online or degraded. • `offline` indicates that all hosts are offline or there are no hosts or members in the host cluster.
host_count	Indicates the number of hosts that are in the host cluster. The value must be a number in the range 0 - 127.
mapping_count	Indicates the number of shared mappings between the host cluster and any existing volumes. The value must be a number in the range 0 - 2047.

Table 108. lshostcluster output (continued)	
Attribute	**Description**
port_count	Indicates the number of host ports that are used for the host cluster mappings to any volumes. The value must be a number in the range 0 - 255.

lshostclustermember

Use the **lshostclustermember** command to generate a list with host information for hosts that belong to the specified host cluster.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

hostcluster_id | hostcluster_name
(Required) Specifies the ID or name of the host cluster. If you do not enter a host cluster ID or name, the command displays a list of all recognized host clusters and volume mappings. The value for the ID must be a number and the value for the name must be an alphanumeric string.

Description

This command information about all hosts that belong to the specified host cluster.

This table provides the attribute values that can be displayed as output view data.

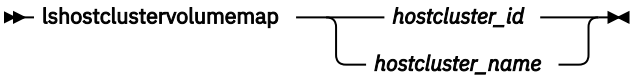
Table 109. **lshostclustermember** output	
Attribute	**Description**
host_id	Indicates the unique ID of the host.
host_name	Indicates the host name.
status	Indicates the status of a host for a host cluster. The values are: • online indicates that all hosts or members are online. • degraded indicates that no hosts are offline but at least one host is degraded. • offline indicates that all hosts are offline or there are no hosts or members in the host cluster.

Table 109. **lshostclustermember** output (continued)	
Attribute	**Description**
type	Indicates the host type. The values are: • generic • tpgs (deprecated) • openvms (deprecated) • adminlun • hpux • hide_secondary (deprecated)
site_id	Indicates the site ID for the host.
site_name	Indicates the site name for the host.
protocol	Indicates the protocol which the hosts in the host cluster use to access the storage.

lshostclustervolumemap

Use the **lshostclustervolumemap** command to display a list of volumes that are mapped to all host clusters (or to a specific host cluster).

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

hostcluster_id | hostcluster_name
(Required) Specifies the ID or name for the host cluster that is being mapped to a volume. The command displays a list of all the volumes that are mapped to the specified host cluster and additionally indicates the Small Computer System Interface (SCSI) ID through which they are mapped. If you do not enter a host cluster ID or name, the command displays a list of all recognized host clusters and volume mappings. The value for the ID must be a number and the value for the name must be an alphanumeric string.

Description

This command displays a list of volumes that are mapped to all host clusters or to a specific host cluster.

This table provides the attribute values that can be displayed as output view data.

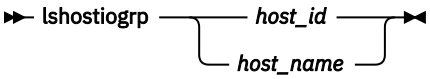
Table 110. **lshostclustervolumemap** output	
Attribute	**Description**
id	Indicates the host cluster ID. The value must be a number in the range 0 - 127.
name	Indicates host cluster name. The value must be an alphanumeric string of no more than 64 characters.

Table 110. **lshostclustervolumemap** output (continued)	
Attribute	**Description**
SCSI_id	Indicates the unique ID (volume ID) that is mapped from a host cluster to an I/O group volume. The value must be a number in the range 0 - 2047.
volume_id	Indicates the unique ID of the volume that is mapped to the host cluster. The value must be a number.
volume_name	Indicates the name for a volume that is mapped to a host cluster. The value must be an alphanumeric string.
volume_UID	Indicates the unique UID of a volume.The value must be an alphanumeric string.
IO_group_id	Indicates the unique ID from the I/O group that the host cluster and volume (from the mapping) are part if. The value must be a number in the range 0 - 3.
IO_group_name	Indicates the I/O group name. The value must be an alphanumeric string.

lshostiogrp

Use the **lshostiogrp** command to display a list the I/O groups that are associated with a specified host.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

host_id | host_name
(Required) The name or ID of the host for which the list of I/O groups is required.

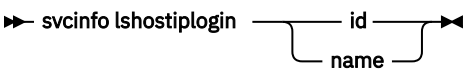
Description

This command displays a list of all the I/O groups that are mapped to the specified host.

lshostiplogin

Use the **lshostiplogin** command to list the login session type such as associated host object, login counts login protocol and other details for hosts that are identified by their iSCSI Qualified Name (IQN) and logged in to Ethernet ports that are configured with IP addresses.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

id
(Optional) Object identifier that indicates a configured host IQN, for which login details are to be displayed. This value must be a number in the range 0 - 2047 (Max number of configured iSCSI IQNs per SVC cluster [512 x 4 = 2048]).

name
(Optional) Object name that indicates a configured host IQN name, for which login details are to be displayed.

Note:

Following are added as ownership rules:

- Global user is allowed to view login information of all host objects.
- Restricted users can view login information only for owned host objects and global host objects.

Description

This command lists the login session type and other details.

This table provides the attribute values that can be displayed as output view data.

Table 111. ***svcinfo lshostiplogin*** *output*

Attribute	Description
id	Object identifier that indicates a configured host IQN for which login details are displayed.
name	Unique IQN for an iSCSI host. This IQN must belong to a configured host object.
host_id	Host object ID to which this IQN belongs.
host_name	Host name corresponding to **host_id**.
login_count	Total number of login sessions that are established by the host to cluster ports to all I/O groups.
login_protocol	Specifies the connection protocol that is used by the configured host to establish sessions with Ethernet ports configured with IP addresses. This value can be one of the following values: • iSCSI • mixed
login	An integer that refers to a single session from the host to any node in the cluster. This value is a number in the range of 0 - 31.
node_id	Node ID of the node to which a login session, referred by **login** field, is established by host IQN. This value is an integer in the range of 1 - 32.
node_name	Node name of the node to which a login session, referred by **login** field, is established by host IQN. This value is a string with a maximum of 16 characters.

*Table 111. **svcinfo lshostiplogin** output (continued)*	
Attribute	**Description**
port_id	Port identifier that indicates the Ethernet port that is displayed by the **lsportethernet** view to which login is established. This value is a number in the range of 1 - 12.
protocol	Indicates the connection protocol that is used by the configured IP host IQN to establish a login session referred by the **login** field. This value can be one of the following values: • iSCSI
ip_version	Indicates the IP addressing version. The value that is displayed is either **IPv4** or **IPv6**.
failover	Identifies whether the session refers to the local or failover IP address. The value that is displayed is either **yes** or **no**.

Note: Individual IQNs within a host object may be associated with different ownership. This command will display host objects' login information based on ownership.

lsiscsiauth

Use the **lsiscsiauth** command to list the Challenge Handshake Authentication Protocol (CHAP) secret that is configured for authenticating an entity to the system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

Description

This command lists the CHAP secret that is configured for authenticating an entity to the system. The command also displays the configured iSCSI authentication method. The iscsi_auth_method field can have values of none or chap.

When you create an iSCSI host by using the **mkhost** command with the **iscsiname** parameter, the host is initially configured with the authentication method as none, and no CHAP secret is set. To set a CHAP secret for authenticating the iSCSI host with the system, use the **chhost** command with the **chapsecret** parameter.

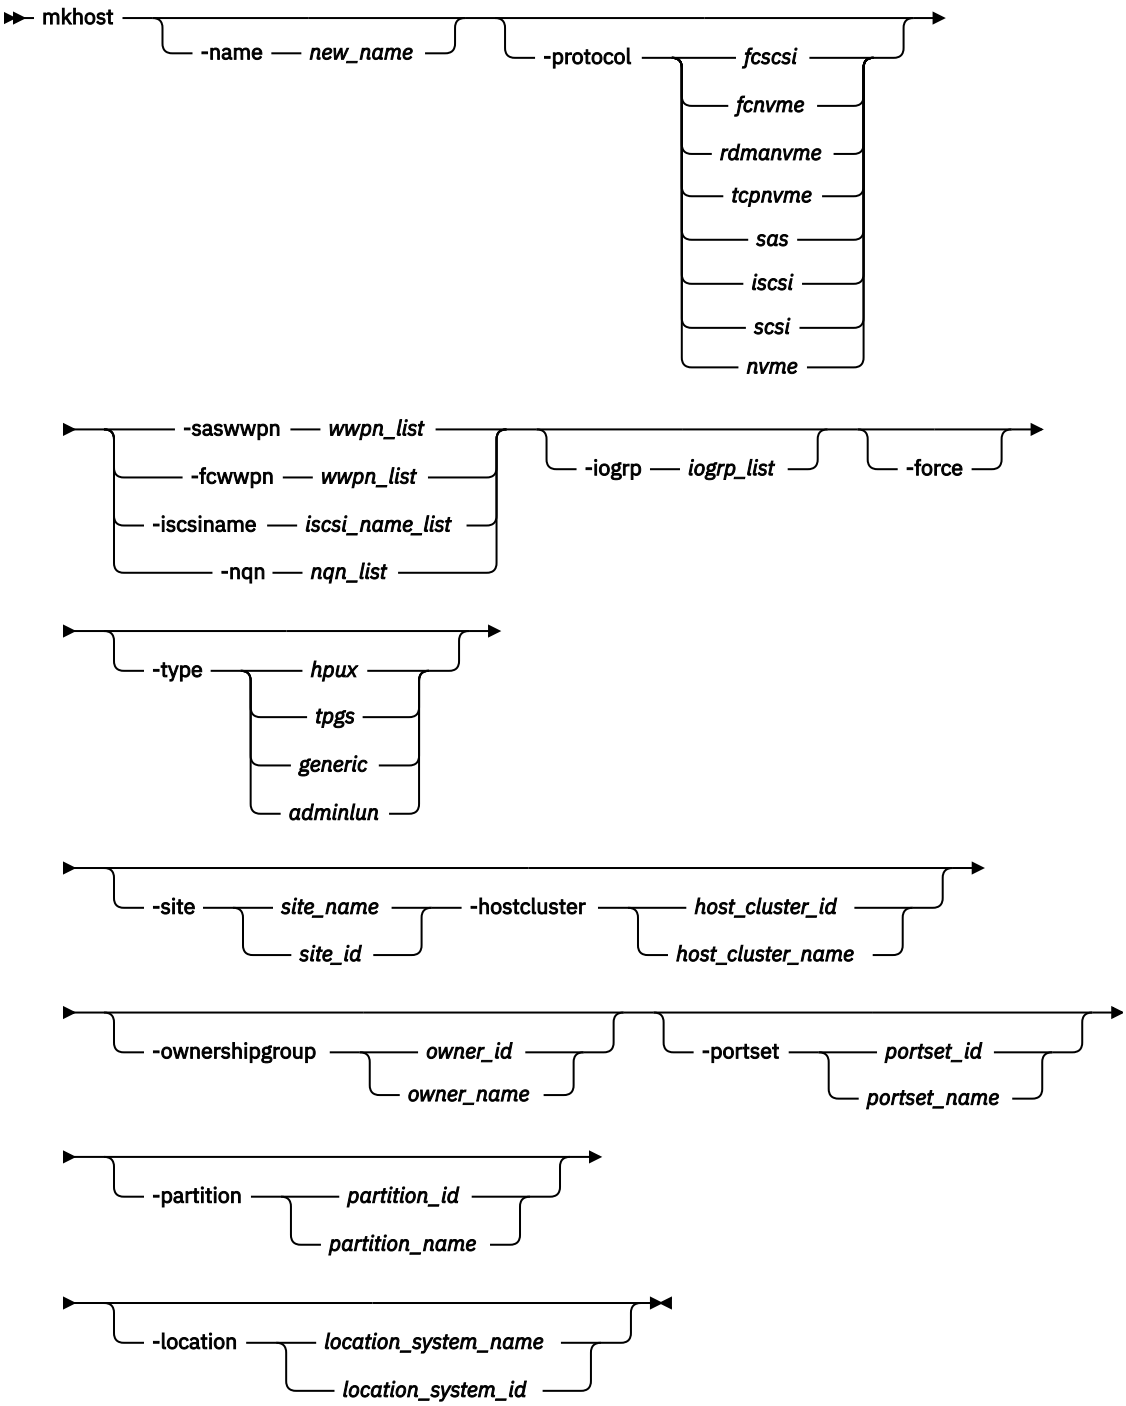
This table provides the attribute values that can be displayed as output view data.

Table 112. lsiscsiauth output	
Attribute	**Description**
type	Indicates the iSCSI system type.
id	Indicates the iSCSI system ID.
name	Indicates the iSCSI system name.
iscsi_auth_method	Indicates the iSCSI authentication method.
iscsi_chap_secret	Indicates whether an iSCSI CHAP secret exists.
cluster_iscsi_auth_method	Indicates clustered system iSCSI authentication method.
cluster_iscsi_chap_secret	Indicates the clustered system iSCSI configured CHAP secret.
iscsiusername	Indicates the iSCSI user name.

mkhost

Use the **mkhost** command to create a logical host object.

Syntax



Parameters

-name *new_name*
(Optional) Specifies a name or label for the new host object.

-protocol *fcscsi | fcnvme | rdmanvme | tcpnvme | sas | iscsi | scsi | nvme*
(Optional) Specifies the protocol that is used by the host to communicate with the storage system. The default is **fcscsi** parameter. The **scsi** and **nvme** options are deprecated.

-saswwpn *wwpn_list*
(Required if you do not use **-iscsiname**, **-fcwwpn**, or **-nqn**). Specifies a list of Serial Attached SCSI (SAS) WWPNs with a 16-character hexadecimal string.

-fcwwpn *wwpn_list*
(Required if you do not use **-saswwpn**, **-iscsiname**, or **-nqn**). Specifies a list of Fibre Channel (FC) WWPNs with a 16-character hexadecimal string.

-iscsiname *iscsi_name_list*
(Required if you do not use **-fcwwpn**, **-saswwpn**, or **-nqn**). Specifies the comma-separated list of iSCSI names to add to the host. At least one WWPN or iSCSI name must be specified. You cannot use this parameter with the **-fcwwpn** or **-saswwpn** parameter.

-nqn *nqn_list*
(Required if you do not use **-fcwwpn**, **-saswwpn**, or **-iscsiname**.) Specifies the comma-separated list of NVMe Qualified Names (NQNs) to add to the host. The host protocol must be **fcnvme**, **rdmanvme**, or **tcpnvme** to use this parameter.

-iogrp *iogrp_list*
(Optional) Specifies a set of one or more input/output (I/O) groups that the host can access the volumes from. I/O groups are specified by 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.

-force
(Optional) Specifies that a logical host object is created without validation of the WWPNs.

-type hpux | tpgs | generic | openvms | adminlun | hide_secondary
(Optional) Specifies the type of host. The following values are the available host types:

- `generic` indicates the default.
- `tpgs` indicates when target port information changes (deprecated).
- `openvms` indicates OpenVMS (deprecated).
- `adminlun` indicates virtual volumes, which are enabled on the host.
- `hpux` indicates HP-UX firmware.
- `hide_secondary` (deprecated)

The host type for NVMe hosts cannot be changed. For more information about hosts that require the **type** parameter, see the host attachment documentation.

-site *site_name | site_id*
(Optional) Specifies the numeric site value or site name of the host. The site name must be an alphanumeric value. The site ID must be 1 or 2. The host site only applies when using the stretched topology. If you are using highly available storage partitions then a location can be set instead of a site.

-hostcluster *host_cluster_id | host_cluster_name*
(Optional) Specifies the host cluster ID (numerical value) or name (alphanumeric value) that the new host object is created in.

-ownershipgroup *owner_id | owner_name*
(Optional) The name or ID of the ownership group to which the object is being added.

-portset *portset_id | portset_name*
(Optional) The name or ID of the portset to associate with the host object.

Note: If a portset that is mapped to the host is not specified, the host is associated with the default portset. For Ethernet portsets, the default portset is *portset0*. For Fiber Channel portsets, the default portset is *portset64*. These are following ownership id matching criteria:

- Restricted owner *A* cannot create and map host object with owner *A* and with portset that is owned by owner B.
- Global owner cannot create a host object with restricted owner *A* and map to portset with restricted owner *B*.
- Restricted or global owner *A* can create and map host object that is completely owned by *A* with portset that is owned by global owner.
- Global owner can create a host object with restricted owner *B* and map to portset with restricted owner *B*.
- Policy to check portset *port_type* must match the type of the host (Fibre Channel, NVMe, or iSCSI) created.

-partition *partition_id | partition_name*
(Optional) The name or ID of the storage partition to associate with the host object. If the host is created with a partition, the **-type** parameter defaults to *generic* and the **-iogrp** parameter defaults to *all*. Only the following other parameters are allowed:

- **-fcwwpn** *wwpn_list*
- (Optional) **-protocol** *fcscsi*
- (Optional) **-portset**
- (Optional) **-name**
- (Optional) **-force**

-location *location_system_name | location_system_id*
(Optional) Specifies the system ID or system name that is co-located with this host. The parameter is only valid if the host is in a highly available storage partition. Setting a host location will result in the I/O submitted by this host being submitted to the system in the same location while high availability is established as this can reduce the latency of I/O operations. If a host location is not set then the host will normally submit I/O operations to the storage partition's preferred system.

Description

The **mkhost** command associates one or more HBA WWPNs, iSCSI names, or NQNs with a logical host object. This command creates a new host. The ID is displayed when the command completes. Subsequently, you can use this object when you map volumes to hosts by using the **mkvdiskhostmap** command. If you create a host directly inside a host cluster by specifying **-hostcluster**, it inherits any shared mappings that exist.

NQNs can be added only if the **protocol** is `fcnvme`, `rdmanvme` or `tcpnvme`. All of the other port types can be added only if the host **protocol** is `fcscsi`. Dual-protocol hosts are not supported. A host cannot be added to a host cluster that has a different protocol to the host. NVMe hosts must have the **type** set to `default`.

Issue the **mkhost** command only once. The clustered system scans the fabric for WWPNs in the host zone. The system itself cannot filter into the hosts to determine which WWPNs are in which hosts. Therefore, you must use the **mkhost** command to identify the hosts.

After you identify the hosts, mappings are created between hosts and volumes. These mappings effectively present the volumes to the hosts to which they are mapped. All WWPNs in the host object are mapped to the volumes.

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 system.

Note: The 100 GB ports are mainly used for NVMe over RDMA and NVMe over TCP connections to host attachments. Any attempt to map an iSCSI host by using **mkhost** to a 100 Gbps portset will fail. This iSCSI mapping using **mkhost** cannot be completed through any plugins or GUI. This fails because a 100 Gbps portset will not match the performance expectations for iSCSI host attach connectivity. However,

users can use the **-force** flag CLI command to map an iSCSI host to a 100 Gbps port for less demanding application or boot from SAN applications.

- While upgrading the CCU, a new host protocol is automatically updated for all the existing hosts based on their port type.

For additional information, see the **mkvdiskhostmap**, **lsibportcandidate**, and **lsfcportcandidate** commands.

For more information about parameter requirements for your specific host, refer to the following support site: http://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss.

An invocation example

```
mkhost -iscsiname iqn.localhost.hostid.7f000001 -name newhost
```

The resulting output:

```
Host, id [10], successfully created
```

An invocation example

```
mkhost -fcwwpn 210100E08B251EE6:210100F08C262EE7 -type openvms
```

The resulting output:

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

An invocation example

```
svctask mkhost -force -name fc-nvmehost -nqn
nqn.2014-08.org.nvmexpress:uuid:449f8291-9c1e-446c-95c1-0942f55fa208 -portset fcnvme -protocol
fcnvme
```

The resulting output:

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

An invocation example to create hostone host with an IQN iqn.localhost.hostid.7f000001 and associate it with myportset

```
mkhost -iscsiname iqn.localhost.hostid.7f000001 -name hostone -portset myportset
```

The resulting output:

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

An invocation example to create hosttwo host with an IQN iqn.localhost.hostid.7f000002 and associate it with default portset0

```
mkhost -iscsiname iqn.localhost.hostid.7f000002 -name hosttwo
```

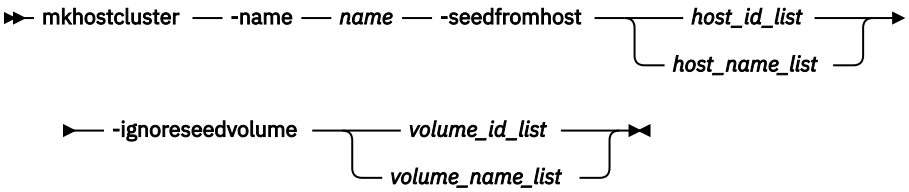
The resulting output:

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

mkhostcluster

Use the **mkhostcluster** command to create a host cluster object.

Syntax



Parameters

-name *name*
(Optional) Specifies the name of the host cluster object.

-seedfromhost *host_id_list | host_name_list*
(Optional) Adds the specified host to the host cluster. The host cluster mappings to existing volumes then become shared host cluster mappings. If a list of hosts is provided, the hosts are mapped to the same volume that uses the same I/O group and with the same SCSI LUN.

Note: These become shared mappings unless explicitly excluded with **-ignoreseedvolume**.

-ignoreseedvolume *volume_id_list | volume_name_list*
(Optional) Specifies volumes that are not part of the shared host cluster mappings. These volumes remain privately mapped to the host or hosts. You must specify **-seedfromhost** if you specify this parameter.

Description

This command is used to create a host cluster object.

Note: This command fails if any of the specified seeding hosts has an associated host throttle.

An invocation example that creates host cluster myhostcluster and obtains its mappings from host myhost1

```
mkhostcluster -name myhostcluster -seedfromhost myhost1
```

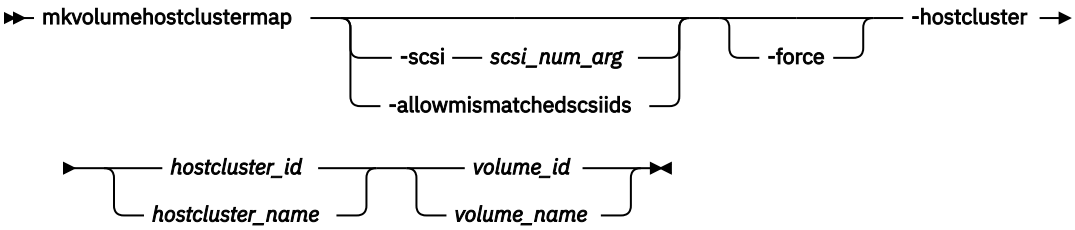
The detailed resulting output:

```
No feedback
```

mkvolumehostclustermap

Use the **mkvolumehostclustermap** command to generate a new mapping between a volume and a host cluster on a clustered system. This volume is then accessible for input or output (I/O) operations to the specified host cluster.

Syntax



Parameters

-scsi *scsi_num_arg*
(Optional) Specifies the Small Computer System Interface (SCSI) logical unit number (LUN) ID to assign to a volume on the specified host cluster. The SCSI LUN ID is assigned to the volume on the host cluster for all I/O groups that provide access to the volume.

Note: You must use the next available SCSI LUN ID for each host in the host cluster.

The **-scsi** parameter is not supported for NVMe hosts. When a volume is mapped to a host that uses the NVMe protocol, the volume namespace ID (NSID) is assigned by the system. This parameter is mutually exclusive with **-allowmismatchedscsiids**. This parameter cannot be used with the **-scsi** parameter and is not supported for NVMe hosts.

-allowmismatchedscsiids
(Optional) When you create a new host cluster mapping, this parameter allows the system to allocate non-identical SCSI LUN IDs in the access I/O groups for the volume. This parameter is mutually exclusive with **-scsi**.

-force
(Optional) Forces a new mapping. Specify this parameter to map a volume to a host cluster and that volume is already mapped to at least one host in a different host cluster.

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

-hostcluster *hostcluster_id | hostcluster_name*
(Required) Specifies the host cluster (by ID or name) to map to the volume. The value for the ID must be a number and the value for the name must be an alphanumeric string.

volume_id | volume_name
(Optional) Specifies the volume by ID or name. The value for the ID must be a number and the value for the name must be an alphanumeric string.

Description

This command generates a new mapping between a volume and a host cluster on a system. This volume is then accessible for input or output (I/O) operations to the specified host cluster. The protocol of the host cluster must be compatible with that of the volume.

If you do not specify a SCSI LUN ID, the system attempts to assign the lowest common SCSI LUN ID in all access I/O groups. The returned ID is the same for all I/O groups to which the volume is mapped. If an identical SCSI LUN ID is not available in all of the access I/O groups, the command fails.

Use the **-allowmismatchedscsiids** parameter to allow the system to allocate non-identical SCSI LUN IDs in the access I/O groups. The lowest value available in each access I/O group is used. It might not

be the same in all access I/O groups. Make sure that the host supports this configuration. This parameter cannot be used with the **-scsi** parameter and is not supported for NVMe hosts. If a volume is mapped to a host that uses the NVMe protocol, the volume namespace ID (NSID) is assigned by the system.

An invocation example that maps volume 0 to host cluster 0

```
mkvolumehostclustermap -hostcluster 0 0
```

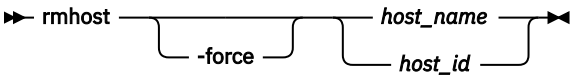
The detailed resulting output:

```
No feedback
```

rmhost

Use the **rmhost** command to delete a host object.

Syntax



Parameters

-force
(Optional) Specifies that you want the system to delete the host object even if mappings still exist between this host and volumes. When the **-force** parameter is specified, the mappings are deleted before the host object is deleted.

host_name | host_id
(Required) Specifies the host object to delete, either by ID or by name.

Description

The **rmhost** command deletes the logical host object. The WWPNs that were contained by this host object (if it is still connected and logged in to the fabric) are returned to the unconfigured state. When you issue the **lsfcportcandidate** or **lssasportcandidate** command, the host objects are listed as candidate ports.

The **rmhost** command deletes the logical host object. The identifiers that were contained by this host object (if it is still connected and logged in to the fabric) are returned to the unconfigured state. If you issue the **lsfcportcandidate**, **lsportip**, or **lsibportcandidate** command, the host objects are listed as candidate ports.

Note: This command deletes the associated host throttle if that host is removed.

Remember: This command is unsuccessful if:

- Volume protection is enabled (using the **chsystem** command)
- The host being deleted is mapped to any volume that has received I/O within the defined volume protection time period

If any mappings still exist between this host and volumes, the command fails unless you specify the **-force** parameter. When the **-force** parameter is specified, the **rmhost** command deletes the mappings before the host object is deleted.

An invocation example

```
rmhost host_one
```

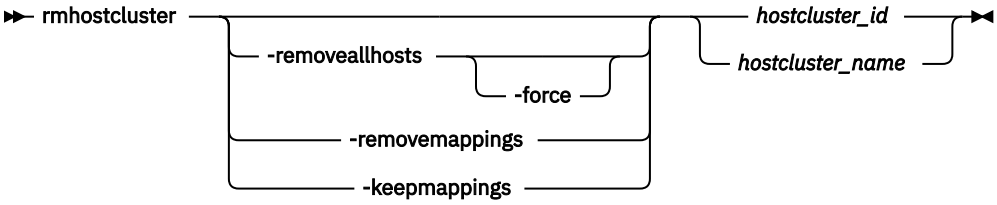
The resulting output:

```
No feedback
```

rmhostcluster

Use the **rmhostcluster** command to remove a host cluster.

Syntax



Parameters

-removeallhosts
(Optional) Specifies the deletion of all hosts in the host cluster and the associated host cluster object. If any volume mappings exist to the hosts that will be deleted, the command fails unless you specify the **-force** parameter.

-force
(Optional) Specifies that hosts should be deleted even if mappings still exist to volumes. Any mappings are deleted when the hosts are deleted.

If volume protection is enabled and the time period is not expired, the command fails even if the **-force** parameter is used.

-removemappings
(Optional) Specifies that all shared volume mappings in the host cluster are removed when the host cluster is removed. The mappings are deleted before the host cluster is deleted.

-keepmappings
(Optional) Specifies that hosts retain the host cluster shared volume mappings (which become private mappings) when the host cluster is removed.

hostcluster_id | hostcluster_name
(Required) Specifies the host cluster to be removed. The value for the ID must be a number and the value for the name must be an alphanumeric string.

Description

This command removes a host cluster.

Note: This command deletes the associated host cluster throttle if that host cluster is removed.

The **-removeallhosts**, **-keepmappings**, and **-removemappings** parameters are mutually exclusive.

An invocation example that removes host cluster hostcluster0 and any related hosts

```
rmhostcluster -removeallhosts hostcluster0
```

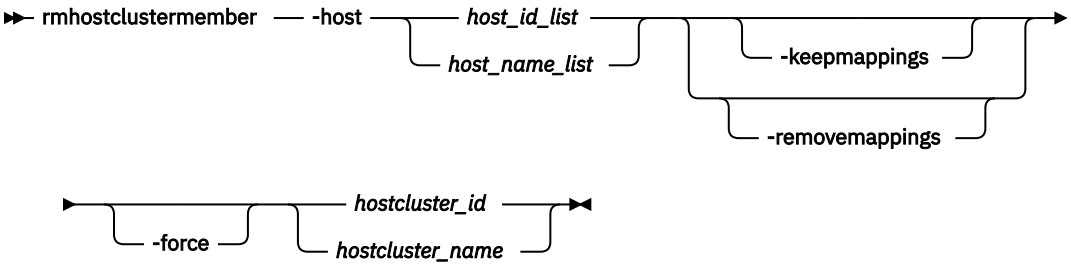
The detailed resulting output:

```
No feedback
```

rmhostclustermember

Use the **rmhostclustermember** command to remove a host from a host cluster object.

Syntax



Parameters

-host ***host_id_list | host_name_list***
(Required) Specifies (by ID or name) the colon-separated host list to remove from the host cluster.

-keepmappings
(Optional) Specifies that the host that is removed from the host cluster retains the host cluster's shared volume mappings. The **-keepmappings** and **-removemappings** parameters are mutually exclusive.

-removemappings
(Optional) Specifies that the host that is removed from the host cluster does not retain the host cluster's shared volume mappings. The **-keepmappings** and **-removemappings** parameters are mutually exclusive.

-force
(Optional) Forces a removal. You must specify this parameter when you remove the last host from a host cluster.

hostcluster_id | hostcluster_name
(Required) Specifies (by ID or name) that the host cluster that the host is removed from. The value for the ID must be a number and the value for the name must be an alphanumeric string.

Description

This command removes a host from a host cluster object.

An invocation example that removes host 0 from host cluster 0 (and also removes the host mappings)

```
rmhostclustermember -host 0 -removemappings 0
```

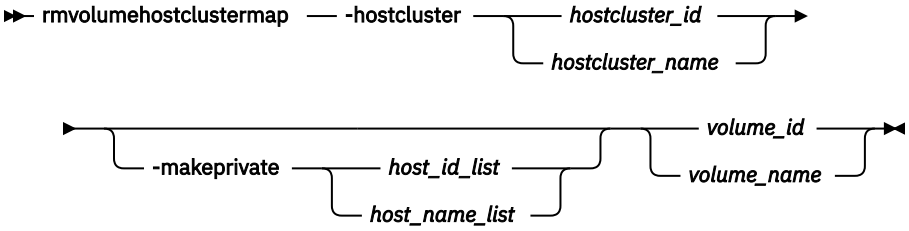
The detailed resulting output:

```
No feedback.
```

rmvolumehostclustermap

Use the **rmvolumehostclustermap** command to remove an existing host cluster mapping to a volume on a clustered system.

Syntax



Parameters

-hostcluster *hostcluster_id | hostcluster_name*
(Required) Specifies the host cluster (by ID or name) to remove from the volume mapping. The value for the ID must be a number and the value for the name must be an alphanumeric string.

-makeprivate *host_id_list | host_name_list*
(Optional) Specifies the host or hosts that acquire private mappings from the volume that is being removed from the host cluster. The value for the ID must be a number and the value for the name must be an alphanumeric string.

volume_id | volume_name
(Required) Specifies the volume by ID or name. The value for the ID must be a number and the value for the name must be an alphanumeric string.

Description

This command removes an existing host cluster mapping on a host cluster. The volume is then inaccessible for input or output (I/O) transactions from the specified host cluster.

An invocation example that moves a mapping from host cluster 0 to volume 0

```
rmvolumehostclustermap -hostcluster 0 0
```

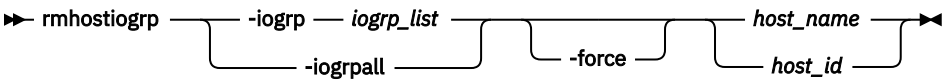
The resulting output:

```
No feedback
```

rmhostiogrp

Use the **rmhostiogrp** command to delete mappings between one or more input/output (I/O) groups and a specified host object.

Syntax



Parameters

-iogrp *iogrp_list*
(Required) Specifies a set of one or more I/O group mappings that will be deleted from the host. You cannot use this parameter with the **iogrpall** parameter.

-iogrpall
(Optional) Specifies that all the I/O group mappings that are associated with the specified host must be deleted from the host. You cannot use this parameter with the **iogrp** parameter.

-force
(Optional) Specifies that you want the system to remove the specified I/O group mappings on the host even if the removal of a host to I/O group mapping results in the loss of host mappings.

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

host_id | host_name
(Required) Specifies the identity of the host either by ID or name from which the I/O group mappings must be deleted.

Description

The **rmhostiogrp** command deletes the mappings between the list of I/O groups and the specified host object.

Remember: This command is unsuccessful if:

- Volume protection is enabled (using the **chsystem** command)
- The host I/O group being removed is mapped to any volume that has received I/O within the defined volume protection time period

If a host is defined in two I/O groups, and has access to a volume through both I/O groups, an attempt to remove the host from just one of those I/O groups fails, even with **-force** specified. To resolve this problem, do one of the following:

- Delete the host mappings that are causing the error
- Delete the volumes or the host

Note: When all I/O groups are removed from an Internet Small Computer System Interface (iSCSI) host, and you want to add an iSCSI port to the host, refer to the **addhostport** and **chhost** commands.

Note: Hosts associated with a partition are by default in all I/O groups and cannot have I/O groups added or removed. The maximum number of I/O groups is hardware type dependent.

An invocation example

```
rmhostiogrp -iogrp 1:2 host0
```

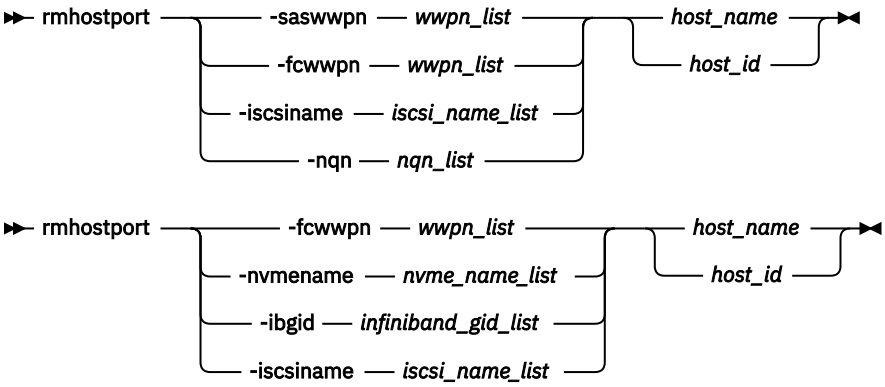
The resulting output:

```
No feedback
```

rmhostport

Use the **rmhostport** command to delete worldwide port names (WWPNs) Internet Small Computer System Interface (iSCSI) names, or NVMe Qualified Names (NQNs) from an existing host object.Use the **rmhostport** command to delete worldwide port names (WWPNs), Internet Small Computer System Interface (iSCSI) names, NVMe Qualified Names (NQNs), or InfiniBand port names from an existing host object.

Syntax



Parameters

-saswwpn *wwpn_list*
(Required if you do not use **-iscsiname**, **-fcwwpn**, or **-nqn**.) Specifies a list of Serial Attached SCSI (SAS) WWPNs with a 16-character hexadecimal string.

-fcwwpn *wwpn_list*
(Required if you do not use **-iscsiname**, **-saswwpn**, or **-nqn**.) Specifies a list of Fibre Channel (FC) WWPNs with a 16-character hexadecimal string.

-iscsiname *iscsi_name_list*
(Required if you do not use **-fcwwpn**, **-saswwpn**, or **-nqn**.) Specifies the comma-separated list of iSCSI names. At least one WWPN or iSCSI name must be specified. You cannot use this parameter with the **-fcwwpn** or **-saswwpn** parameter.

-nqn *nqn_list*
(Required if you do not use **-fcwwpn**, **-saswwpn**, or **-iscsiname**.) Specifies the comma-separated list of NVMe Qualified Names (NQNs). The host protocol must be **nvme** to use this parameter.

host_name | host_id
(Required) Specifies the host name or the host ID.

-fcwwpn *wwpn_list*
(Required if you do not use **-ibgid** or **-nvmename**.) Specifies a list of Fibre Channel (FC) WWPNs with a 16-character hexadecimal string. You cannot use this parameter with **-ibgid** or **-nvmename**.

-ibgid *infiniband_gid_list*
(Required if you do not use **-fcwwpn** or **-nvmename**.) Specifies a list of globally unique identifiers (GUIDs) with a 32-character hexadecimal string. You cannot use this parameter with **-fcwwpn** or **-nvmename**.

-nvmename *nvme_name_list*
(Required if you do not use **-ibgid** or **-fcwwpn**.) Specifies a list of a list of NVMe-qualified names that are defined for a NVMe-attached host. You cannot use this parameter with **-ibgid** or **-fcwwpn**.

host_id | host_name
(Required) Specifies the host object to remove ports from, either by ID or by name.

Description

This command deletes the list of host-bus adapter (HBA) WWPNs, iSCSI names, or NQNs 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.

This command deletes the list of host-bus adapter (HBA) WWPNs, iSCSI port names, NVMe Qualified Names (NQNs), or InfiniBand port names from the specified host object. If the WWPN ports are still logged in to the fabric, they become unconfigured and are listed as candidate WWPNs.

Any volumes that are mapped to this host object are automatically unmapped from the ports.

Remember: This command is unsuccessful if:

- Volume protection is enabled (using the **chsystem** command)
- The last host port being deleted is mapped to any volume that has received I/O within the defined volume protection time period

If multiple hosts are mapped to the same active volume, the host port removal is allowed if the host is offline. This allows for the removal of ports from hosts that might be part of the same system.

An invocation example

```
rmhostport -saswwpn 210100E08B251DD4 host1
```

The resulting output:

```
No feedback
```

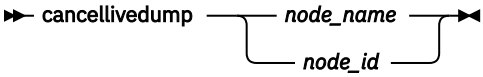
Chapter 22. Livedump commands

Use the **livedump** commands to manage the node livedumps.

cancellivedump

Use the **cancellivedump** command to cancel a live dump.

Syntax



Parameters

node_name|node_id
(Required) Identifies the node name or ID.

Description

Use this command if you issue a **preplivedump** command, but then decide not to issue a **triggerlivedump** command. This releases the resources you allocated for the livedump. This event is recorded in the node trace (.trc) file. For this command to succeed, the node must be in a livedump `prepared` state.

An invocation example

```
cancellivedump node1
```

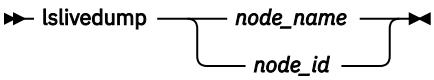
The resulting output:

```
No feedback
```

lslivedump

Use the **lslivedump** command to query the live dump state of a node.

Syntax



Parameters

node_name|node_id
(Required) Identifies the node name or ID.

Description

You can issue this command repeatedly to determine if a live dump is in progress for the node. This table provides the possible values that are applicable to the attributes that are displayed as data in the output views.

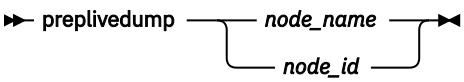
© Copyright IBM Corp. 2003, 2025

Table 113. **lslivedump** outputs	
Attribute	**Description**
inactive	The node has no live dump activity.
prepared	The node is ready to be triggered.
dumping	The node is writing the dump file.

preplivedump

Use the **preplivedump** command to reserve the system resources that are required for livedump.

Syntax



Parameters

node_name|node_id
(Required) Identifies the node name or ID.

Description

You can prepare more than one node for livedump at a time by issuing the **preplivedump** command consecutively. However, you can only trigger one livedump at a time, with an automatic lag time of 30 seconds between each trigger event. This helps maintain node stability.

You can issue multiple **preplivedump** commands on the same node; however, only a **preplivedump** command followed by a **triggerlivedump** command results in output.

Because the livedump resource allocation can take time to execute, you can issue this command to prepare the livedump but trigger it at a later time. This command times out after 60 seconds. The preplivedump event is located in the node trace (.trc) file.

An invocation example

```
preplivedump node1
```

The resulting output:

```
No feedback
```

triggerlivedump

Use the **triggerlivedump** command to capture the metadata that you want to dump, and write the dump file to the internal disk on the node.

Syntax



Parameters

node_name|node_id
(Required) Identifies the node name or ID.

Description

You can issue this command to trigger a **livedump** command. Only one **triggerlivedump** action can be in progress at one time, with an automatic lag time of 30 seconds between each trigger event. The node must have a live dump state of `prepared` for this command to succeed. Output is recorded in the node trace (.trc) file.

After you issue the **triggerlivedump** command, the command captures data and returns you to the CLI interface so that you can issue more commands. While you issue more commands, the live dump disk file is written to the disk in the background, and the live dump state shows as `dumping`. After the write is complete, the state shows as `inactive`.

An invocation example

```
triggerlivedump node1
```

The resulting output:

```
No feedback
```

Chapter 23. Managed disk commands

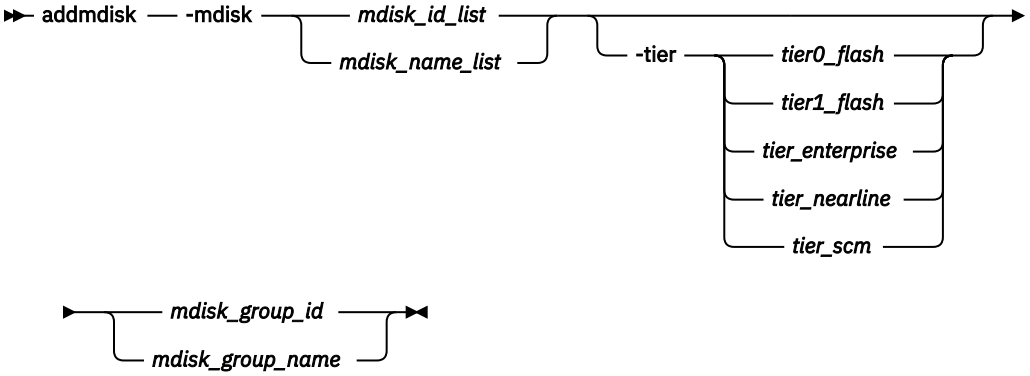
Use the managed disk commands to work with managed disk options on a system.

If the system detects an MDisk, it automatically adds it to the list of known MDisks. If you delete the array that corresponds to the MDisk, the system only deletes the MDisk from the list if the MDisk is offline and it has a mode of unmanaged (it does not belong to a storage pool).

addmdisk

Use the **addmdisk** command to add one or more managed disks to an existing storage pool..

Syntax



Parameters

-mdisk ***mdisk_id_list | mdisk_name_list***
(Required) Specifies one or more managed disk IDs or names to add to the storage pool.

-tier ***tier0_flash | tier1_flash | tier_enterprise | tier_nearline | tier_scm***
(Optional) Specifies the tier of the MDisk or MDisks being added. Unless otherwise specified, the current tier value that is associated with the MDisk is retained. The values are:

tier0_flash
Specifies a `tier0_flash` hard disk drive or an external MDisk for the newly discovered or external volume.

tier1_flash
Specifies an `tier1_flash` (or flash drive) hard disk drive or an external MDisk for the newly discovered or external volume.

tier_enterprise
Specifies a `tier_enterprise` hard disk drive or an external MDisk for the newly discovered or external volume.

tier_nearline
Specifies a `tier_nearline` hard disk drive or an external MDisk for the newly discovered or external volume.

tier_scm
Specifies a `tier_scm` hard disk drive or an external MDisk for the newly discovered or external volume.

The default value for a newly discovered unmanaged MDisk is *enterprise*. You can change this value by using the **chmdisk** command.

© Copyright IBM Corp. 2003, 2025

The tier of external managed disks is not detected automatically and is set to `enterprise`. If the external managed disk is made up of flash drives or nearline Serial Attached SCSI (SAS) drives and you want to use Easy Tier, you must either specify the tier when adding the managed disk to the storage pool or use the **chmdisk** command.

mdisk_group_id | mdisk_group_name
(Required) Specifies the ID or name of the storage pool to add the disks to. When a MDisk is added, the warning threshold for the storage pool is automatically scaled.

Description

This command adds the managed disks that you specify to the storage pool.

If you are adding an MDisk to a pool that uses a provisioning policy, the capabilities of the MDisk must match the capacity savings method that the provisioning policy defines. For example, an MDisk that uses self-compressing drive technology cannot be added to a data reduction pool that defines a provisioning policy that uses compression and deduplication.

Remember: This command cannot be used for child pools.

The disks can be specified in terms of the managed disk ID or the managed disk name. The managed disks must be in unmanaged mode.

Disks that already belong to a storage pool cannot be added to another storage pool until they have been deleted from their current storage pool. You can delete a managed disk from a storage pool under the following circumstances:

- If the managed disk does not contain any extents in use by a volume.
- If you can first migrate the extents in use onto other free extents within the storage pool.

Remember: Do not include a MDisk in a storage pool if it can be only used in image mode.

An invocation example

```
addmdisk -mdisk mdisk13:mdisk14 -tier tier_nearline Group0
```

The resulting output:

```
No feedback
```

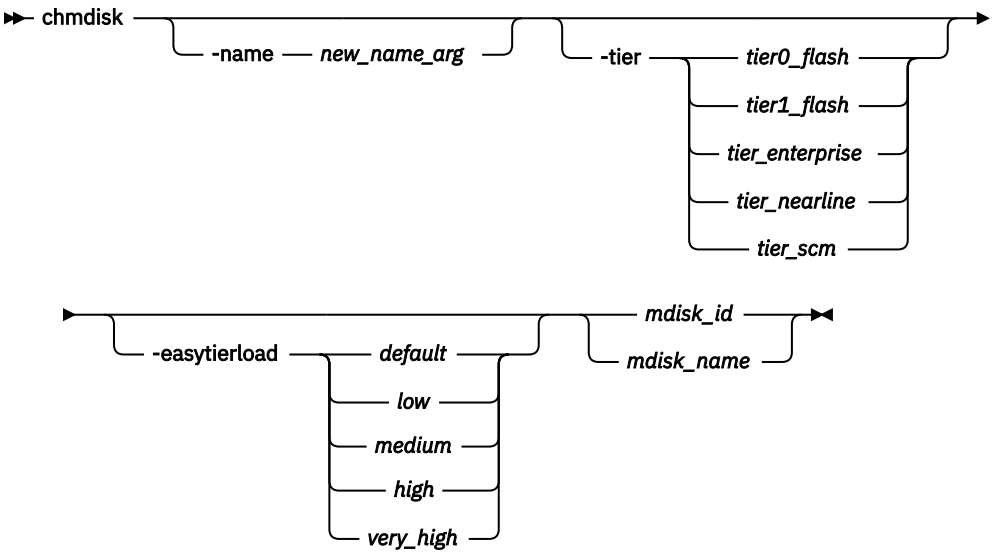
applymdisksoftware (Discontinued)

Attention: The **applymdisksoftware** command has been discontinued. Use the **applydrivesoftware** command to update drives.

chmdisk

Use the **chmdisk** command to modify the name or IBM Easy Tier settings for a managed disk (MDisk).

Syntax



Parameters

-name *new_name_arg*
(Optional) Specifies the new name to be applied to the managed disk.

-tier *tier0_flash | tier1_flash | tier_enterprise | tier_nearline | tier_scm*
(Optional) Specifies the new tier of the MDisk. The values are:

tier0_flash
Specifies a `tier0_flash` hard disk drive or an external MDisk for the newly discovered or external volume.

tier1_flash
Specifies an `tier1_flash` (or flash drive) hard disk drive or an external MDisk for the newly discovered or external volume.

tier_enterprise
Specifies a `tier_enterprise` hard disk drive or an external MDisk for the newly discovered or external volume.

tier_nearline
Specifies a `tier_nearline` hard disk drive or an external MDisk for the newly discovered or external volume.

tier_scm
Specifies a `tier_scm` hard disk drive or an external MDisk for the newly discovered or external volume.

Note: Since Release 8.1.0, the tier of MDisks mapped from certain types of IBM System Storage Enterprise Flash is fixed to `tier0_flash`, and cannot be changed.

-easytierload *default | low | medium | high | very_high*
(Optional) Specifies the Easy Tier load (amount) to place on a non-array MDisk within its tier.

If Easy Tier is either overusing or under-utilizing a particular MDisk, modify the `easy_tier_load` value to change the load size.

Note: Specifying *default* returns the performance capability to the value used by the system. Specify *very_high* only if the MDisk `tier` is `ssd`.

mdisk_id | mdisk_name
(Required) Specifies the ID or name of the managed disk to modify.

Description

This command modifies the attributes of a managed disk.

If you are upgrading your system and the back end of the system uses encrypted storage, you must indicate which MDisks are self-encrypting before you add MDisks to a storage pool. If those MDisks are part of a storage pool, the system assumes that the back-end is not self-encrypting (even if it might be).

If you create encrypted storage pools, the system encrypts locally before it sends data to the back-end. So, the back-end of the system could encrypt again and cannot compress data because the data is random and not compressible.

Note: You must upgrade the system first.

To use encryption on the system that already has encryption that is enabled on the back-end, upgrade the back-end of the system before you enable encryption on the system.

An invocation example

```
chmdisk -tier tier0_flash mdisk13
```

The resulting output:

```
No feedback
```

detectmdisk

Use the **detectmdisk** command to manually rescan the Internet Small Computer Systems Interface (iSCSI) or Fibre Channel (FC) network for any new managed disks (MDisks) that might have been added, and to rebalance MDisk access across all available controller device ports.

Syntax



Parameters

-scope ***scope_id***
(Optional) Specifies the domain index. The value must be a number from `0` to `7`. For example, the value `0` indicates FC and `6` indicates iSCSI and a value of `7` indicates cloud.

Description

This command causes the system to rescan the FC or iSCSI network. The rescan discovers any new MDisks that have been added to the system and rebalances MDisk access across the available controller device ports. This command also detects any loss of controller port availability, and updates the system 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 show if any discovery operations are running.

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

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

When back-end controllers are added to the FC SAN and are included in the same switch zone as a system, the system automatically discovers the back-end controller and determines what storage is presented to it. The SCSI LUs that are presented by the back-end controller are displayed as unmanaged MDisks. However, if the configuration of the back-end controller is modified after this has occurred, the system might be unaware of these configuration changes. Run this command to rescan the FC or iSCSI network and update the list of unmanaged MDisks.

Note: The automatic discovery that is performed by the system does not write to an unmanaged MDisk. Only when you add an MDisk to a storage pool, or use an MDisk to create an image mode volume, is the storage actually used.

To identify the available MDisks, issue the **detectmdisk** command to scan the FC or iSCSI network for any MDisks. When the detection is complete, issue the **lsmdiskcandidate** command to show the unmanaged MDisks; these MDisks have not been assigned to a storage pool. Alternatively, you can issue the **lsmdisk** command to view all of the MDisks.

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

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

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

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

You must issue the **detectmdisk** command to force the system to update its configuration and accept the changes to the controller ports.

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

An invocation example

```
detectmdisk
```

The resulting output:

```
No feedback
```

dumpallmdiskbadblocks

Use the **dumpallmdiskbadblocks** command to dump bad block counts to a dump file used by the fix procedures and the satask **snap** command.

Syntax



Parameters

None

Description

Use the **dumpallmdiskbadblocks** command to dump bad block counts to a readable ASCII dump file for use by fix procedures and the satask **snap** command. The output contains bad blocks for which an error log has been raised.

Use **lsdumps -prefix** /dumps/mdisk to list the output files. Use **cleardumps -prefix** /dumps/mdisk to clear the output files.

The maximum number of dump files is 20.

An invocation example

```
dumpallmdiskbadblocks
```

The resulting output if MDisk 2 and MDisk 5 have bad blocks:

```
Cluster name: my_cluster
Timestamp of dump: Fri Oct 31 11:27:33 2009 UTC

Mdisk id: 2
Mdisk name: mdisk2
Number of bad blocks: 4

Mdisk id: 5
Mdisk name: mdisk 5
Number of bad blocks: 1

Total mdisks with bad blocks: 2
Total number of bad blocks: 5
```

The resulting output if the MDisks have no bad blocks

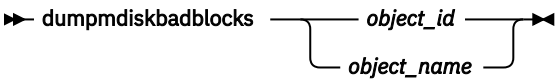
```
Cluster name: my_cluster
Timestamp of dump: Fri Oct 31 11:27:33 2009 UTC

Total mdisks with bad blocks: 0
Total number of bad blocks: 0
```

dumpmdiskbadblocks

Use the **dumpmdiskbadblocks** command to write the bad block counts and locations that are on a specified MDisk to a dump file for use by fix procedures.

Syntax



Parameters

object_id* | *object_name
(Required) Specifies the MDisk for which you need to dump the bad block record table.

Description

Use the **dumpmdiskbadblocks** command to write the bad block counts and locations that are on a specified MDisk to a readable ASCII dump file for use by fix procedures. The output consists of bad blocks for which an event log has been raised.

Use **lsdumps -prefix** /dumps/mdisk to list the output files. Use **cleardumps -prefix** /dumps/mdisk to clear the output files.

The reported event log sequence numbers correspond to the first event seen in the bad block record, which is a 512-block region.

- If there are multiple event logs in the same region, the earliest event sequence is used.
- If there are event logs of different types in the same region, event sequence numbers for bad blocks caused by medium errors on RAID member drives take precedence.
- If a range of bad blocks runs across record boundaries, the sequence number corresponding to the last record is used.

The maximum number of dump files is 20.

An invocation example

```
dumpmdiskbadblocks 3
```

The resulting output if the MDisk has bad blocks:

```
Cluster name: my_cluster
Timestamp of dump: Fri Oct 31 11:27:33 2017 UTC

Mdisk id: 3
Mdisk name: mdisk3
Number of bad blocks: 6

Start LBA: 0x1234123412341234
Length: 2
Event log sequence number: 1

Start LBA: 0x5678568102341234
Length: 4
Event log sequence number: 2
```

The resulting output if the MDisk has no bad blocks:

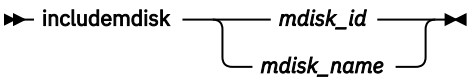
```
Cluster name: my_cluster
Timestamp of dump: Fri Oct 31 11:27:33 2017 UTC

Mdisk id: 3
Mdisk name: mdisk3
Number of bad blocks: 0
```

includemdisk

Use the **includemdisk** command to include a disk that has been excluded by the system.

Syntax



Parameters

mdisk_id | mdisk_name
(Required) Specifies the ID or name of the managed disk to add back into the system.

Description

The specified managed disk is included in the system.

You might exclude a disk from the system because of multiple I/O failures. These failures might be caused by noisy (or unstable) links. Once a fabric-related problem has been fixed, the excluded disk can be added back into the system.

Running this command against an MDisk might change its state, whether the state is reported as excluded.

Note: If an MDisk is in the excluded state, is offline, and does not belong to a storage pool, issuing an include command for this MDisk results in the MDisk record being deleted from the system.

An invocation example

```
includemdisk mdisk5
```

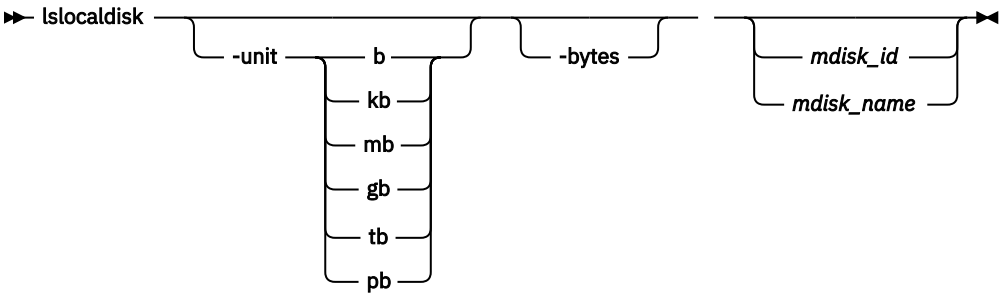
The resulting output:

```
No feedback
```

lslocaldisk

Use the **lslocaldisk** command to display the cloud managed disks (MDisks) visible to the system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-unit b | kb | mb | gb | tb | pb
(Optional) Specifies the data units for the **-filtervalue** parameter.

Note: -unit must be used with **-filtervalue**.

-bytes
(Optional) Specifies that you want the report to display all capacities as bytes. Capacity values that are displayed in units other than bytes might be rounded. When you filter on capacity, use a unit of bytes, **-unit b**, for exact filtering.

mdisk_id | mdisk_name
(Optional) Specifies the name or ID of an MDisk. When you use this parameter, the detailed view of the specific MDisk is returned and any value that is specified by the **-filtervalue** parameter is ignored. If you do not specify the *mdisk_id | mdisk_name* parameter, the concise view displays all objects that match the filtering requirements that are specified by the **-filtervalue** parameter.

Description

This command returns a concise list or a detailed view of MDisks visible to the system. This table provides the potential output.

*Table 114. **lslocaldisk** output*

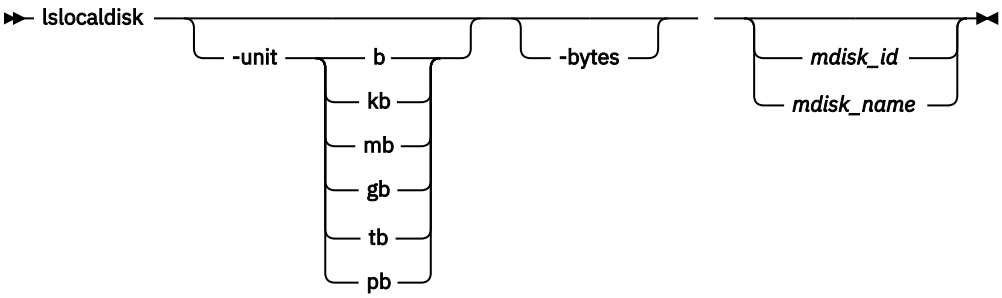
Attribute	Values
mdisk_id	Displays the local MDisk ID.
mdisk_name	Displays the local MDisk name.
status	Displays the local MDisk status (online, offline, or degraded).
mode	Displays the local MDisk mode (ummanaged, managed, or image).
mdisk_grp_id	Displays the ID of the MDisk group that the local disk belongs to.
mdisk_grp_name	Displays the name of the MDisk group that the local disk belongs to.
capacity	Displays the capacity of the disk.
encrypt	Displays the encryption status of the disk. The value is either yes or no.
fabric_type	Displays the type of local MDisk fabric type. This value is the same as lsmdisk fabric_type. The values are: • Fibre Channel (fc) • sas • cloud • iscsi
disk_id	Displays the disk ID from the cloud provider, if any.
type	Displays the local MDisk type. The values are: • gp2 • io1 • st1 • sc1 • standard • standardSSD • premium
node_id	Displays the node ID that the local MDisk is attached.
node_name	Displays the node name that the local MDisk is attached.

Table 114. **lslocaldisk** output (continued)	
Attribute	**Values**
state	Displays the state of the local MDisk. The values are: • available • in-use • attaching • attached • detaching • detached • creating • deleting • deleted • error
iops	Displays the number of I/O operations per second (IOPS) that the local MDisk supports. For Provisioned IOPS SSD volumes, this value represents the number of IOPS that are provisioned for the volume. For General Purpose SSD volumes, this value represents the baseline performance of the volume and the rate at which the volume accumulates I/O credits for bursting.
zone	Displays the local MDisk zone.

lslocaldisk

Use the **lslocaldisk** command to display the cloud managed disks (MDisks) visible to the system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-unit b | kb | mb | gb | tb | pb
(Optional) Specifies the data units for the **-filtervalue** parameter.

Note: -unit must be used with **-filtervalue**.

-bytes
(Optional) Specifies that you want the report to display all capacities as bytes. Capacity values that are displayed in units other than bytes might be rounded. When you filter on capacity, use a unit of bytes, **-unit b**, for exact filtering.

mdisk_id | mdisk_name
(Optional) Specifies the name or ID of an MDisk. When you use this parameter, the detailed view of the specific MDisk is returned and any value that is specified by the **-filtervalue** parameter is ignored. If you do not specify the *mdisk_id | mdisk_name* parameter, the concise view displays all objects that match the filtering requirements that are specified by the **-filtervalue** parameter.

Description

This command returns a concise list or a detailed view of MDisks visible to the system. This table provides the potential output.

*Table 115. **lslocaldisk** output*

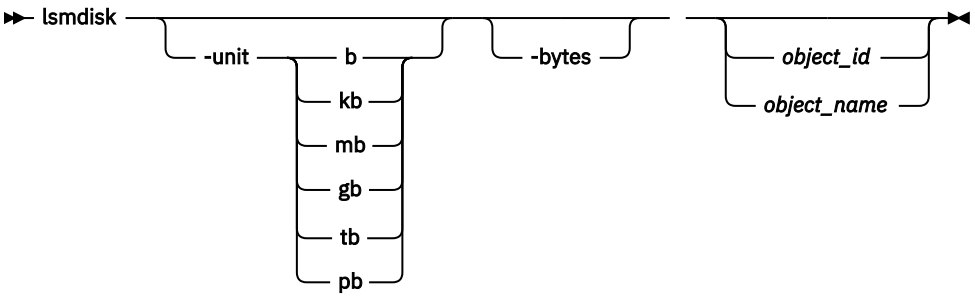
Attribute	Values
mdisk_id	Displays the local MDisk ID.
mdisk_name	Displays the local MDisk name.
status	Displays the local MDisk status (online, offline, or degraded).
mode	Displays the local MDisk mode (ummanaged, managed, or image).
mdisk_grp_id	Displays the ID of the MDisk group that the local disk belongs to.
mdisk_grp_name	Displays the name of the MDisk group that the local disk belongs to.
capacity	Displays the capacity of the disk.
encrypt	Displays the encryption status of the disk. The value is either yes or no.
fabric_type	Displays the type of local MDisk fabric type. This value is the same as lsmdisk fabric_type. The values are: • Fibre Channel (fc) • sas • cloud • iscsi
disk_id	Displays the disk ID from the cloud provider, if any.
type	Displays the local MDisk type. The values are: • gp2 • io1 • st1 • sc1 • standard • standardSSD • premium
node_id	Displays the node ID that the local MDisk is attached.

Table 115. **lslocaldisk** output (continued)	
Attribute	**Values**
node_name	Displays the node name that the local MDisk is attached.
state	Displays the state of the local MDisk. The values are: • available • in-use • attaching • attached • detaching • detached • creating • deleting • deleted • error
iops	Displays the number of I/O operations per second (IOPS) that the local MDisk supports. For Provisioned IOPS SSD volumes, this value represents the number of IOPS that are provisioned for the volume. For General Purpose SSD volumes, this value represents the baseline performance of the volume and the rate at which the volume accumulates I/O credits for bursting.
zone	Displays the local MDisk zone.

lsmdisk

Use the **lsmdisk** command to display a concise list or a detailed view of managed disks (MDisks) visible to the system. It can also list detailed information about a single MDisk.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-unit b | kb | mb | gb | tb | pb
(Optional) Specifies the data units for the **-filtervalue** parameter.

Note: -unit must be used with **-filtervalue**.

-bytes
(Optional) Specifies that you want the report to display all capacities as bytes. Capacity values that are displayed in units other than bytes might be rounded. When you filter on capacity, use a unit of bytes, **-unit b**, for exact filtering.

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 that match the filtering requirements that are specified by the **-filtervalue** parameter.

Description

This command returns a concise list or a detailed view of MDisks visible to the system. This table provides the potential output for MDisks.

Note: Some of the attributes may not be applicable to your system.

Table 116. MDisk output

Attribute	Values
status	• online • offline • excluded • degraded_paths • degraded_ports • degraded (This value applies only to internal MDisks.)
mode	unmanaged, managed, image, array
quorum_index	0, 1, 2, or blank if the MDisk is not being used as a quorum disk.
block_size	512, 524 bytes in each block of storage
ctrl_type	4, 6, where 6 is a flash drive that is attached inside a node and 4 is any other device.
tier	The tier this MDisk is assigned to by auto-detection (for internal MDisks) or by the user: • tier0_flash • tier1_flash • tier_enterprise • tier_nearline • tier_scm **Note:** You can change this value by using the **chmdisk** command.

Table 116. MDisk output (continued)	
Attribute	**Values**
easy_tier_load	This value controls Easy Tier settings, and is either blank (for arrays) or one of the following values (for MDisks): • `low` • `medium` • `high` • `very_high`
raid_status	**offline** Array is offline on all nodes. **degraded** Array has deconfigured or offline members; the array is not fully redundant. **syncing** Array members are all online. The array is syncing parity or mirrors to achieve redundancy. **initting** Array members are all online. The array is initializing; the array is fully redundant. **online** Array members are all online and the array is fully redundant. **expanding** The array members are all online and the array is expanding.
raid_level	The RAID level of the array (`RAID1`, `RAID5`, `RAID6`, `RAID10`).
redundancy	The number of how many member disks that fail before the array fails.
strip_size	The strip size of the array (in KB).
spare_goal (Deprecated)	The value of this attribute is always `blank`.
spare_protection_min (Deprecated)	The value of this attribute is always `blank`.
balanced	Describes if the array is balanced to its spare goals: **exact** All populated members have exact capability match, exact location match. **yes** All populated members have at least exact capability match, exact chain, or different enclosure or slot. **no** Anything else that is not included for `yes` or `exact`.
site_id	Indicates the site value for the MDisk. This numeric value is 1, 2, 3, or blank.
site_name	Indicates the site name for the MDisk. This is an alphanumeric value or is blank.

Table 116. MDisk output (continued)	
Attribute	**Values**
fabric_type	Indicates the type of MDisk. The values are: • `fc` indicates that it is an MDisk from an FC controller. • `sas_direct` indicates that it is an MDisk from an SAS direct-attached controller. • `iscsi` indicates that it is an iSCSI controller.
distributed (Deprecated)	Indicates whether the array is distributed. The value of this attribute is always `yes`.
drive_class_id	Indicates the drive class that makes up this array. If **`-allowsuperior`** was used during array creation, the lowest used drive class ID is displayed. This value is blank for traditional arrays.
drive_count	Indicates the total width of the array, including rebuild areas. The value is a number in the range 4 - 128. The minimum value for RAID-6 and RAID-10 arrays is 6.
stripe_width	Indicates the width of a single unit of redundancy within a distributed set of drives. The values are: • Any number in the range 3 - 16 for RAID-5 arrays. • Any number in the range 4 - 16 for RAID-6 arrays. • An even number in the range 2 - 16 for RAID-10 arrays.
rebuild_areas_total	Indicates the total number of rebuild areas set at array creation time. These rebuild areas provide performance but no capacity. The value is a number in the range 1 - 4 for distributed array RAID-5 and RAID-6, and the value is a number in the range 2 - 4 for distributed array RAID-10 (the value is blank for traditional arrays).
rebuild_areas_available	Indicates the number of remaining build areas within the set of arrays. The value is a number in the range 1 - 4 for distributed array RAID-5 and RAID-6, and the value is a number in the range 2 - 4 for distributed array RAID-10 (the value is blank for traditional arrays).
rebuild_areas_goal	Indicates the rebuild areas threshold (minimum limit) at which point the array logs an error. The value is a number in the range 1 - 4 for distributed array RAID-5 and RAID-6, and a number in the range 2 - 4 for distributed array RAID-10 (the value is blank for traditional arrays).
dedupe	Indicates that dedupe is enabled. If dedupe is enabled, duplicate copies of repeating data are compressed or removed.
ctrl_WWNN	Indicates the control worldwide node name (WWNN).
preferred_WWPN	Indicates the preferred worldwide port name (WWPN).
active_WWPN	Indicates the active WWPN.
preferred_iscsi_port_id	Indicates the preferred I/O port identifier, which has the same value as the `preferred_WWPN` value in the Fibre Channel (FC) domain. The Internet Small Computer System Interface (iSCSI) port ID value is displayed, but the value is blank for non-iSCSI domains. This value must be a numeric value that can range in the range 0 - 1023.

Table 116. MDisk output (continued)	
Attribute	**Values**
active_iscsi_port_id	Indicates the active I/O port identifier, which has the same value as the `active_WWPN` value in the FC domain. The Internet Small Computer System Interface (iSCSI) port ID value is displayed, but the value is blank for non-iSCSI domains. This value must be a numeric value that can range in the range 0 - 1023.
over_provisioned	Indicates whether the MDisk is thin-provisioned. The value is no if the MDisk is marked as standard-provisioned. resource-provisioned, the information cannot be verified at the backend. The value is `yes` or no.
supports_unmap	Indicates whether the mdisk is provided by a controller that indicates that it supports unmapping. The value is no if the MDisk indicates that it does not support unmapping or this information cannot be verified at the backend. The value is `yes` or no.
provisioning_group_id	Indicates the allocated identifier for the provisioning group affiliated with the MDisk. The identifier lists the MDisks that are contained in the same provisioning group. The value must be an integer (number).
physical_capacity	Indicates the total physical storage capacity of the provisioning group that contains this MDisk. If this MDisk is not over-provisioned the logical capacity is displayed here. The value must be a number (indicated in units) that is rounded to two decimal places.
physical_free_capacity	Indicates the amount of formatted available physical space in the provisioning group that contains this MDisk. If this MDisk is not over-provisioned the remaining logical capacity is displayed instead. The value must be a number (indicated in units) that is rounded to two decimal places.
write_protected	Shows `yes` if the MDisk is in the write protected state, otherwise no.
effective_used_capacity	The amount of logical data that is written to the MDisk. If this MDisk is not `over_provisioned`, then this field is blank.
allocated_capacity	The amount of logical capacity that is allocated (by a pool) from this MDisk for volume data and quorum.
fips_enabled	Indicates whether the mdisk follows FIPS standards. The value is `yes` or no.

Note: The automatic discovery that is performed by the system does not write anything to an unmanaged MDisk. It is only when you add an MDisk to a storage pool, or use an MDisk to create an image mode volume, that the system uses the storage.

To see which MDisks are available, issue the **`detectmdisk`** command to manually rescan the Fibre Channel or iSCSI network for any new MDisks. Issue the **`lsmdiskcandidate`** command to show the unmanaged MDisks. These MDisks are not assigned to a storage pool.

Notes:

1. A system connection from a node or node canister port to a storage controller port for a single MDisk is a path. The Mdisk *path_count* value is the number of paths currently being used to submit input/output (I/O) to this MDisk.
2. The MDisk *max_path_count* value is the highest value *path_count* reaches since the MDisk was last fully online.
3. The *preferred_WWPN* is one of the World Wide Port Names (WWPNs) the storage controller specifies as a preferred WWPN. If the controller has nothing that is specified, this is a blank field.
4. The *active_WWPN* indicates the WWPN of the storage controller port currently being used for I/O.
   a. If no storage controller ports are available for I/O, this is a blank field.

b. If multiple controller ports are actively being used for I/O, this field's value is `many`.

The following define the status fields:

online
The MDisk is online and available.

degraded
(Internal MDisks only) The array has members that are `degraded`, or the `raid_status` is `degraded`.

degraded_ports
There are one or more MDisk port errors.Errors are detected on the Fibre Channel (FC) network that connects the storage. The ports that are causing the errors are excluded from the configuration.

degraded_paths
One or more paths to the MDisk are lost; the MDisk is not online to every node in the system.One or more nodes in the configuration lost contact with the specified MDisk or array because of an FC connectivity issue (or because of the LUN mapping configuration).

offline
All paths to the MDisk are lost.

excluded
The MDisk is excluded from use by the system; the MDisk port error count exceeded the threshold.

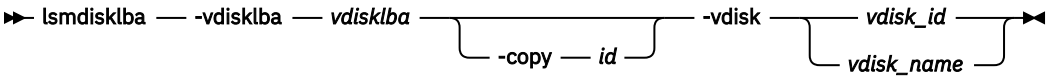
lsmdiskdumps (Deprecated)

The **lsmdiskdumps** command is deprecated. Use the **lsdumps** command to display a list of files in a particular dumps directory.

lsmdisklba

Use the **lsmdisklba** command to list the MDisk and logical block address (LBA) for the specified volume LBA.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-vdisklba *vdisklba*
(Required) Specifies the 64-bit hexadecimal logical block address (LBA) on the volume. The LBA must be specified in hex, with a 0x prefix.

-copy *id*
(Optional) Specifies the volume copy ID to list the MDisk and LBA for. If this parameter is not specified, the command lists MDisks and LBAs for all volume copies.

vdisk_id | vdisk_name
(Required) Specifies the volume name or ID.

Description

The **lsmdisklba** command returns the logical block address (LBA) of the MDisk that is associated with the volume LBA. For mirrored volume, the command lists the MDisk LBA for both the primary and the copy.

If applicable, the command also lists the range of LBAs on both the volume and MDisk that are mapped in the same extent, or for thin-provisioned disks, in the same grain. If a thin-provisioned volume is offline and the specified LBA is not allocated, the command displays the volume LBA range only.

The mdisk_lba field provides the corresponding LBA on the real capacity for the input LBA. For compressed volume copies it is empty, and the system displays only the range of physical LBAs where the compressed input LBA is located.

This table summarizes the data that can be returned with this command.

*Table 117. **lsmdisklba** command output*

Field	standard-provisioned, single copy volume	LBA not allocated on thin-provisioned volume	Mirrored volume with one normal copy and one offline thin-provisioned copy	
			Normal copy	Thin-provisioned copy
copy_id	yes	yes	yes	yes
mdisk_id	yes	no	yes	no
mdisk_name	yes	no	yes	no
type	allocated	unallocated	allocated	offline
mdisk_lba	yes	no	yes	no
mdisk_start	yes	no	yes	no
mdisk_end	yes	no	yes	no
vdisk_start	yes	yes	yes	yes
vdisk_end	yes	yes	yes	yes

lsmdiskcandidate

Use the **lsmdiskcandidate** command to list all unmanaged MDisks by MDisk ID.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

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 occurs, 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.

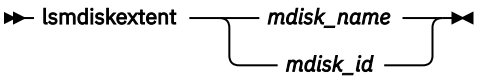
Note: The automatic detection that is performed by the cluster does not write anything to an unmanaged MDisk. It is only when you instruct the cluster to add an MDisk to a storage pool or use a MDisk to create an image mode volume that the storage is used.

Check to see which MDisks are available by issuing the **detectmdisk** command to manually scan the Fibre Channel network for any MDisks. Issue the **lsmdiskcandidate** command to show the unmanaged MDisks. These MDisks are not assigned to a storage pool. Alternatively, you can issue the **lsmdisk** command to view all of the MDisks.

lsmdiskextent

Use the **lsmdiskextent** command to display the extent allocation between managed disks and volumes. The output lists a volume ID, volume copy ID, and the number of extents.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

mdisk_name | mdisk_id
(Required) Specifies the specific object ID or name of the specified type.

Description

The command displays a list, in which each entry contains a volume ID, volume copy ID, and the number of extents. These volume copies are using extents on the specified MDisk. The number of extents that are being used on each MDisk is also shown.

Note: You cannot specify this command for MDisks that are in data reduction pools. It means that for:

- Thin-provisioned or compressed volumes, the number extents that are shown is not accurate.
- standard-provisioned volumes, the number of extents that are shown is accurate.

A thin-provisioned or compressed volume in a data reduction pool cannot display how many extents are on an MDisk that is in a data reduction pool.

Every volume copy is constructed from one or more MDisks. At times, you might have to determine the relationship between the two objects.

To determine the relationship between volume copies and MDisks, issue the following command for each volume copy:

```
lsvdiskmember vdisk_name | vdisk_id
```

where *vdisk_name* | *vdisk_id* is the name or ID of the volume copy. It displays a list of IDs that correspond to the MDisks that make up the volume copy.

To determine the relationship between volume copies and MDisks and the number of extents that are provided by each MDisk, you must use the command-line interface. For each volume copy, issue the following command:

```
lsvdiskextent vdisk_name | vdisk_id
```

where *vdisk_name* | *vdisk_id* is the name or ID of the volume copy. It displays a table of MDisk IDs and the corresponding number of extents that each MDisk is providing as storage for the specified volume copy.

To determine the relationship between MDisks and volume copies, issue the following command for each MDisk:

```
lsmdiskmember mdisk_name | mdisk_id
```

where *mdisk_name* | *mdisk_id* is the name or ID of the MDisk. It displays a list of IDs that correspond to the volume copies that are using this MDisk.

To determine the relationship between MDisks and volume copies and the number of extents that are used by each volume copy, you must use the command-line interface. For each MDisk, issue the following command:

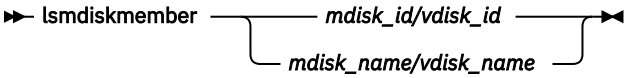
```
lsmdiskextent mdisk_name | mdisk_id
```

where *mdisk_name* | *mdisk_id* is the name or ID of the MDisk. This command displays a table of volume copy IDs and the corresponding number of extents that are being used by each volume copy. In the output, *number_of_extents* displays either a number (for standard-provisioned volumes in data reduction pools or volumes in regular pools) or a 1 (for thin-provisioned/compressed volumes in data reduction pools).

lsmdiskmember

Use the **lsmdiskmember** command to display a list of volumes that use extents on the specified MDisk. That is, the volumes use extents on the managed disk that are specified by the MDisk ID.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

mdisk_id/vdisk_id | mdisk_name/vdisk_name
(Required) Specifies the ID or name of the MDisk or volume copy for which you want a list of volumes that use extents of that MDisk.

Description

This command displays a list of volumes that use extents on the managed disk that are specified by the ID. The list displays members of the respective object and is independent of the state of the individual members. That is, if they are in offline state, they are still displayed.

- A volume in a data reduction storage pool cannot display how many members are on an MDisk that is in a data reduction pool.
- If the MDisk specified is in a data reduction pool, output includes all thin-provisioned and compressed volumes in the pool.
- standard-provisioned volume copies in data reduction pools are displayed correctly.

Every volume is constructed from one or more MDisks. To determine the relationship between volume copies and MDisks, issue the following command:

`lsvdiskmember` *vdisk_id | vdisk_name*

where *vdisk_id | vdisk_name* is the name or ID of the volume copy. This action displays a list of IDs that correspond to the MDisks that make up the volume copy.

To determine the relationship between volume copies and MDisks and the number of extents that are provided by each MDisk, you must use the command-line interface. For each volume copy, issue the following command:

`lsvdiskextent` *vdisk_id | vdisk_name*

where *vdisk_id | vdisk_name* is the name or ID of the volume copy. This command displays a table of MDisk IDs and the corresponding number of extents that each MDisk provides as storage for the volume copy.

To determine the relationship between MDisks and volume copies, issue the following command:

`lsmdiskmember` *mdisk_id | mdisk_name*

where *mdisk_id | mdisk_name* is the name or ID of the MDisk. This command displays a list of IDs that correspond to the volume copies that are using this MDisk.

To determine the relationship between MDisks and volume copies and the number of extents that are used by each volume copy, you must use the command-line interface. For each MDisk *mdisk_id | mdisk_name*, issue the following command:

`lsmdiskextent` *mdisk_id | mdisk_name*

where *mdisk_id | mdisk_name* is the name or ID of the MDisk. This command displays a table of volume copy IDs and the corresponding number of extents that are being used by each volume copy.

setquorum (Deprecated)

The **setquorum** command is deprecated. Use the **chquorum** command to change the quorum association.

triggermdiskdump (Discontinued)

Attention: The **triggermdiskdump** command is discontinued. Use the **triggerdrivedump** command to collect support data from a disk drive.

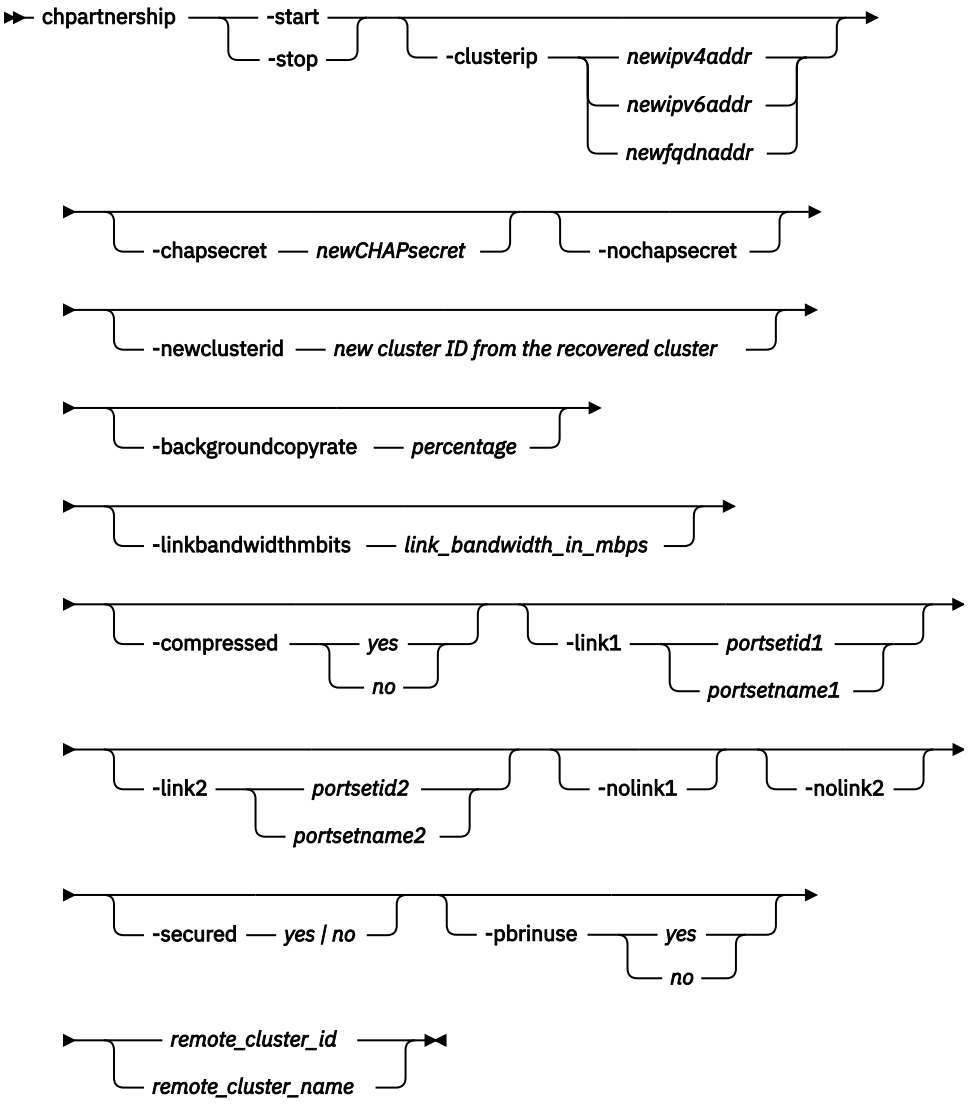
Chapter 24. Copy Service commands

Use the Copy Service commands to work with the Metro Mirror, Global Mirror, and HyperSwap® services that are provided by the system.

chpartnership

Use the **chpartnership** command to update properties that are related to a partnership between the local system and the remote system.

Syntax



Parameters

-start | -stop
(Optional) Starts or stops a partnership. To start or stop a partnership, run the **chpartnership** command from either system.

© Copyright IBM Corp. 2003, 2025

-clusterip *newipv4addr | newipv6addr | newfqdnaddr*
(Optional) Specifies the new partner system IP address, which can be IPv4, IPv6 or FQDNs. Systems that are connected over IP links are not displayed by **lspartnershipcandidate** before executing **mkippartnership**. This does not apply to FC-based connections.

Specify this parameter when creating partnerships with systems connected over native IP links. To change the partner system IP address, first specify **chpartnership -stop** to stop the partnership.

-chapsecret *newCHAPsecret*
(Optional) Specifies the new Challenge-Handshake Authentication Protocol (CHAP) secret of the partner system. The maximum size of the CHAP secret is eighty alphanumeric characters.

-nochapsecret
(Optional) Resets the CHAP secret used to authenticate with the partner system. Specify **chpartnership -stop** to stop the partnership. Reset the CHAP secret of the partner system when authentication of discovery requests is turned off on the partner system .

-newclusterid *new cluster ID from the recovered cluster*
(Optional) Specifies the new cluster ID for a partnered system. Use this parameter on the local system to recover the partnership to a remote system after completing a system recovery, also known as Tier 3 (T3) recovery, on the remote system.

Note: This parameter is mutually exclusive with all other parameters. Use this parameter only under the direction of IBM Support Center.

-backgroundcopyrate *percentage*
(Optional) Deprecated. This value is ignored by the system.

-linkbandwidthmbits *link_bandwidth_in_mbps*
(Optional) Specifies the aggregate bandwidth of the replication link between two clustered systems (systems) in megabits per second (Mbps). It is a numeric value from 1 to 100000, specified in Mbps. Asynchronous replication will aim to not exceed this value when replicating to the partnered systems. If multiple I/O groups exist in the partnered system then the system will expect to have this amount of bandwidth to each I/O group in the remote system. High availability will consume as much bandwidth as needed to maintain two copies of the storage partition and is not limited by this value.

Note: when using IP compression, on hardware platforms that do not have hardware compression capabilities, **-linkbandwidthmbits** *link_bandwidth_in_mbps* has limitation up to 1120 megabits per second (Mbps)

This parameter can be specified without stopping the partnership. This command is mutually-exclusive with all parameters other than **-backgroundcopyrate**.

-compressed *yes | no*
(Optional) Specifies whether compression is enabled for this partnership. The default value is no.

The following parameters are mutually exclusive with IP partnerships:

-link1 *portsetid1 | portsetname1*
(Optional) Specifies the portset id or name to be used for WAN link 1.

-link2 *portsetid2 | portsetname2*
(Optional) Specifies the portset id or name to be used for WAN link 2.

-nolink1
(Optional) Specifies the removal of the portset id or name specified for WAN link 1.

-nolink2
(Optional) Specifies the removal of the portset id or name specified for WAN link 2.

Note: A partnership is required to have at lease one active link between systems. You must have a portset defined in either -link1 or -link2 parameters for the systems. You cannot delete both portsets from -link1 and -link2.

-secured *yes | no*
(Optional) Specifies whether to use encryption to secure the connection between systems in the IP partnership. If **-secured** parameter is not specified, the connection between systems in the IP partnership will not be secured. If **-secured** parameter is specified, the value yes or no is mandatory.

-pbrinuse *yes | no*
(Optional) Specifies whether policy-based replication will be used on the partnership. The value can be yes or no.

remote_cluster_ID | remote_cluster_name
(Required) Specifies the remote system ID or name of a partnership. The specified value must match one of the system IDs or names returned after issuing **lspartnershipcandidate**. The specified value must match one of the system IDs or names listed by **lspartnership**. Specifying a remote system ID or name with **chpartnership** does not affect the remote system. To change the system name, specify **chsystem**

Description

The **chpartnership** command modifies the properties of a partnership. Properties that can be changed include:

- The bandwidth of the partnership (specified using the **-linkbandwidthmbits** option). This affects the bandwidth available replication between the local system and the remote system specified in the partnership. To modify the background copy bandwidth from the remote system to local system, issue the **chpartnership** command a second time for the remote system.
- Whether the partnership is started or stopped (specified using the **-start** and **-stop** options). This starts or stops replication between the local system and the remote system specified in the partnership. System partnership must be in either the `partially_configured_stopped` or `fully_configured_stopped` states to be started.
- The CHAP secret (specified using the **-chapsecret** option) or the system IP (specified using the **-clusterip** option) for partnerships created using IP links. Before changing the partner CHAP secret or system IP, stop the partnership; and after making the necessary changes, start the partnership.

Note: The local and remote systems in an IP partnership must use the same IP address types, IPv4 or IPv6.

An invocation example to configure a new CHAP secret for a partner.

```
chpartnership -stop remote-system-2
chpartnership -chapsecret newpassword remote-system-2
chpartnership -start remote-system-2
```

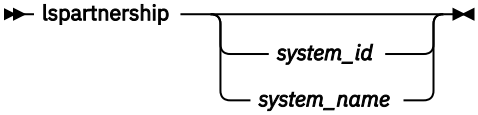
The resulting output:

```
No feedback
```

lspartnership

Use the **lspartnership** command to display a concise or detailed view of the current clustered systems (systems) that are associated with the local system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

system_id **|** ***system_name***
(Optional) Specifies the name or ID of a system. This parameter displays the detailed view of the specified system.

Description

This table described attribute values.

Table 118. ***lspartnership*** *attribute values*

Attribute	Value
id	Indicates the system ID.
name	Indicates the system name.
location	Indicates the system location.
code_level	Indicates the code level.

*Table 118. **lspartnership** attribute values (continued)*	
Attribute	**Value**
partnership	Indicates the current state of the partnership; not applicable for the local system and is blank. The partnership field can show the following values: **fully_configured** The **mkfcpartnership** or **mkippartnership** command is issued in both directions and the remote system is online and available. **partially_configured_local** The **mkfcpartnership** or **mkippartnership** command is issued from the local system to the remote system. The remote system is online and available for partnership. **partially_configured_local_stopped** The **mkfcpartnership** or **mkippartnership** command is issued from the local system to the remote system. The **chpartnership** command with the stop parameter is issued from the local system, and the remote system is online and available. Issue **chpartnership -start** on the local system, and **mkfcpartnership** or **mkippartnership** on the remote system **not_present** The **mkfcpartnership** or **mkippartnership** command is issued from the local system to the remote system, and the remote system is not available. The remote system is either offline or not connected to the local system. **fully_configured_stopped** The **mkfcpartnership** or **mkippartnership** command is issued in both directions and the remote system is online and available. The **chpartnership** command with the stop parameter is issued from the local system. **fully_configured_remote_stopped** The **mkfcpartnership** or **mkippartnership** command is issued in both directions and the remote system is online and available. The **chpartnership** command with the stop parameter is issued from the remote system. **fully_configured_local_excluded** The **mkfcpartnership** or **mkippartnership** command is issued in both directions. The local system excludes the connection to the remote system. **fully_configured_remote_excluded** The **mkfcpartnership** or **mkippartnership** command is issued in both directions. The local system excludes the connection to the remote system. **fully_configured_exceeded** There are too many systems in the system network and the partnership from the local system to the remote is disabled. **mid_recovery** This state means that the partnership is left in a partially completed state.  **Attention:** Contact IBM Support Center, if this field indicates the mid_recovery.

*Table 118. **lspartnership** attribute values (continued)*	
Attribute	**Value**
relationship_bandwidth_limit (Deprecated)	
type	Indicates the type of partnership: • Fibre Channel (FC) • Internet Protocol Version 4 (IPv4) or Internet Protocol Version 6 (IPv6) • DNS FC partnerships are created connecting two systems over FC fabrics. IPv4 or IPv6 partnerships are created connecting two systems over native IP links.
cluster_ip	Indicates the partner system IP address, which can be IPv4, IPv6 or FQDNs. This information is displayed for IP-based partnerships. For IP-based partnership this field displays the system IP address that is specified while the partnership using **mkippartnership** is created.
chap_secret	Indicates the Challenge-Handshake Access Protocol (CHAP) secret (up to 80 alphanumeric characters) for the partner system. The CHAP authenticates the local system with the partner system during discovery, and Internet Small Computer System Interface (iSCSI) system session creation. For FC-based relationships this field is always blank.
link_bandwidth_mbits	Indicates the aggregate bandwidth for the replication link in megabits per second (Mbps). This is a numeric value from 0 to 100000. If there are multiple links between the local and remote systems, this paramepater is set to the sum of the link bandwidths for these links.
background_copy_rate	Indicates the bandwidth allocation for background copy operations that are performed over the replication link. It is expressed as a percentage of the link bandwidth value, and is the maximum rate at which background copy operations are performed. This value is a number from 0 to 100 .
event_log_sequence	Indicates the last sequence number (indicating the last event) from event log for this partnership. This is a numeric value in the range 100 - 8000000. For FC-based relationships this field is always blank.
max_replication_delay (Deprecated)	
compressed	Indicates whether compression is enabled. The value can be yes or no (default).
link1	Indicates the local portset name used for WAN Link 1. The value can be string with up to 63 characters.
link2	Indicates the local portset name used for WAN Link 2. The value can be string with up to 63 characters.
link1_ip_id	Indicates the IP address index which is currently in use for IP Partnership for WAN Link1 from **lsip** view. The value can be numeric from 0 - 2047.
link2_ip_id	Indicates the IP address index which is currently in use for IP Partnership for WAN Link 2 from **lsip** view. The value can be numeric from 0 - 2047.
secured	(Optional) Specifies whether to use an encrypted data path for secured IP replication. Default value is no as unsecured IP Partnership is created.
pbr_in_use	Indicates whether policy-based replication is used on the partnership. The value can be yes or no.

lspartnershipcandidate

Use the **lspartnershipcandidate** command to list the clustered systems available for setting up a partnership with the local system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

Description

This table provides the attribute values that can be displayed as output view data.

Table 119. ***lspartnershipcandidate*** *output*

Attribute	Description
id	Indicates the system ID.
configured	Indicates whether system is configured.
system_name	Indicates the system name.
IP_address_4	Indicates the IPv4 address of the system.
IP_address_6	Indicates the IPv6 address of the system.

Description

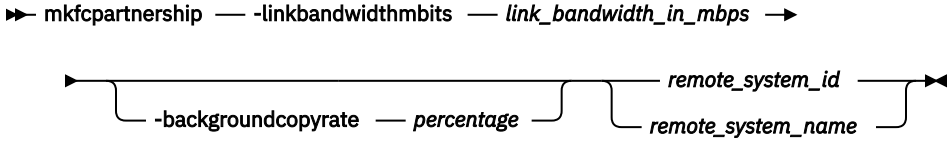
This command displays a list of systems that are available as candidate partner systems to form a partnership between two systems.

Output from the command shows the system ID, name, and configured status of the remote candidate system. The remote candidate system forms a partnership with the local system when you use the **mkippartnership** or **mkfcpartnership** command. The remote system shows the partnership status as `partially_configured_local_stopped` or `partially_configured_local` when you use the **lssystem** command. The **lspartnershipcandidate** command displays the configured status of those remote systems that form a partnership with the local system.

mkfcpartnership

Use the **mkfcpartnership** command to define partnerships using Fibre Channel (FC).

Syntax



Parameters

-linkbandwidthmbits *link_bandwidth_in_mbps*
(Required) Specifies the aggregate bandwidth of the replication link between two clustered systems (systems in megabits per second (Mbps). It is a numeric value from 1 to 100000.

-backgroundcopyrate *percentage*

(Optional) Deprecated. This value is ignored by the system.

remote_system_id | remote_system_name
(Required) Specifies the remote system ID or name of a partnership. The specified value must match one of the system IDs or names returned after issuing **lspartnershipcandidate**.

Description

This command creates a partnership between the local system and the remote system using Fibre Channel. A partnership is required to enable replication, high availability or migration between the systems.

An invocation example

```
mkfcpartnership –linkbandwidthmbits 10000 remote-system-2
```

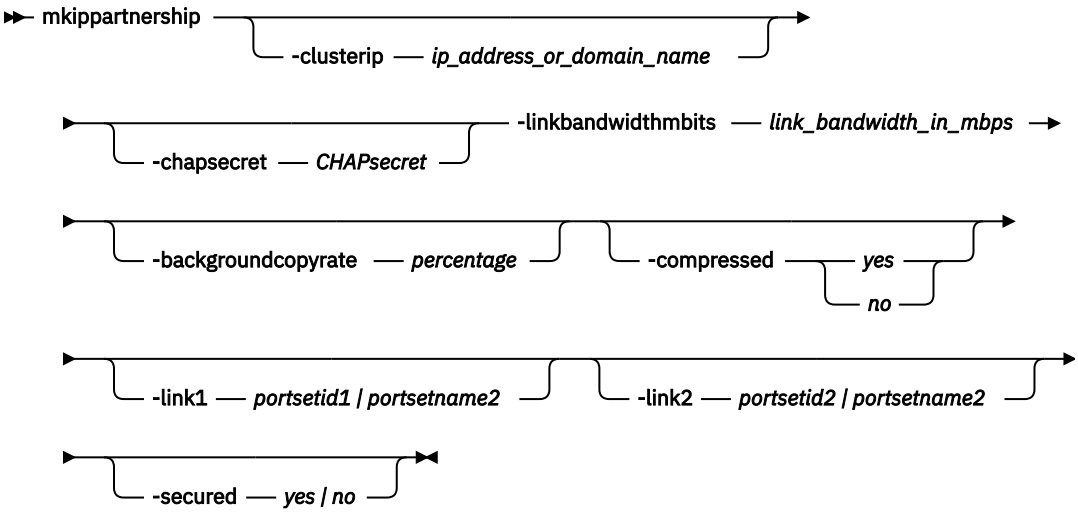
The resulting output:

```
No feedback
```

mkippartnership

Use the **mkippartnership** command to define a new partnership created over Internet Protocol (IP) links.

Syntax



Parameters

-clusterip *ip_address_or_domain_name*
(Required) Specifies the partner system IP address, which can be IPv4, IPv6 or FQDNs. Systems that are connected over IP links are not displayed by **lspartnershipcandidate** before running **mkippartnership**.

-chapsecret *CHAPsecret*
(Optional) Specifies the Challenge-Handshake Authentication Protocol (CHAP) secret of the partner system. The maximum size of the CHAP secret is eighty alphanumeric characters.

-linkbandwidthmbits *link_bandwidth_in_mbps*
(Required) Specifies the aggregate bandwidth of the replication link between two clustered systems (systems) in megabits per second (Mbps). It is a numeric value from 1 to 100000.

Important: For partnerships over IP links with compression, this parameter specifies the aggregate bandwidth before the data has been compressed. Do not set this parameter higher than the physical link bandwidth multiplied by the (carefully rounded down) compression factor.

This parameter can be specified without stopping the partnership.

-backgroundcopyrate *percentage*
(Optional) Deprecated. This setting is not used by the system.

-compressed *yes | no*
(Optional) Specifies whether compression is enabled for this partnership. The default value is no.

-link1 *portsetid1 | portsetname1*
(Optional) Specifies the portset ID or name to be used for WAN link 1.

-link2 *portsetid2 | portsetname2*
(Optional) Specifies the portset ID or name to be used for WAN link 2.

Note: Even though link1 and link2 parameters are optional, it is required to provide at least one of these links to establish an IP partnership.

Note: There are two options to specify up to 2 portset IDs or names indicating the local Ethernet ports to be used by the partnership. The system supports up to 2 WAN links per partnership. These 2 portsets indicate the WAN links to which local ports are connected. All the ports that are part of the portset must be connected to or belong to the same WAN link. The values specified for the link1 and link2 parameters indicate the portset ID or name that specifies the group of IP addresses on the local system that can be used to access intersite WAN links.

-secured *yes | no*
(Optional) Specifies whether to use encryption to secure the connection between systems in the IP partnership. If **-secured** parameter is not specified, the connection between systems in the IP partnership will not be secured. If **-secured** parameter is specified, the value yes or no is mandatory.

Description

This command defines a new partnership that is created over Internet Protocol (IP) links. A remote system IP must be specified so its IP ports are enabled for data replication.

All TCP connections are established that uses either IPv4 or IPv6, and it cannot be a mix of the two IP address types.

An invocation example

```
mkippartnership -clusterip xxx.xxx.xx.xx –chapsecret mychapsecret -linkbandwidthmbits 100
-link1 5 -link2 portset6
```

The resulting output:

```
No feedback
```

An invocation example

```
mkippartnership –clusterip xxxx::xxx:xxxx:xxxx:xxxx –chapsecret mychapsecret –
linkbandwidthmbits 1024 -link1 5 -link2 portset6
```

The resulting output:

```
No feedback
```

An invocation example

The following example shows an IP partnership that uses compression:

```
mkippartnership -clusterip xxx.xxx.xx.xx –chapsecret mychapsecret -linkbandwidthmbits 100 -link1 5 -link2
portset6
```

The resulting output:

```
No feedback
```

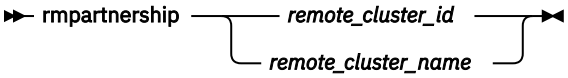
mkpartnership (Discontinued)

The **mkpartnership** command is deprecated. Use either the **mkfcpartnership** or **mkippartnership** command instead.

rmpartnership

Use the **rmpartnership** command to remove a partnership to a remote system.

Syntax



Parameters

remote_cluster_id | remote_cluster_name
(Required) Specifies the system ID or the name of the remote system.

Description

This command removes a partnership from the local system to a remote system. To completely remove the partnership, you must run the command on both the local system and the remote system.

Before you run the **rmpartnership** command, you must remove any replication policies that exist between the two systems.

An invocation example

```
rmpartnership cluster1
```

The resulting output:

```
No feedback
```

Related reference

“ mkreplicationpolicy ” on page 364
Use the **mkreplicationpolicy** command to create a replication policy.

“lsreplicationpolicy” on page 356
Use the **lsreplicationpolicy** command to display a list of all replication policies on the system.

“rmreplicationpolicy” on page 366
Use the **rmreplicationpolicy** command to delete a replication policy.

Chapter 25. Migration commands

Use the migration commands to work with migration options for the system storage system.

lsmigrate

Use the **lsmigrate** command to display the progress of all current data migration operations.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

If you use multiple threads to migrate data, the progress increases when all threads complete the migration of an extent. For large extent sizes with many threads, which can result in large increments in the percentage progress.

Description

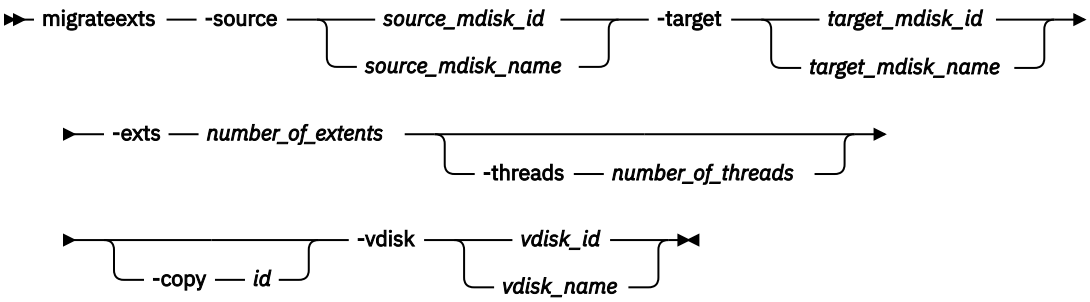
This command displays information of all the migrations that are currently in progress.

Note: Only user-initiated migrations are reported by using this command. Easy Tier migrations are not included in the output.

migrateexts

Use the **migrateexts** command to migrate extents from one managed disk to another.

Syntax



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.

© Copyright IBM Corp. 2003, 2025

-threads *number_of_threads*
(Optional) Specifies the number of threads to use while migrating these extents. You can specify 1 - 4 threads. The default number of threads is 4.

-copy *id*
(Required if the specified volume has more than one copy) Specifies the volume copy that the extents belong to.

-vdisk *vdisk_id | vdisk_name*
(Required) Specifies the volume that the extents belong to.

Note: The command cannot be used with image-mode volumes.

Description

This command migrates a given number of extents from the source volume and the managed disk that contains extents that are used to make up the volume. The target is a managed disk within the same storage pool.

You cannot specify this command for thin-provisioned or compressed volume copies that are in data reduction storage pools.

If a large number of extents are being migrated, you can specify 1 - 4 threads. You can issue the **lsmigrate** command to check the progress of the migration.

The **migrateexts** command fails if there are insufficient free extents on the target managed disk. To avoid this problem, do not issue new commands that use extents until the extents migration is completed.

The **migrateexts** command fails if the target or source volume is offline, or if Easy Tier is active for the volume copy. Correct the offline condition before attempting to migrate the volume.

Note: Migration activity on a single managed disk is limited to a maximum of 4 concurrent operations. This limit does not take into account whether the managed disk is the source or the destination target. If more than four migrations are scheduled for a particular managed disk, further migration operations are queued pending the completion of one of the currently running migrations. If a migration operation is stopped for any reason, a queued migration task can be started. However, if a migration is suspended, the current migration continues to use resources and a pending migration is not started. For example, the following setup is a possible initial configuration:

- `MDiskGrp` 1 has volume 1 created in it
- `MDiskGrp` 2 has volume 2 created in it
- `MDiskGrp` 3 has only one MDisk

With the previous configuration, the following migration operations are started:

- Migration 1 migrates volume 1 from `MDiskGrp` 1 to `MDiskGrp` 3, running with 4 threads.
- Migration 2 migrates volume 2 from `MDiskGrp` 2 to `MDiskGrp` 3, running with 4 threads.

Due to the previous limitations, the two migration operations do not always run at the same speed. `MDiskGrp` 3 has only one MDisk and the two migration operations have a total of 8 threads that are trying to access the one MDisk. Four threads are active. The remaining threads are in standby mode waiting to access the MDisk.

Remember: This command cannot be used if the source MDisk is a SAS MDisk (which works in image mode only).

An invocation example

```
migrateexts -vdisk vdisk4 -source mdisk4 -exts
    64 -target mdisk6 -threads 4
```

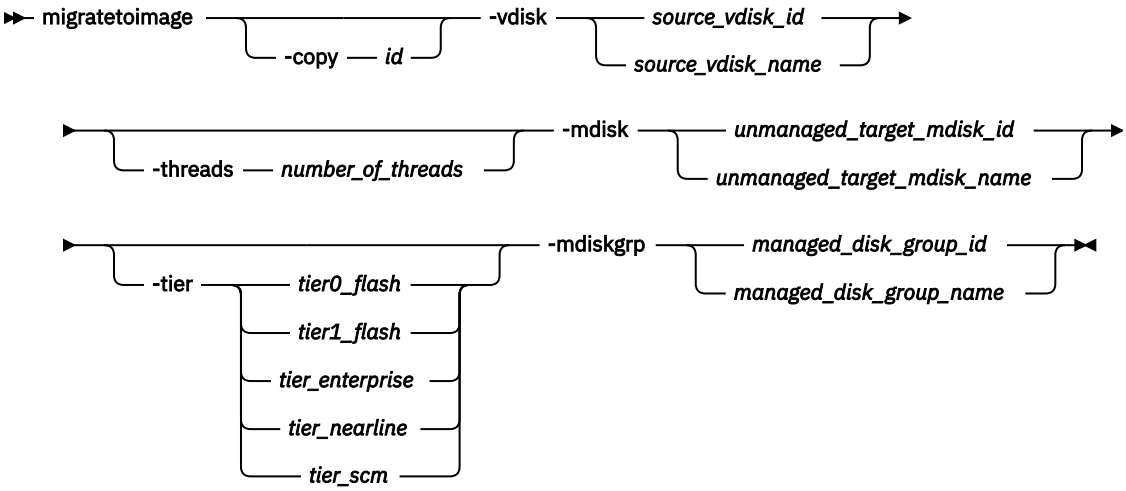
The resulting output:

```
No feedback
```

migratetoimage

Use the **migratetoimage** command to migrate data from a volume (image mode or managed mode) onto a new image mode volume copy. The target disk does not have to be in the same storage pool as the source disk.

Syntax



Parameters

-vdisk ***source_vdisk_id | name***
(Required) Specifies the name or ID of the source volume to be migrated.

-copy ***id***
(Required if the specified volume has more than one copy) Specifies the volume copy to migrate from.

-threads ***number_of_threads***
(Optional) Specifies the number of threads to use during the migration of extents. You can specify 1 - 4 threads. The default number of threads is 4.

-mdisk ***unmanaged_target_mdisk_id | name***
(Required) Specifies the name of the MDisk to which the data must be migrated. This disk must be unmanaged and large enough to contain the data of the disk that is being migrated.

-mdiskgrp ***managed_disk_group_id | name***
(Required) Specifies the storage pool into which the MDisk must be placed, after the migration has completed.

-tier ***tier0_flash | tier1_flash | tier_enterprise | tier_nearline | tier_scm***
(Optional) Specifies the tier of the MDisk being added.

tier0_flash
Specifies a `tier0_flash` hard disk drive or an external MDisk for the newly discovered or external volume.

tier1_flash
Specifies an `tier1_flash` (or flash drive) hard disk drive or an external MDisk for the newly discovered or external volume.

tier_enterprise
Specifies a `tier_enterprise` hard disk drive or an external MDisk for the newly discovered or external volume.

tier_nearline
Specifies a `tier_nearline` hard disk drive or an external MDisk for the newly discovered or external volume.

tier_scm
Specifies a `tier_scm` hard disk drive or an external MDisk for the newly discovered or external volume.

Description

This command cannot be used to if the source volume copy is in a child pool or if the MDisk group that is specified is a child pool. This command does not work if the volume is fast formatting.

Note: You cannot migrate a volume or volume image between storage pools if cloud snapshot feature is enabled on the volume.

The **migratetoimage** command migrates the data of a user-specified volume by consolidating its extents (which might reside on one or more MDisks) onto the extents of the target MDisk that you specify. After migration is complete, the volume is classified as an image type volume, and the corresponding mdisk is classified as an image mode MDisk.

The managed disk that is specified as the target must be in an unmanaged state at the time that the command is run. Running this command results in the inclusion of the MDisk into the user-specified storage pool.

You cannot specify **migratetoimage** if the target or source volume is offline. Correct the offline condition before you migrate the volume.

You cannot specify **migratetoimage** to migrate a thin or compressed volume in a data reduction storage pool to an image mode volume. You must replicate the volume to a standard-provisioned, thin, or compressed volume copy in a standard storage pool before you migrate a thin or compressed volume that is in a data reduction pool from one clustered system to another clustered system.

The new image mode volume copy which is created by the 'migratetoimage' command contains the same capacity savings as the source volume copy. A thin or compressed image mode copy can only be used for importing the volume into another Storage Virtualize system.

To use the new image mode volume copy directly with a host, you must remove any capacity savings from the volume prior to perform the migrate to image mode process.

An invocation example

The following example specifies that the user wants to migrate the data from `vdisk0` onto `mdisk5` and that the MDisk must be put into the storage pool `mdgrp2`.

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

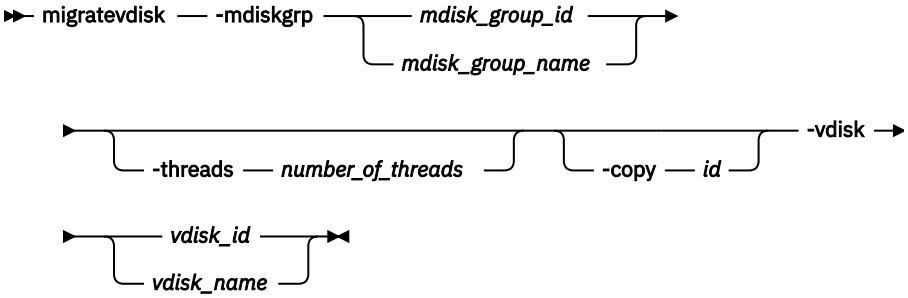
The resulting output:

```
No feedback
```

migratevdisk

Use the **migratevdisk** command to migrate an entire volume from one storage pool to another storage pool.

Syntax



Parameters

-mdiskgrp *mdisk_group_id | mdisk_group_name*
(Required) Specifies the new storage pool ID or name.

-threads *number_of_threads*
(Optional) Specifies the number of threads to use during the migration of these extents. You can specify 1 - 4 threads. The default number of threads is 4.

-copy *id*
(Required if the specified volume has more than one copy) Specifies the volume copy to migrate.

-vdisk *vdisk_id | vdisk_name*
(Required) Specifies the volume ID or name to migrate in to a new storage pool.

Note: The command cannot be used with image-mode volumes.

Description

The **migratevdisk** command migrates the specified volume into a new storage pool; all the extents that make up the volume are migrated onto free extents in the new storage pool.

You can reassign a volume from a:

- Child pool to its parent pool
- Parent pool to one of its child pools
- Between the child pools in the same parent pool
- Between two parent pools

Note: You cannot migrate a volume between storage pools if cloud snapshot is enabled on the volume.

You can issue the **lsmigrate** command to view the progress of the migration.

The process can be prioritized by specifying the number of threads to use during the migration. Using only one thread puts the least background load on the system.

The **migratevdisk** command fails if there are insufficient free extents on the targeted storage pool for the duration of the command. To avoid this problem, do not issue new commands that use extents until the volume migration is completed.

The **migratevdisk** command fails if the target volume or source volume is offline. Correct the offline condition before you attempt to migrate the volume.

Remember: You cannot specify this command:

- If the source MDisk is a SAS MDisk (which works in image mode only).

- If the volume that is being migrated is thin-provisioned or compressed and in a data reduction pool.
- If the target pool is a data reduction pool and the volume that is being migrated is either thin-provisioned or compressed.

For these volume types, you must create a volume copy in the destination pool by using volume mirroring to perform the migration. For more information, see the **addvdiskcopy** or **addvolumecopy** command.

An invocation example

```
migratevdisk -vdisk 4 -mdiskgrp Group0 -threads 2
```

The resulting output:

```
No feedback
```

Chapter 26. Multifactor authentication commands

Use the following set of commands to configure the system-wide IBM Security Verify and Duo Security multifactor authentication settings.

lsauthmultifactorverify

Use the **lsauthmultifactorverify** command to list the system-wide IBM Security Verify multifactor authentication configuration.

Syntax



Parameters

None.

Description

This table provides the attribute values that can be displayed as output view data.

*Table 120. **lsauthmultifactorverify** output*

Attribute	Description
status	Indicates the MFA enablement status using IBM Security Verify. Value can be *enabled* or *disabled*.
failmode	Indicates how the system should behave for the user logins when the MFA server is unavailable. Value can be *insecure* or *secure*. • A secure failmode indicates the failure of MFA-enabled users if the server is unavailable. • An insecure failmode indicates the succeed of MFA-enabled users if the server is unavailable.
hostname	Indicates the hostname of the IBM Security Verify tenant. It must consist of a string up to 255 printable ASCII characters.
port	Indicates the port number of the IBM Security Verify tenant. The value must be an integer 1-65535.
openid_client_id	Indicates the IBM Security Verify OpenID Connect client ID for the system. It must consist of a string up to 64 printable ASCII characters.
openid_client_secret_set	Indicates whether the IBM Security Verify OpenID Connect client secret has been set. Value can be *yes* or *no*.
cli_client_id	Indicates the IBM Security Verify CLI client ID for the system. It must consist of a string up to 64 printable ASCII characters.
cli_client_secret_set	Indicates whether the IBM Security Verify CLI client secret is set or not. Value can be *yes* or *no*.
admin_enrolled_factors	Indicates whether factors enrolled by the IBM Security Verify admin are displayed on the choice list of factors when logging into the command line.

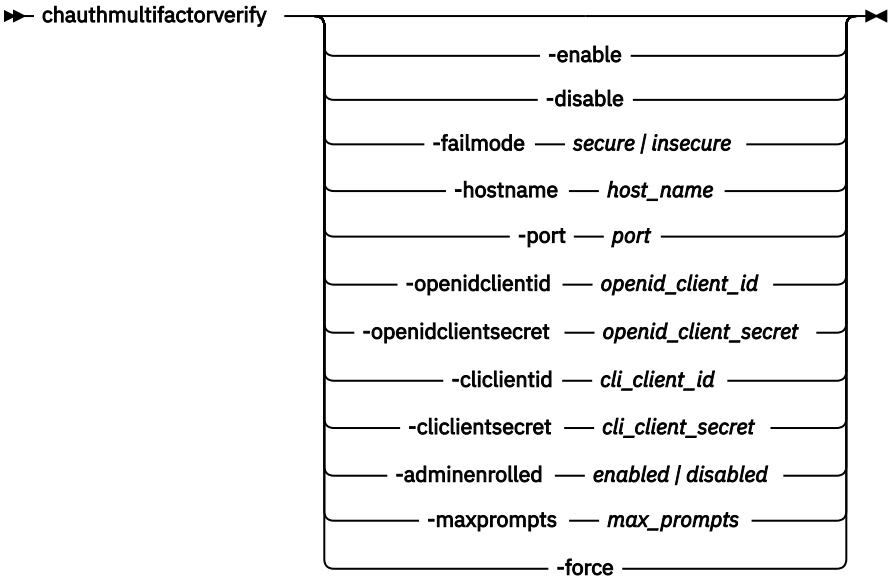
© Copyright IBM Corp. 2003, 2025

Table 120. **lsauthmultifactorverify** output (continued)	
Attribute	**Description**
max_prompts	Indicates the maximum number of prompts displayed when logging in with a second factor on the CLI.

chauthmultifactorverify

Use the **chauthmultifactorverify** command to change the system-wide IBM Security Verify multifactor authentication configuration.

Syntax



Parameters

-enable
(Optional) Specifies the enablement of the multifactor authentication service.

-disable
(Optional) Specifies the disablement of the multifactor authentication service. This parameter is mutually exclusive with `-enable`, `-failmode`, `-hostname`, `-port`, `-openidclientid`, `-openidclientsecret`, `-cliclientid`, `-cliclientsecret`.

-failmode
(Optional) Specifies the system behavior for the user logins when the MFA server is unavailable. Value can be *insecure* or *secure*

- A failmode of *secure* means login to MFA-enabled users will fail if the server is unavailable.
- A failmode of *insecure* means logins to MFA-enabled users will succeed if the server is unavailable.

-hostname ***host_name***
(Optional) Specifies the hostname of the IBM Security Verify tenant. It must consist of a string up to 255 printable ASCII characters.

-port ***port***
(Optional) Specifies the port number of the IBM Security Verify tenant. The value must be an integer 1-65535.

-openidclientid *openid_client_id*
(Optional) Specifies the IBM Security Verify OpenID Connect client ID for the system. It is required to enable the multifactor authentication for login to the Storage Virtualize GUI. It must consist of a string up to 64 printable ASCII characters.

-openidclientsecret *openid_client_secret*
(Optional) Specifies the IBM Security Verify OpenID Connect client secret of the system. It must consist of a string up to 64 printable ASCII characters.

-cliclientid *cli_client_id*
(Optional) Specifies the IBM Security Verify CLI client ID for the system. It is required to enable the multifactor authentication for login to the Storage Virtualize CLI. It must consist of a string up to 64 printable ASCII characters.

-cliclientsecret *cli_client_secret*
(Optional) Specifies the IBM Security Verify CLI client secret key of the system. It must consist of a string up to 64 printable ASCII characters.

-adminenrolled *enabled | disabled*
(Optional) Specifies whether to display factors enrolled by the IBM Security Verify administrator on the choice list of factors when logging into the command line.

-maxprompts *max_prompts*
(Optional) Specifies the maximum number of prompts displayed when attempting to log in with a second factor on the CLI.

-force
(Optional) Specifies that MFA will be removed for superuser and any user groups with MFA enabled, when used with the **-disable** parameter.

An invocation example

```
chauthmultifactorverify -hostname xxxxx.xxxxxx.ibm.com -openidclientid
xxxxxxx-xxxxx-xxxx-xxxxx-xxxxx -openidclientsecret xxxxxx -cliclientid
xxxxxxx-xxxxx-xxxxx-xxx-xxxx -cliclientsecret xxxxxxx -enable
```

The resulting output:

```
No feedback.
```

lsauthmultifactorduo

Use the **lsauthmultifactorduo** command to list the system-wide Duo Security multifactor authentication configuration.

Syntax



Parameters

None.

Description

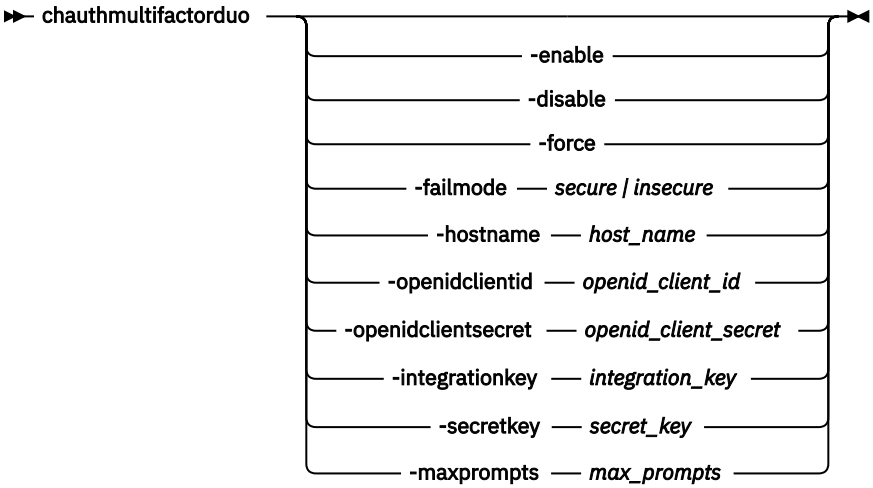
This table provides the attribute values that can be displayed as output view data.

*Table 121. **lsauthmultifactorduo** output*	
Attribute	**Description**
status	Indicates the MFA enablement status using Duo Security. Value can be *enabled* or *disabled*.
failmode	Indicates how the system should behave for the user logins when the MFA server is disabled. Value can be *insecure* or *secure*. • A secure failmode indicates the failure of the MFA-enabled users if the server is unavailable. • An insecure failmode indicates the success of the MFA-enabled users if the server is unavailable.
hostname	Indicates the hostname of the Duo Security tenant. It must consist of a string up to 255 printable ASCII characters.
openid_client_id	Indicates the Duo Security OpenID Connect client ID for the system. It must consist of a string up to 64 printable ASCII characters.
openid_client_secret_set	Indicates whether the Duo Security OpenID Connect client secret has been set. Value can be *yes* or *no*.
integration_key	Indicates the Duo Security integration key for the system. It must consist of a string up to 64 printable ASCII characters.
secret_key_set	Indicates whether the Duo Security secret key for the system is set or not. Value can be *yes* or *no*.
max_prompts	Indicates the maximum number of prompts displayed when logging in with a second factor on the CLI.

chauthmultifactorduo

Use the **chauthmultifactorduo** command to change the system-wide Duo Security multifactor authentication configuration.

Syntax



Note:

- The **-disable** parameter is mutually exclusive with all other parameters.
- The **-force** parameter is mutually exclusive with all other parameters except **-disable** parameter.

Parameters

-enable
(Optional) Specifies the enablement of the multifactor authentication service.

-disable
(Optional) Specifies the disablement of the multifactor authentication service.

-force
(Optional) When used with the **-disable** parameter, removes MFA from all user groups and the superuser profile that currently use MFA.

-failmode
(Optional) Specifies the system behavior for the user logins when the MFA server is unavailable. Value can be *insecure* or *secure*

- A secure failmode indicates the failure of the MFA-enabled users if the server is unavailable.
- An insecure failmode indicates the success of the MFA-enabled users if the server is unavailable.

-hostname ***host_name***
(Optional) Specifies the hostname of the Duo Security tenant. It must consist of a string up to 255 printable ASCII characters.

-openidclientid ***openid_client_id***
(Optional) Specifies the Duo Security OpenID Connect client ID for the system. It is required to enable the multifactor authentication for login to the GUI. It must consist of a string up to 64 printable ASCII characters.

-openidclientsecret ***openid_client_secret***
(Optional) Specifies the Duo Security OpenID Connect client secret of the system. It must consist of a string up to 64 printable ASCII characters.

-integrationkey ***integration_key***
(Optional) Specifies the Duo Security key of the system. It is required to enable the multifactor authentication for login to the CLI. It must consist of a string up to 64 printable ASCII characters.

-secretkey ***secret_key***
(Optional) Specifies the Duo Security secret key of the system. It must consist of a string up to 64 printable ASCII characters.

-maxprompts ***max_prompts***
(Optional) Specifies the maximum number of the prompts displayed when attempting to log in with a second factor on the CLI.

An invocation example

```
chauthmultifactorduo -hostname api-12345678.duosecurity.com -openidclientid
xxxxxxxxxxxxxxx -openidclientsecret xxxxxxxxxxxxxxxxxxxxxxxxxxx
-integrationkey xxxxxxxxxxxx
-secretkey xxxxxxxxxxxx -maxprompts 3 -enable
```

The resulting output:

```
No feedback.
```

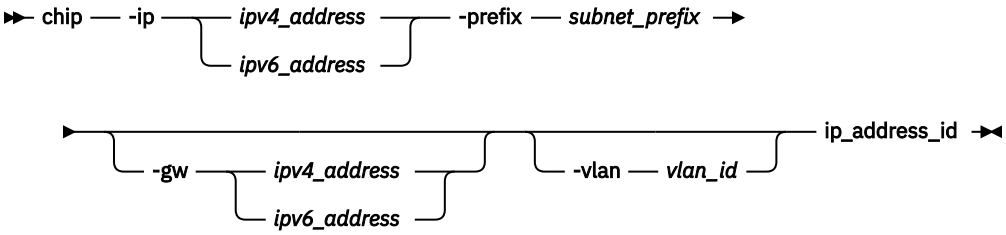
Chapter 27. IP addresses commands

Use the following set of commands to configure or create portset, IP objects and associate them with host objects.

chip

Use the **chip** command to change the system IP address.

Syntax



Parameters

-ip *ipv4_address ipv6_address*
(Required) Specifies the Internet Protocol Version 4 (IPv4) or Internet Protocol Version 6 (IPv6) address.

-prefix *subnet_prefix*
(Required) Specifies the IPv4 or IPv6 prefix of the subnet.

-gw *ipv4_address ipv6_address*
(Optional) Specifies the IPv4 or IPv6 gateway address.

-vlan *vlan_id*
(Optional) Specifies the VLAN ID for a IPv4 or IPv6 address.

ip_address_id
(Required) Specifies the row index of IP address to be changed from **lsip** command output.

Description

This command modifies IP configuration parameters for the system.

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

To modify an IP address, issue the **chip** command and list the IP address of the system by issuing the **lsip** command.

An invocation example

The following example uses specific **-ip**, **-gw**, and **-mask** parameters.

```
chip -ip 9.122.123.193 -prefix 24 -gw 9.122.123.1  -vlan 100 2
chip -ip 9.122.123.193 -prefix 24 2
```

The resulting output:

```
lsip
id node_id node_name port_id portset_id portset_name       IP_address       prefix vlan
gateway        owner_id owner_name
0                   1       72         SystemManagement   192.168.1.52     24
```

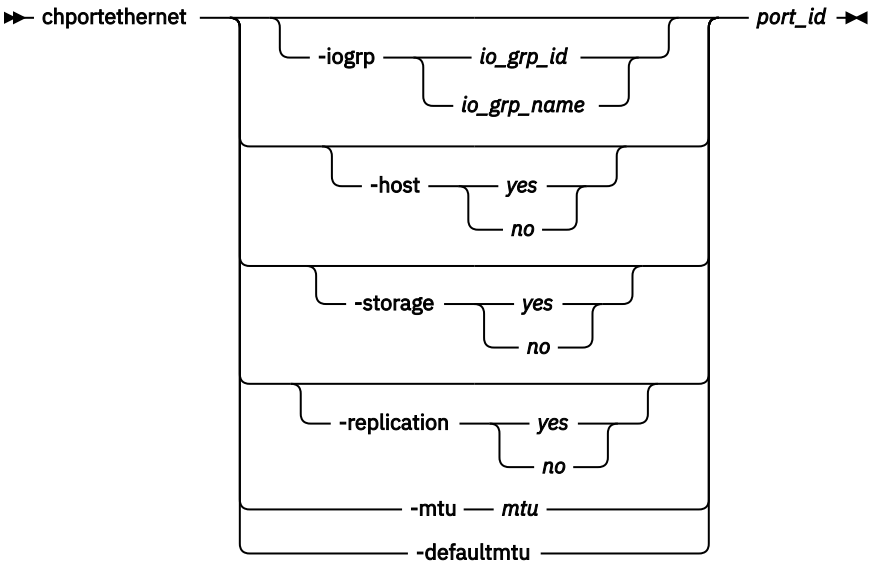
© Copyright IBM Corp. 2003, 2025

```
192.168.1.1
1                    3       72         SystemManagement   192.168.10.52    24
192.168.10.1
$ chip -ip 9.122.123.193 -prefix 24 -gw 9.122.123.1 0
$ lsip
id node_id node_name port_id portset_id portset_name       IP_address       prefix vlan
gateway        owner_id owner_name
0                    3       72         SystemManagement   192.168.10.52    24
192.168.10.1
1                    1       72         SystemManagement   9.122.123.193    24
9.122.123.1
```

chportethernet

Use the **chportethernet** command to configure the permitted uses of ethernet ports and the maximum transmission unit (MTU) setting for the ethernet. Changing MTU value of a ethernet port through **chportethernet** command takes some time to configure MTU in the background. If you want to change the MTU of the ethernet ports, you need to wait for five seconds to reflect the MTU change on the system.

Syntax



Parameters

-iogrp *iogrp*
(Optional) Specifies the I/O group that contains the nodes to modify.

-host *yes* | *no*
(Optional) Specifies port that is used for host attach (the existing system settings are retained). Specifying:

- *yes* turns on IP host attached functionality for that port (default).
- *no* turns off IP host attached functionality for that port.

-storage *yes* | *no*
(Optional) Specifies whether the port can be used for backend storage functionality. The values are *yes* and *no*.

- The value *yes* indicates that this port can be used for backend storage connectivity (default).
- You must specify *no* if you are not using backend storage.

-replication *yes* | *no*
(Optional) Specifies whether the port can be used for replication. The values are *yes* and *no*.

-mtu *mtu* **| -defaultmtu**
(Optional) Specifies the maximum transmission unit (MTU). The default is 1500, with a maximum of 9000. Specify **-defaultmtu** to use the default value.

port_id
(Required) Specifies the port on which changes are applicable.

Note:

- Management-only port refers to dedicated ethernet ports for storage management.
- The changes in port attributes are not allowed for a management-only ports.

An invocation example to disable host attachment for Ethernet port 1 of the nodes in I/O group iogrp0

```
chportethernet -iogrp iogrp0 -host no 1
```

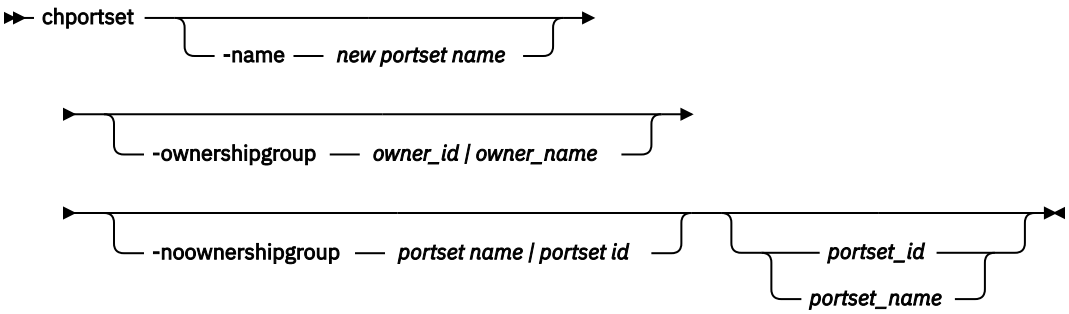
The resulting output:

```
No feedback
```

chportset

Use the **chportset** command to modify the attributes of a portset.

Syntax



Parameters

-name ***new portset name***
(Optional) Specifies a name or label for the new portset. The value can be an ASCII string in the range 0 - 64 characters.

-ownershipgroup ***owner_id | owner_name***
(Optional) Specifies the owner id or own name to assign to portset object. The value can be an ASCII string or decimal in the range 0 - 64 characters.

-noownershipgroup ***portset name | portset id***
(Optional) Remove the portset from the ownership group. It is mutually exclusive with ownershipgroup.

portset_id | portset_name
(Optional) Specifies the portset ID or name.

Note: Restricted users are permitted to view and use globally owned portsets, but are not permitted to make any modifications. Restricted users are permitted to modify portsets that are assigned to their ownership group. Portsets of type storage and replication are always globally owned and an attempt to change its ownership will fail.

An invocation example

```
chportset -name test <portset name|portset id>
```

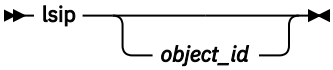
The resulting output:

```
No feedback.
```

lsip

Use **lsip** command to list the currently configured IP addresses. IP addresses can be assigned to different portset types which determine if the IP address is used for host attachment, storage virtualization, IP replication or management.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

object_id
(Optional) Specifies the ID of an object.

Note: Restricted users can view only those IP addresses that they own or IP addresses that are globally owned, but cannot view IP addresses assigned to other owners.

Description

This table provides the attribute values that can be displayed as output view data.

Table 122. ***lsip*** *output*

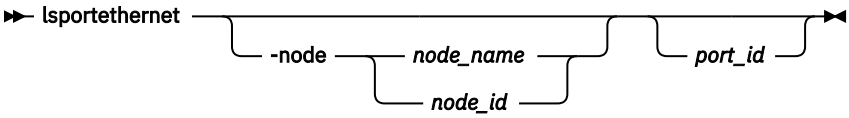
Attribute	Description
port_id	Indicates the ID of the port. The ID can be a value in the range 1 - 28.
io_group_id	Indicates the I/O group ID of the ethernet port. The value must be a number in the range 0 - 3.
io_group_name	Indicates the name of the I/O group ID of the ethernet port. The I/O group ID can consist of a string of 1 - 64 ASCII characters.
node_id	Indicates the ID of the node. The value can be set to blank or can consist of 64 bits.
node_name	Indicates the name of the node. The value can be set to blank or can consist of a string of 1 - 64 ASCII characters.

*Table 122. **lsip** output (continued)*	
Attribute	**Description**
portset_id	Indicates the ID of the portset. The value must be a number in the range 0 - 63 or 72.
portset_name	Indicates the name of the portset. The portset name can consist of a string of 1 - 64 ASCII characters.
ip_address_id	Indicates an IPv4 address or IPv6 address. You can set the address 192.168.1.52 or fc00:0000:0000:0000:445a:0a17:fcf7:0237 format.
prefix	Indicates the prefix of subnet. The value must be a number in the range 1 - 128.
vlan	Displays the virtual local area network (VLAN) ID associated with the IP address on this port (a numeric character in the range 1 - 4094).
gateway	Indicates the IPv4 or IPv6 gateway. The value is set to blank (no gateway) or string in either `192.168.1.52` format or `fc00:0000:0000:0000:445a:0a17:fcf7:0237` format.
owner_id	Indicates the ID of ownership group. The value must be a number in the range 0 - 63.
owner_name	Indicates the name of the ownership group. The ownership group name can consist of a string of 1 - 64 ASCII characters.
lossless	Displays whether priority flow control (PFC) is enabled for this IP address for the portset type or the reason of lossless not enabled. String can be on or off (off-eth_port_not_enabled_for_pfc, off-vlan_not_configured off-not_supported) **Note:** Setting COS value for IP replication is not supported. Hence for IP replication portsets, the value will always be off-not_supported.

lsportethernet

Use the **lsportethernet** command to list the properties and status of the storage system ethernet ports.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-node ***node_id | node_name***
(Optional) Specifies the node ID or node name that contains the port. The ID must be of 64 bits and node name can be a string of 63 ASCII characters.

port_id
(Optional) Indicates the ID of the ethernet port within the node. The ID can be a value in the range 1 - 28.

Description

This command lists the configuration of the Ethernet ports for each node in the IBM Storage Virtualize system.

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.

This table provides the attribute values that can be displayed as output view data.

Table 123. ***`lsportethernet`*** *output*

Attribute	Description
port_id	Indicates the ID of the ethernet port within the node. You can specify a value in the range 1 - 28.
node_id	Indicates the node id. Value can be in decimal of 64 bits.
node_name	Indicates the name of the node that contains the port. Value can be a string of 63 ASCII characters.
MAC	Indicates the current MAC address (blank if unknown).
duplex	Indicates the current duplex state of the port (blank if unknown).
speed	Indicates the negotiated speed of the port.
max_speed	Indicates the maximum speed of the port.
link_state	Indicates the link state of ethernet port. The values are `active` and `inactive`.
mtu	Indicates the maximum transmission unit (MTU) for the port. The default value is `1500`.
rdma_type	Indicates the colon separated list of RDMA protocol capabilities of the port.
adapter_location	Displays the location of the adaptor that contains the Ethernet port (any number in the range 0 - 8). Where 1 - 8 is the PCIe expansion slot number and 0 means that the adaptor is part of the system board or not in a PCIe expansion slot.
adapter_port_id	Displays the location of the ethernet port that is in the adapter (any number in the range 1 - 4).

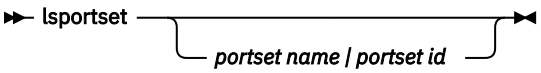
*Table 123. **lsportethernet** output (continued)*	
Attribute	**Description**
dcbx_state	Displays the Data Center Bridging Capability Exchange protocol (DCBx) state of the port. A value of: • `unsupported` indicates that the port does not support priority flow control (PFC) on that port. All ports that are less than 10 Gbps have this value. • `enabled` indicates that the connected switch port is enabled for DCBx and the port state is online. • `disabled` indicates that the connected switch port turned off for DCBx or the port state is offline. For 2 Port 25GbE RoCE Adapter Card, the **lsportethernet** command does not display correct DCBx values.
pfc_enabled_host_attach	Displays whether a PFC tag for host attach is enabled on the connected switch port.
pfc_enabled_storage	Displays whether a PFC tag for storage virtualization is enabled on the connected switch port.
pfc_enabled_clustering	Displays whether a PFC tag for clustering is enabled on the connected switch port.
priority_group_0	Displays the set of priority tags that are within the priority group zero. This value is either blank or is a colon-separated list of numbers from `0` to `7`. This field is part of the Enhanced Transmission Selection (ETS) settings.
priority_group_1	Displays the set of priority tags that are within the priority group one. This value is either blank or is a colon-separated list of numbers from `0` to `7`. This field is part of the ETS settings.
priority_group_2	Displays the set of priority tags that are within the priority group two. This value is either blank or is a colon-separated list of numbers from `0` to `7`. This field is part of the ETS settings.
priority_group_3	Displays the set of priority tags that are within the priority group three. This value is either blank or is a colon-separated list of numbers from `0` to `7`. This field is part of the ETS settings.
priority_group_4	Displays the set of priority tags that are within the priority group four. This value is either blank or is a colon-separated list of numbers from `0` to `7`. This field is part of the ETS settings.
priority_group_5	Displays the set of priority tags that are within the priority group five. This value is either blank or is a colon-separated list of numbers from `0` to `7`. This field is part of the ETS settings.
priority_group_6	Displays the set of priority tags that are within the priority group six. This value is either blank or is a colon-separated list of numbers from `0` to `7`. This field is part of the ETS settings.
priority_group_7	Displays the set of priority tags that are within the priority group seven. This value is either blank or is a colon-separated list of numbers from `0` to `7`. This field is part of the ETS settings.
bandwidth_allocation	Displays a list of guaranteed bandwidth allocation percentages for priority groups zero through seven. The value is either blank or a colon-separated numeric list of bandwidth allocation percentages for each priority group, where each percentage is a whole number integer. This field is part of the ETS settings.

Table 123. **lsportethernet** output (continued)	
Attribute	**Description**
management	Indicates that this port can be used for management. Values are yes for ports 1 and 2, no for the rest.
host	Indicates if the port can be used for host attachment. The values are yes (default) and no.
storage	Indicates if the port can be used for storage virtualization. The values are yes (default) and no.
replication	Indicates if the port can be used for replication. The values are yes (default) and no.
eth_clustering	Indicates if the port is configured for ethernet clustering. The values are yes and no (default).

lsportset

Use the **lsportset** command to list the currently configured portsets. This list shows the ID, portset name, and the number of IP addresses and hosts associated with each portset along with ownership group information.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

portset_name | portset_id
(Optional) Specifies the ID and name of the portset. The name value is an alphanumeric string of 64 characters long and ID value is in decimal of 0 - 71.

Description

This table provides the attribute values that can be displayed as output view data.

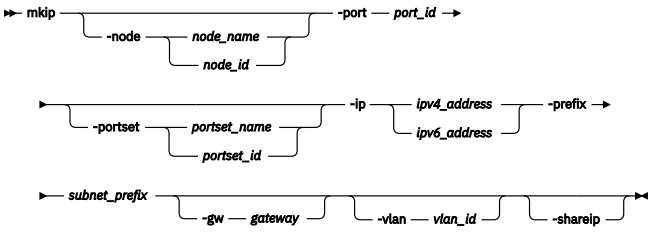
Table 124. lsportset output	
Attribute	**Description**
id	Indicates the ID of the portset.
name	Indicates the name of the portset.

Table 124. lsportset output (continued)	
Attribute	**Description**
type	Indicates the type of portset. The value can be • host • storage • replication • highspeedreplication • management
port_count	Indicates the number of ports (IP addresses or WWPNs) associated with the portset. The value can be a number in the range 0 - 255.
host_count	Indicates the number of host or partnership objects associated with this portset. Default value is set to 1 for the storage portset. Use **lshost** and **lspartnership** to view the associations of portset to host and partnership objects. The value can be a number in the range 0 - 2048.
lossless	Indicates whether all ports (IP addresses or WWPNs) for this portset are **lossless**. Even if a single IP address is not configured as **lossless**, the portset is not a **lossless** portset. Use **lsip** to view **lossless** parameter per IP address.
owner_id	Indicates the ID of the ownership group.
owner_name	Indicates the name of the ownership group.
port_type	Indicates the type of ports in portset. Supported values: **ethernet** or **fc**.
is_default	Indicates whether the portset is the default portset for type or port_type combinations. Supported values: **yes** or **no**.

mkip

Use the **mkip** command to create a new IP address and associate it with an existing portset. You can create one IP address for each node, port, and portset.

Syntax



Parameters

-node ***node_name | node_id***
(Optional) Specifies which node has the ethernet port that the IP address is being assigned to. Value can be in decimal of 64 bits.

Note: Node parameter is optional for system IP addresses and required for all other IP addresses.

port_id
(Required) Specifies the port on which changes are applicable. The ID can be a value in the range 1 - 28.

Note: Management-only port refers to dedicated ethernet ports for storage management.

-portset *portset_id | portset_name*
(Optional) Specifies the portset ID or name that the IP address is being assigned to. The portset ID value is of 0 - 63, 72 and the portset name consists of a string of 64 characters.

-ip *ipv4_address | ipv6_address*
(Required) Specifies the IPv4 address or IPv6 address. The value consists in the standard format for IPv4 address, or IPv6 address of a string of 40 characters.

-prefix *subnet_prefix*
(Required) Specifies the prefix of the subnet. The value for IPv4 addresses can range from 1 to 31. The value for IPv6 addresses can range from 1 to 127.

-gw *gateway*
(Optional) Specifies the IPv4 or IPv6 gateway address. The value consists in the standard format for IPv4 address, or IPv6 address of a string of 40 characters.

-vlan *vlan_id*
(Optional) Sets the virtual local area network (VLAN) ID for a IPv4 or IPv6 address. The value can be set between 1 - 4095.

-shareip
(Optional) Specifies that an IP address is being shared between multiple portsets.

Note: Restricted users can create, modify, or delete IP addresses for portsets that they own. Restricted users can view but are not allowed to create, modify, or delete IP addresses for portsets that have global ownership.

Note: The portset that is specified must be of the type ethernet.

Note: The **mkip** command can now be used to configure the second system IP address.

An invocation example

```
mkip -node node1 -port 1 -portset myportset -ip 10.0.1.1 -prefix 24 -vlan 1001
```

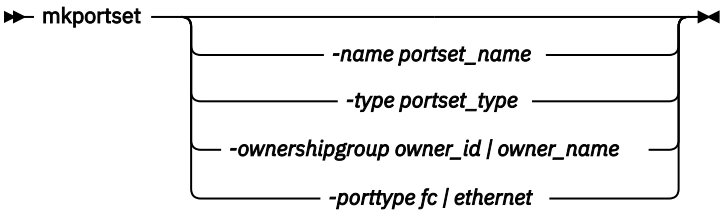
The resulting output:

```
IP Address, id [0], successfully created
```

mkportset

Use the **mkportset** command to create a new portset.

Syntax



Parameters

-name *portset_name*
(Optional) Specifies the name of the portset.

-type *portset_type*
(Optional) Specifies the type of portset. It can be host, replication, or highspeedreplication type portset.

-ownershipgroup *owner_id | owner_name*
(Optional) Specifies the name or ID of the ownership group to which the object is added. This option is available only to the global user.

-porttype *fc | ethernet*
(Optional) Specifies the type of port that can be mapped to the portset. By default **mkportset** creates an ethernet portset if **-porttype** parameter is not provided. High-speed replication portset has to be of type ethernet only

Note: By default there are 4 portsets defined on a system **-portset0**, **-portset1**, **-portset2** and **-portset3**. The administrator must create a new portsets and associate them with host or partnership objects as necessary. The system supports up to 72 portsets. By default a portset type is host attachment. Ownership rules:

- Global owner is allowed to create a portset with restricted owner ID A.
- Restricted owner ID A can create portset with owner ID A.

Note:

- The **mkip** command can map an IP address with a portset of **-porttype ethernet** only.
- The **addfcportsetmember** command can map a Fibre Channel port with a portset of **-porttype fc** only.
- The **mkhost** command for Fibre Channel or NVMe over Fibre Channel hosts can map the host object with a portset of **-porttype fc**.
- The **mkhost** command for iSCSI, NVMe over RDMA , and NVMe over TCP can map the host object with a portset of **-porttype ethernet**.

Note: A global admin by default create portsets with no ownership which can be used by restricted users. Alternatively a global owner may use **-ownershipgroup** to assign the new portset to a specific ownership group. A restricted user cannot use the **-ownershipgroup** parameter.

A concise invocation example

```
mkportset
```

The resulting output:

```
Portset, id [7], successfully created
```

rmip

Use **rmip** command to delete an IP address specified by the IP address ID.

Syntax

Parameters

ip_address_id
(Required) Specifies the row ID (lsip) of the IP address to be deleted.

CAUTION:

- User is not allowed to delete the last IP address in a portset if there are any hosts associated with the portset. Delete hosts or assign to different portsets first and then delete IP addresses.
- Cannot remove an IP address if it is a part of the storage portset or replication portset and is in use.
- The **rmip** command can be used to remove system IP address, except the last system IP address.

Note:

An invocation example

```
rmip 1
```

The resulting output:

```
No feedback.
```

rmportset

Use the **rmportset** command to delete the portset. Before deleting the Fibre Channel (FC) port make sure if any Fibre Channel host or Fibre Channel IO port under the portset is defined.

Syntax



Parameters

-rmportset *portset_name | portset_id*
(Required) The name or ID of the portset.

Note: The **portset0**, **portset3** and **portset64** are system pre-defined portsets of type host and storage respectively and cannot be deleted. If portset has port count and host count set then portset cannot be deleted.

Note: Depending on the portset type (fc or ethernet), check for any existence of host or ports under portset before deleting it.

- If you are deleting Fibre Channel portset, check for any Fibre Channel host or Fibre Channel IO ports under the portset.
- If you are deleting Ethernet portset, check for any iSCSI host or Ethernet ports under the portset.

An invocation example

```
rmportset portset1
```

The resulting output:

```
No feedback
```

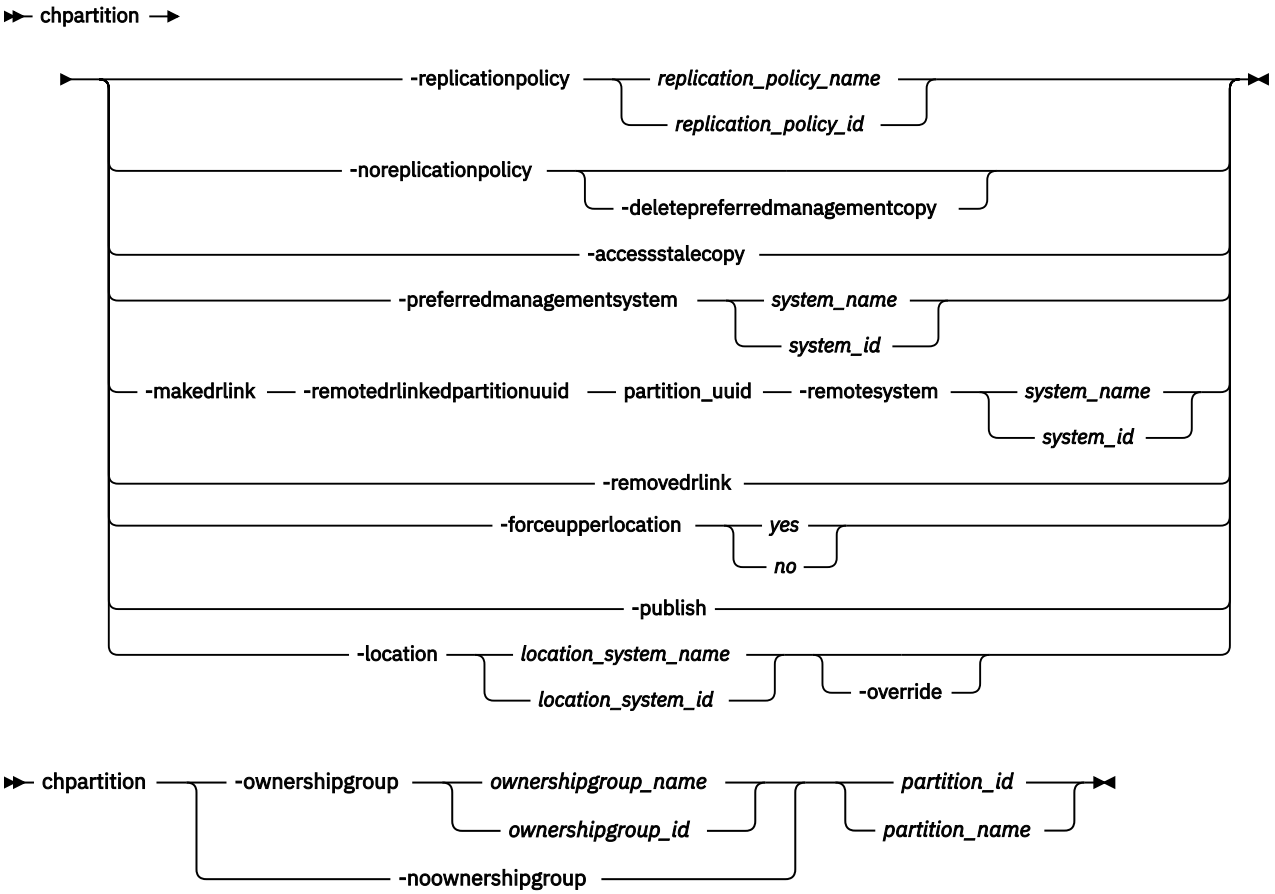
Chapter 28. Partition commands

The High Availability (HA) solution uses Storage Partitions, a configuration object that is designed to be the single point of management for HA objects. These partitions contain volumes, volume groups, hosts and host-to-volume mappings.

chpartition

Use the **chpartition** command to modify a storage partition.

Syntax



Parameters

-replicationpolicy *replication_policy_name | replication_policy_id*

(Optional) Specifies to assign a replication policy for the storage partition. This operation can be performed whether the remote systems in the policy are connected or disconnected. If the command is successful, the change is applied asynchronously to the remote system. The replication policy must be a HA replication policy.

A replication policy cannot be applied to a storage partition if any volume group that is associated with the storage partition contains volumes that are:

- cache disabled
- image mode
- configured to use Transparent Cloud Tiering (TCT)

© Copyright IBM Corp. 2003, 2025

- zero size
- in the process of removing associated system resources for replication
- a volume with a copy in a standard pool and a copy in a data reduction pool.

A replication policy cannot be assigned to a partition if any associated volume groups have a replication policy with a 2-site-async-dr topology assigned.

When making a disaster recovery link, if there are any volume groups in the partition with a 2-site-async-dr replication policy assigned, the specified remote system must match the remote system of the 2-site-async-dr replication policies.

Additionally, any 2-site-async-dr volume groups that are being migrated to async-dr replication policies must have been added to both the local and remote partition before a disaster recovery link can be made.

Note: This parameter is mutually exclusive with **-noreplicationpolicy**, **-accessstalecopy** and **-preferredmanagementsystem** parameters.

-noreplicationpolicy
(Optional) Specifies to remove the replication policy from the storage partition. This operation can be performed whether the remote systems in the policy are connected or disconnected. The remote storage partition and associated objects are removed asynchronously.

If the **-deletepreferredmanagementcopy** flag is specified, the command is only permitted on the active management system, and the active management system should not be equal to the preferred management system. This allows for the replication policy to be removed in response to the complete loss of the remote system.

Note: This parameter is mutually exclusive with **-replicationpolicy**, **-accessstalecopy** and **-preferredmanagementsystem** parameters.

-deletepreferredmanagementcopy
(Optional) Specifies to remove the preferred management copy of the storage partition. Specify **-noreplicationpolicy** parameter with this parameter and the active management system must not be the same as the preferred management system.

-accessstalecopy
(Optional) Specified to modify the UUID of the volumes associated with the storage partition. This parameter is only permitted when the command runs on the non-active management system and the preferred management system is ignored. The UUID of all volumes in the storage partition will be changed and **vdiskhostmaps** will be deleted. The replication policy is removed from the storage partition resulting in a single-system storage partition, and all fields that refer to a replication policy in the storage partition view will be blank.

If volumes are offline on the system where this command is run, the command completes successfully, but all volumes in each group remain offline until a consistent image can be made available. Volumes may be deleted while in this state and a new replication policy cannot be applied until the consistent image is restored.

Note: This parameter is mutually exclusive with **-replicationpolicy**, **-noreplicationpolicy** and **-preferredmanagementsystem** parameters.

-preferredmanagementsystem
(Optional) Specifies to modify the preferred management system for the storage partition. This operation can be performed while the remote systems in the policy are connected, or under the following conditions while disconnected:

- If the storage partition has HA and there is management access to the remote system, the command succeeds and the change will be applied to both systems.
- If the storage partition has HA and there is no management access to the remote system then the command will fail.
- If the storage partition does not have HA, then the operation succeeds and the configuration change will be applied to the remote system in line with other configuration and data syncing that is

required to regain HA, which may mean to wait for management connectivity to resume if there is currently no connectivity.

Note: This parameter is mutually exclusive with **-replicationpolicy**, **-accessstalecopy** and **-noreplicationpolicy** parameters.

-makedrlink
(Optional) Creates the disaster recovery link between this partition and a partition on a remote system. This operation can only be performed when the partition does not already have a disaster recovery link.

Note: This parameter requires **-remotesystem** and **-remotedrlinkedpartitionuuid** and is mutually exclusive with **-removedrlink**.

-remotesystem
(Optional) Specifies the remote system where the partition to link exists.

Note: This parameter requires **-makedrlink** and **-remotedrlinkedpartitionuuid**.

-remotedrlinkedpartitionuuid
(Optional) Specifies the uuid of the remote partition to link.

Note: This parameter requires **-makedrlink** and **-remotesystem**.

-removedrlink
(Optional) Removes the disaster recovery link from this partition.

Note: This parameter is mutually exclusive with **-makedrlink**

-forceupperlocation
(Optional) Specifies this partition should be shown in the upper location of the topology diagram in the GUI, overriding the default layout for the view.

Note: This parameter is only used for partitions that are configured for 2-site replication for disaster recovery.

-publish
(Optional) Specify to publish a draft storage partition.

-location ***location_system_name | location_system_id***
(Optional) Specifies a desired target system location to migrate an existing storage partition.

-override
(Optional) This parameter should be specified along with the **-location** parameter to trigger any operations on an ongoing storage partition migration.

When used with **-location** at source system:

- If the location specified is the source system, it aborts any ongoing active or queued storage partition migration.
- If a new location is specified for any ongoing storage partition migration, the ongoing migration gets aborted and the migration to the new location gets initiated or queued.
- Abort is supported on source until the user completes the host rescan action and responds to the host rescan event.

When used with **-location** at target system:

- Rolls back the migrating storage partition to its original state at source system without deleting the target copy.
- Rollback is supported on target until the user confirms the source copy delete event.

partition_name | partition_id
(Required) Specifies the name or ID of the storage partition.

Description

This command modifies the attributes of a storage partition.

Invocation examples

- The following example changes the preferred management system for a storage partition.

```
chpartition -preferredmanagementsystem myTargetFlashSystem mypartition1
```

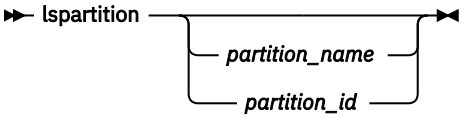
The resulting output:

```
No Feedback
```

lspartition

Use the **lspartition** command to generate a list with concise or detailed information on storage partitions on the system.

Syntax



Parameters

partition_name | partition_id
(Optional) Specifies the name or ID of the partition. When this parameter is used, a detailed view of the storage partition is displayed. If this parameter is omitted, a concise view of all storage partitions is displayed.

This table describes possible outputs.

*Table 125. **lspartition** output*

Attribute	Value
preferred_management_location	Indicates the location number (1 or 2) for the preferred management system. This identifier can be used to determine which of the location1 or location2 fields correspond to preferred management system.
preferred_management_system_id	Indicates the ID of the system for the preferred management system.
preferred_management_system_name	Indicates the name of the system for the preferred management system.
active_management_location	Indicates the location number (1 or 2) for the active management system. This identifier can be used to determine which of the location1 or location2 fields correspond to active management system.
active_management_system_id	Indicates the ID of the system for the active management system.
active_management_system_name	Indicates the name of the system for the active management system.
replication_policy_id	Indicates the ID of the replication policy that is associated with the partition.

*Table 125. **lspartition** output (continued)*	
Attribute	**Value**
replication_policy_name	Indicates the name of the replication policy that is associated with the partition.
location1_system_id	Indicates the ID of the system in location1 of the replication policy that is associated with the partition.
location1_system_name	Indicates the name of the system in location1 of the replication policy that is associated with the partition.
location1_status	Indicates the health status of the system in location1 of the replication policy that is associated with the partition. The following statuses can be displayed: • healthy • unhealthy • disconnected • suspended • blank A status of *unhealthy* indicates that there is an error that affects replication. A status of *suspended* indicates that user intervention is required.
location2_system_id	Indicates the ID of the system in location 2 of the replication policy that is associated with the partition.
location2_system_name	Indicates the name of the system in location 2 of the replication policy that is associated with the partition.
location2_status	Indicates the health status of the system in location 2 of the replication policy that is associated with the partition. The following statuses can be displayed: • healthy • unhealthy • disconnected • suspended • blank A status of *unhealthy* indicates that there is an error that affects replication. A status of *suspended* indicates that user intervention is required.
host_count	Indicates the number of hosts that are associated with the partition.
host_offline_count	Indicates the number of hosts in the storage partition that are offline on one or more systems.
volume_group_count	Indicates the number of volume groups that are associated with the partition.
volume_group_synchronized_count	Indicates the number of volume groups that are associated with the partition that are synchronized.
volume_group_synchronizing_count	Indicates the number of volume groups that are associated with the partition that are synchronizing.
volume_group_stopped_count	Indicates the number of volume groups that are associated with the partition that are not synchronized or actively synchronizing.

*Table 125. **lspartition** output (continued)*	
Attribute	**Value**
ha_status	Indicates the HA status between the two locations in the replication policy that is associated with the partition. The following statuses can be displayed: • established • establishing • suspended • problem • blank A status of *establishing* indicates that HA is in the process of being established and no user action is required.
link_status	Indicates the status of HA replication between the systems in location1 and location 2 of the replication policy assigned to the storage partition. The following statuses can be displayed: • synchronized • synchronizing • suspended • disconnected • blank
desired_location_system_id	Indicates the cluster ID of the remote system to which the storage partition is migrated.
desired_location_system_name	Indicates the cluster name of the remote system to which the storage partition is migrated.
migration_status	Indicates the storage partition migration status. One of the following status will be displayed: • queued • in progress • user action required • aborting • blank • suspended
uuid	Indicates the universally unique identifier for this partition.
dr_link_status	Indicates the status of the DR replication link for the partition. This field is only populated when the partition is linked for DR replication. One of the following status will be displayed. • unused - The partition is linked for DR replication, but no volume groups are configured for DR replication. • healthy - All volume groups configured for DR replication are either running or independent. • problem - There is a problem affecting the DR replication for one or more volume groups. Check the volume groups in the partition. • disconnected - A connectivity problem prevents DR replication for the volume groups in this partition.

*Table 125. **lspartition** output (continued)*	
Attribute	**Value**
dr_link_direction	Indicates the direction of replication using the DR replication link for the partition. This field is only populated when the partition is linked for DR replication.
dr_linked_partition_uuid	The UUID of the remote partition linked for disaster recovery replication. This field is only populated when the partition has a disaster recovery link.
dr_linked_partition_name	Indicates the name of the remote partition linked for disaster recovery replication.
dr_linked_remote_system_1_id	Indicates the ID of the first system in the DR-linked partition.
dr_linked_remote_system_1_name	Indicates the name of the first system in the DR-linked partition.
dr_linked_remote_system_2_id	Indicates the ID of the second system in the DR-linked partition.
dr_linked_remote_system_2_name	Indicates the name of the second system in the DR-linked partition.
force_upper_location	Indicates if this partition should be shown in the upper location of the topology diagram in the GUI, overriding the default layout for the view. This attribute indicates any user-specified preference for the layout. The **local_location** field below determines the actual layout. It is only used for partitions configured for 2-site replication for disaster recovery.
local_location	Indicates if this partition is in the upper or lower location of the topology diagram. This attribute is only used for partitions configured for disaster recovery replication.
draft	Indicates that this partition is a draft.
draft_volume_group_count	Indicates the number of volume groups that are associated with this draft partition.
draft_host_count	Indicates the number of hosts that are associated with this draft partition.
user_action_sequence_number	Indicates the event sequence number for which the user action is requested.
user_action_type	Indicates the type of the user action. The following statuses can be displayed: • host_rescan_requested • commit_or_rollback
default_volume_group_id	Displays the ID of the default volume group for the partition.
default_volume_group_name	Displays the name of the default volume group for the partition.
ownership_group_id	Displays the ID of the ownership group for the partition.
ownership_group_name	Displays the name of the ownership group for the partition.
vcenter_id	Display the ID of the vcenter for the partition.
vcenter_name	Display the name of the vcenter for the partition.

Description

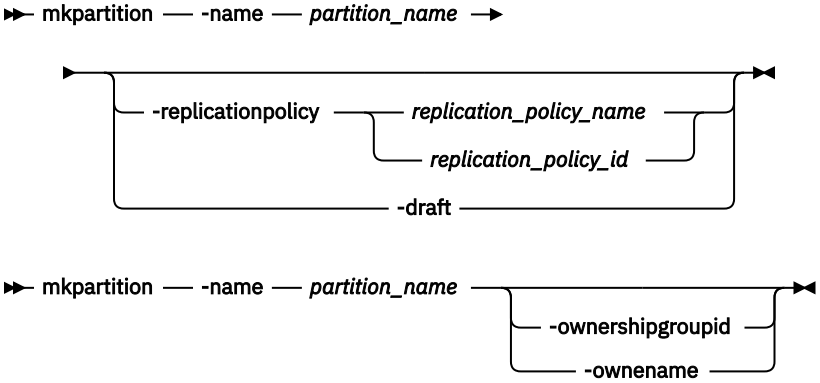
The following fields are blank for a single-system partition:

- replication_policy_id
- replication_policy_name
- preferred_management_location
- preferred_management_system_id
- preferred_management_system_name
- active_management_location
- active_management_system_id
- active_management_system_name
- location2_system_id
- location2_system_name
- location2_status
- host_count
- host_offline_count
- volume_group_count
- volume_group_synchronized_count
- volume_group_synchronizing_count
- volume_group_stopped_count
- ha_status
- link_status

mkpartition

Use the **mkpartition** command to create a storage partition.

Syntax



Parameters

-name *partition_name*
(Required) Specifies the name for the partition. The name must be unique on all systems where the partition is used and cannot be changed later.

-replicationpolicy *replication_policy_name | replication_policy_id*
(Optional) Specifies the replication policy for the partition. This operation can be performed whether the remote systems in the policy are connected or disconnected. The partition is created on the remote systems when connectivity allows.

-draft
(Optional) Creates a draft partition.

Description

This **mkpartition** command is used to create a partition. A system can have a maximum of four partitions and can be created only on single IO Group systems.

Note: If a replication policy is specified, the policy must be a HA policy. If no policy is specified, the partition is a single-system partition.

An invocation example

The following example creates a partition with no replication policy:

```
mkpartition -name mypartition1
```

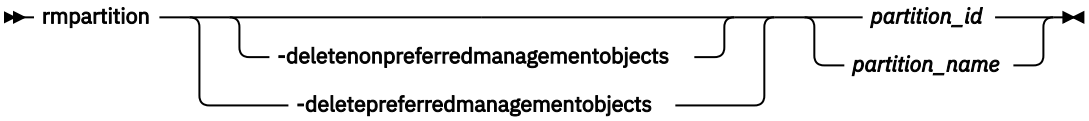
The resulting output:

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

rmpartition

Use the **rmpartition** command to remove a partition.

Syntax



Parameters

-name partition_name | partition_id
(Required) Specifies the name of the partition to be removed.

-deletenonpreferredmanagementobjects
(Optional) Specified to remove the objects associated with the partition on the non-preferred management system. If the partition has a replication policy and associated objects, then the **rmpartition** command requires the **-deletenonpreferredmanagementobjects** or **-deletepreferredmanagementobjectsone** flag to succeed. Otherwise, the command will fail.

If the **-deletenonpreferredmanagementobjects** flag is specified, the command is only allowed on the active management system, and requires that the active management system is the same as the preferred management system. Objects associated with the partition on the non-preferred management system will be removed. The associated objects will remain on the preferred management system but will no longer be associated with the partition or the replication policy that was assigned to the partition.

The command fails if the **-deletenonpreferredmanagementobjects** parameter is specified for a partition that cannot be managed at the preferred management system, or if the partition has no replication policy.

Note: This parameter is mutually exclusive with **-deletepreferredmanagementobjects** parameter.

-deletepreferredmanagementobject
(Optional) Specified to remove the objects associated with the partition on the preferred management system.

If the partition has a replication policy and associated objects, then the **rmpartition** command requires **-deletenonpreferredmanagementobjects** or **-deletepreferredmanagementobjectsone** flag to succeed otherwise the command will fail.

If the partition cannot be managed at the preferred management system then use **-deletepreferredmanagementobjects** parameter to remove the partition and unassign the replication policy. This will remove the partition and associated objects at the preferred management system.

The command will fail if **-deletepreferredmanagementobjects** is specified for a partition that cannot be managed at the non-preferred management system, or if the partition has no replication policy.

Note: This parameter is mutually exclusive with **-deletenonpreferredmanagementobjects** parameter.

Description

Removing a partition does not delete the objects in the partition. If a replication policy is assigned then there will be a copy of the objects in each of the locations associated with the policy, and one of the optional flags **-deletenonpreferredmanagementobjects** or **-deletepreferredmanagementobjects** are required. The optional flag **-deletenonpreferredmanagementobjects** or **-deletepreferredmanagementobjects** is required to track the location to remove the associated objects from both the locations. These flags are not permitted if the partition does not have a replication policy.

Note:

A storage partition cannot be removed if it contains DR replicated **volumegroups**.

A storage partition can be removed if it contains a DR replication policy. When a storage partition is removed any DR replication policy in the storage partition will also be removed.

An invocation example

The following example removes a partition with no replication policy:

```
rmpartition mypartition1
```

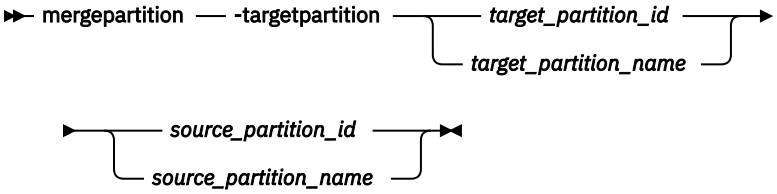
The resulting output:

```
No Feedback
```

mergepartition

Use the **mergepartition** command to merge the source partition into the target partition and delete the source partition.

Syntax



Parameters

-targetpartition ***target_partition_id | target_partition_name***
(Required) Specifies the target storage partition, either by ID or by name.

source_partition_id | source_partition_name
(Required) Specifies the source storage partition, either by ID or by name.

Description

The **mergepartition** command merges the source partition into the target partition, and the source partition will be deleted. The source partition is identified by the command object and the target is identified as a parameter. The command works with both local-only and HA storage partitions; if the storage partition is configured as HA then the changes are applied to both systems.

Note: It is not permitted to merge from a storage partition with a DR link, the command will fail if the source storage partition has a DR link.

Note: A storage partition that contains volume groups with 2-site-async-dr replication policies cannot be merged into a storage partition that has a DR link. If both the source and target partitions contain 2-site-async-dr volume groups, they can only be merged if they are all configured for replication to the same remote system.

An invocation example

Command to merge 2 partitions:

```
mergepartition -targetpartition mypartition1 mypartition2
```

Chapter 29. Provisioning policies commands

Use the provisioning policies commands to simplify the provisioning of storage.

chprovisioningpolicy

Use the **chprovisioningpolicy** command to change the name of a provisioning policy on the system.

Syntax



Parameters

-name ***new_name***
(Required) Specifies a new name for the provisioning policy.

policy_name | policy_id
(Required) Name or ID of the provisioning policy.

Note: Only the name of a provisioning policy can be changed.

An invocation example

```
chprovisioningpolicy -name dbVolumes provisioningpolicy5
```

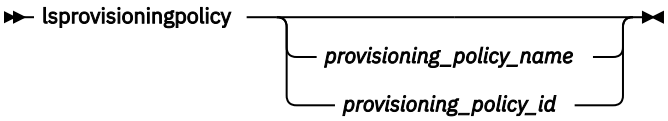
The resulting output:

```
No feedback
```

lsprovisioningpolicy

Use the **lsprovisioningpolicy** command to list the provisioning policies available in the system.

Syntax



Parameters

provisioning_policy_name | provisioning_policy_id
(Optional) The name or ID of the provisioning policy.

Description

Use this command to display a concise list or a detailed view of provisioning policies that are configured on the system.

This table provides the attribute values that can be displayed as output view data.

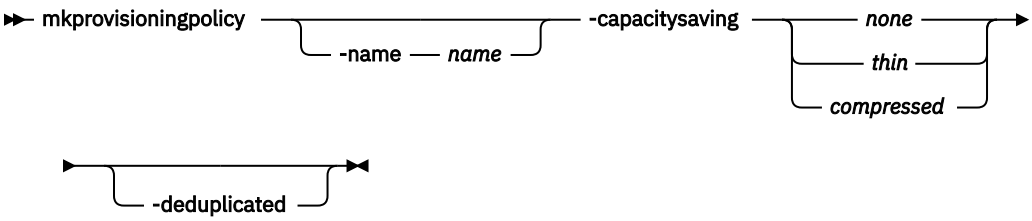
© Copyright IBM Corp. 2003, 2025

Table 126. **lsprovisioningpolicy** output	
Attribute	**Description**
id	Indicates the system defined ID of the provisioning policy.
name	Indicates the name of the provisioning policy.
capacity_saving	Indicates the capacity saving that is specified by the provisioning policy.
deduplication	Indicates whether deduplication is specified by the provisioning policy.
in_use	Indicates whether the policy is associated with one or more storage pools.

mkprovisioningpolicy

Use the **mkprovisioningpolicy** command to create a new provisioning policy on the system.

Syntax



Parameters

-name
(Optional) Name of the provisioning policy.

-capacitysaving ***none | thin | compressed***
(Required) Specifies the capacity savings for the provisioning policy.

The capacity savings specified by the provisioning policy are applied to new volumes in addition to any built-in capacity saving functions of the storage pool. For example, pools containing certain drive types, such as FlashCore Module drives, are self-compressing.

Note:

- A provisioning policy specifying thin-provisioning can only be used with standard storage pools.
- A provisioning policy specifying compression can only be used with data reduction storage pools.

-deduplicated
(Optional) Specifies the provisioning policy should use deduplication.

Note: A provisioning policy specifying deduplication can only be used with data reduction pools.

An invocation example

```
mkreplicationpolicy -capacitysaving thin -warning 25% -buffersize 5 -unit gb
```

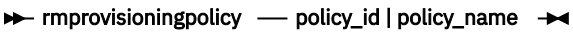
The resulting output:

```
Provisioning policy, id [2], successfully created
```

rmprovisioningpolicy

Use the **rmprovisioningpolicy** command to remove a provisioning policy on the system.

Syntax



Parameters

policy_id | policy_name
(Required) The name or ID of the provisioning policy to remove.

Description

A provisioning policy cannot be removed while it is assigned to a storage pool. You must remove the provisioning policy from the pool before removing the policy itself.

An invocation example

```
rmprovisioningpolicy 0
```

The resulting output:

```
No feedback
```

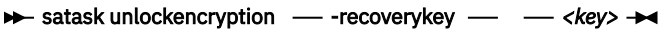
Chapter 30. Recovery access key command

Use the recovery access key command to unlock the system storage.

satask unlockencryption

Use the **satask unlockencryption** command to allow a security administrator to unlock the system storage in scenarios when the system lacks the required quorum and the nodes are unable to process a configuration node.

Syntax



Parameters

-recoverykey *<key>*
(Required) Specifies the 44 character base64-encoded recovery key.

An invocation example

```
satask unlockencryption -recoverykey
Enter the recovery key for the system:
```

A confirmation message is displayed when the recovery key has been entered correctly.

© Copyright IBM Corp. 2003, 2025

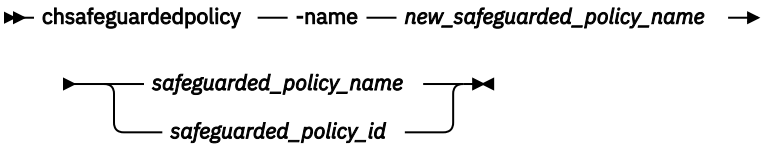
Chapter 31. Legacy Safeguarded policy commands

These Safeguarded policy commands are used for a method of simplifying copies that are made on IBM Copy Services Manager. This is legacy functionality that has been replaced by the Safeguarded Snapshots feature.

chsafeguardedpolicy

Use the **chsafeguardedpolicy** command to change the name of an existing Safeguarded policies.

Syntax



Parameters

-name new_safeguarded_policy_name
(Required) Specifies a new name for the Safeguarded policy.

safeguarded_policy_name | safeguarded_policy_id
Specifies the ID or name of an existing Safeguarded policy.

Note: Only the name of an existing Safeguarded policy can be changed.

An invocation example

```
chsafeguardedpolicy -name mysgpolicy3 3
```

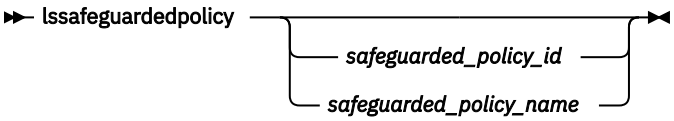
The resulting output:

```
No feedback
```

lssafeguardedpolicy

Use the **lssafeguardedpolicy** command to list the Safeguarded policies.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

© Copyright IBM Corp. 2003, 2025

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

Description

Use this command to display a concise list or a detailed view of the Safeguarded policies.

Note: The concise view list displays predefined Safeguarded policies first.

This table provides the attribute values that can be displayed as output view data.

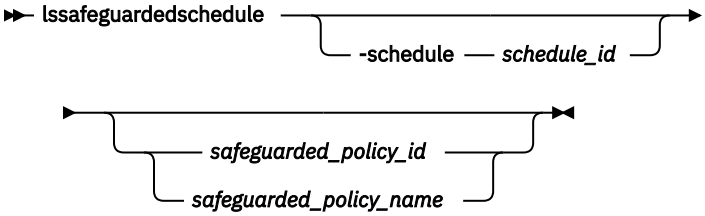
Table 127. ***lssafeguardedpolicy*** *output*

Attribute	Description
id	Indicates the system defined ID of the Safeguarded policies.
name	Indicates the name of the Safeguarded policies.
volume_group_count	Indicates the number of volume groups that are associated with the Safeguarded policies.
schedule_id	Indicates the *schedule_id* of the Safeguarded policies.
backup_unit	Indicates the unit of time of the *backup_interval*.
backup_interval	Indicates the interval between the Safeguarded backups.
backup_start_time	Indicates the time of the Safeguarded backups.
retention_days	Indicates the number of days to keep the Safeguarded backups.
external	Indicates the Safeguarded Policy frequency and retention is stored in external scheduling applications.

lssafeguardedschedule

Use the **lssafeguardedschedule** command to display the Safeguarded backup schedule that is associated with a Safeguarded policies.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-schedule schedule_id
(Required) Specifies a schedule ID for the Safeguarded policies. Only allowed value is 1.

safeguarded_policy_id | safeguarded_policy_name
Specifies the ID or name of an existing Safeguarded policy. Integer value from 0 - 31.

Note:

The **-schedule** parameter requires that a Safeguarded policy is specified.

The **backup_start_time** parameter is displayed in the standard *YYMMDDHHMMSS* format, but the seconds always be 00.

Note: If a Safeguarded policy is external the following field will be blank:

- backup_unit
- backup_interval
- backup_start_time
- retention_days

Description

This table provides the attribute values that can be displayed as output view data.

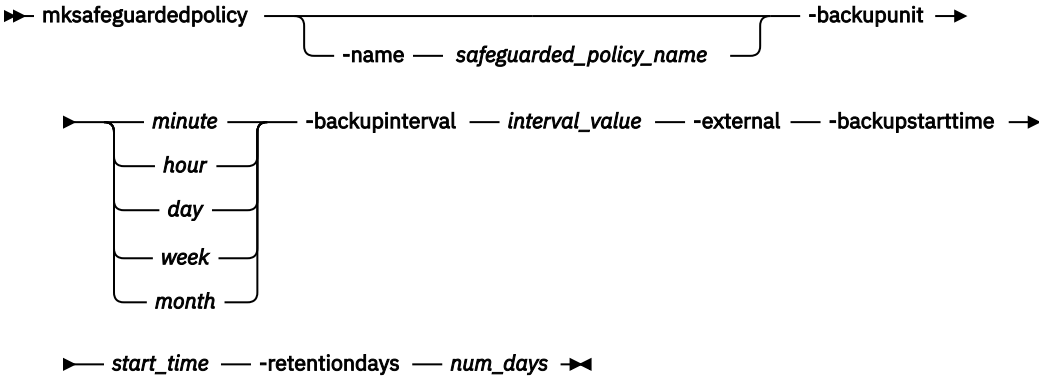
Table 128. ***lssafeguardedschedule*** *output*

Option	Description
name	Indicates the name of the Safeguarded policies.
backup_interval	Indicates the interval between the Safeguarded backups.
backup_start_time	Indicates the time the first safeguarded backup will be taken. If this **backup_start_time** is in the past, internally a new start time will be calculated by moving the original **backup_start_time** forward in increments of the **backup_interval** value until a current start time is found. **Note:** The **backup_start_time** value can be overriden at the time a volume group is associated with the safeguarded policy. Both **mkvolumegroup** and **chvolumegroup** commands have an input option to add a policy start time.
retention_days	Indicates the number of days to keep the Safeguarded backups.

mksafeguardedpolicy

Use the **mksafeguardedpolicy** command to create a new Safeguarded policies to associate with the volume groups.

Syntax



Parameters

-name ***safeguarded_policy_name***
(Optional) Indicates a unique name for the Safeguarded policy. If name is not specified then a unique name is automatically generated with prefix *sgpolicy* and assigned to the Safeguarded policy.

-backupunit minute | hour | day | month | year
(Required) Specifies the unit of time of the backup interval.

-backupinterval interval_value
(Required) Indicates the interval between Safeguarded backups. The value must be an integer (1 - 65535).

-external
(Required) Indicates if schedule is described by an external software.

Note:

(Required) The **-external** parameter indicates the schedule and retention information stored in external copy management software. One of the following parameter must be specified:

Mutually exclusive parameter with **-external**

- **-backupunit**
- **-backupinterval**
- **-backupstarttime** and **-retentiondays**

-backupstarttime start_time
(Required) Indicates the time when the Safeguarded backups are started. The value must be in the format YYMMDDHHmm.

- YY is the year.
- MM is the month.
- DD is the day.
- HH is the hour.
- mm is the minute.

Note: If the backup start time is in the past, it will be advanced according to the backup interval until a time in the future is reached. Backup start time does not include seconds.

-retentiondays num_days
(Required) Indicates the number of days to keep the Safeguarded backups. The value must be an integer (1 - 365).

The system supports a maximum of 32 Safeguarded policies with 3 predefined policies and 29 user-defined policies. Predefined policies are named **predefinedsgpolicy0**, **predefinedsgpolicy1**, and **predefinedsgpolicy2** and cannot be changed in either the command-line interface or the management GUI. User-defined Safeguarded policies have IDs that start following the predefined Safeguarded policies.

An invocation example

```
mksafeguardedpolicy -backupunit day -backupinterval 1 -backupstarttime 2109010100
-retentiondays 7
```

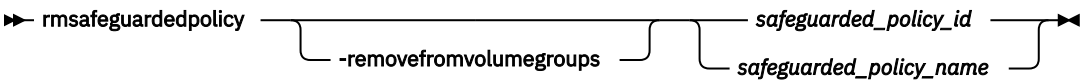
The resulting output:

```
Safeguarded Policy, id [3], successfully created
```

rmsafeguardedpolicy

Use the **rmsafeguardedpolicy** command to remove the existing Safeguarded policies.

Syntax



Parameters

-removefromvolumegroups
(Optional) Removes the volume group associations to this Safeguarded policy. This option is available to Security Administrator users only.

safeguarded_policy_id | safeguarded_policy_name
(Required) The ID or name of a Safeguarded policy to remove.

Note: The predefined Safeguarded policies cannot be removed.

An invocation example in which the security administrator removes a Safeguarded policy that has volume group associations:

```
rmsafeguardedpolicy -removefromvolumegroups 2
```

The resulting output:

```
No feedback.
```

Chapter 32. Service information commands

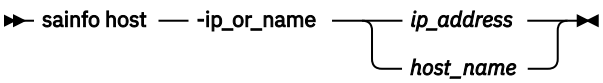
Use the service information commands to view hardware status and to report hardware errors.

Note: You can find out the *panel_name* for each of your nodes by issuing `sainfo lsservicenodes`.

host

Use the **host** command to change an Internet Protocol (IP) address to a host name or change a host name to an IP address.

Syntax



Parameters

-ip_or_name ***ip_address | host_name***
(Required) Specifies the host system IP address or host system name. The value for the IP address must be a standard IPv4 or IPv6 address. The value for the host name must be an alphanumeric string.

Description

This command changes an IP address to a host name or change a host name to an IP address.

This command handles Domain Name System (DNS) lookup and assists with DNS configuration.

Use this command to convert host names to IP addresses or IP address to host names.

An invocation example

```
sainfo host -ip_or_name vardy
```

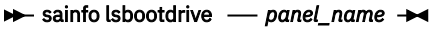
The resulting output:

```
Host vardy not found: 3(NXDOMAIN)
```

lsbootdrive

Use the **lsbootdrive** command to return the drive information of the internal boot drives for the specified node.

Syntax



Parameters

panel_name
(Optional) Identifies the node that is being used.

© Copyright IBM Corp. 2003, 2025

Description

The command displays information about internal boot drives for the specified node (if applicable).

This table provides the possible values for attributes that are displayed as data in the output views.

*Table 129. **lsbootdrive** attribute values*

Attribute	Value
panel_id	Identifies the panel ID of the node that contains the boot drive. The value is a 7-character alphanumeric string.
node_id	Identifies the ID (in decimal format) of the node that contains the boot drive.
node_name	Identifies the name of the node that contains the boot drive.
can_sync	Indicates when synchronization is needed and is not prevented by another problem. The values are yes or no. **Remember:** This value must be yes when status is out_of_sync. If the status is not out_of_sync the value is no.
slot_id	Identifies the ID (in decimal format) of the slot within the node.
booted	Indicates whether the node started from the specified drive. The values are yes, no.
status	Indicates the slot status. The values are: • missing indicates that the slot must be occupied, but the software cannot detect a drive. It also indicates the serial number of the expected drive. • empty indicates that the slot is supposed to be empty and is empty. • unsupported indicates that the slot is supposed to be empty but is not. • failed indicates that the drive in the slot is not working. • uninitialized indicates that the drive is not formatted for the system. • wrong_node indicates that the drive is working but is not in the correct node. It also indicates the serial number of the node the drive must be in and might indicate the serial number of the drive that must be in the slot. • wrong_slot indicates that the drive is working correctly - it is in the right node but the wrong slot. It also indicates which drives belong in which slots. • out_of_sync indicates that the drive is working correctly but must be resynchronized. Make sure the value of can_sync is yes before you specify the command chbootdrive -sync. • online indicates that the drive in the slot is working correctly. • unknown indicates that the node is not an active member of the system, and the state of the drive in that slot is unknown.
actual_drive_sn	Indicates the serial number of the drive in the slot. The value is an alphanumeric string or blank if no drive is in the slot.
configured_drive_sn	Identifies the serial number of the drive that must be in the slot. The value is an alphanumeric string or blank if no drive is in the slot.
actual_node_sn	Indicates the serial number of the node the drive (currently in the slot) belongs in. The value is an alphanumeric string or blank if no drive is in the slot.

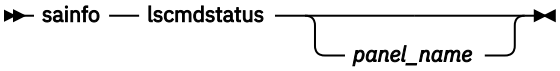
Table 129. **lsbootdrive** attribute values (continued)	
Attribute	**Value**
identify	Indicates whether `chnodebootdrive -identify` is specified. The following values are possible for this attribute: • on indicates `chbnodeootdrive -identify yes -slot` is specified • `off` indicates `chnodebootdrive -identify no -slot` is specified
FRU_part_number	Identifies the field-replaceable unit (FRU) part number of the drive. The value is a 7-character alphanumeric string or blank if no drive exists.
FRU_identity	Identifies the 11S number that combines the manufacturing part number and serial number. The value is a 22-character alphanumeric string.

Note: If the status is `out_of_sync` and `can_sync` is set to `no`, look for an indication to determine what must be fixed to make resynchronization possible.

lscmdstatus

Use the **lscmdstatus** command to display the status of any currently running service-aid task.

Syntax



Parameters

panel_name
The name of the panel. This command fails if the *panel_name* ID is not in the list that is returned by **lsservicenodes**.

Note: If `panel_name` is not supplied, this value applies to the node on which the command is running.

Description

This command displays the status of any currently running service-aid task. If no task is running, then the completion status of the last task is displayed. Any user can specify this CLI command.

If no service-aid tasks run since the node was last restarted, the command returns immediately with no output. Otherwise, it display something similar to example below.

This table provides the attribute values that can be displayed as output view data.

Table 130. **lscmdstatus** output	
Attribute	**Description**
T3_status	Indicates the T3 recovery status.
T3_status_data	Indicates the T3 recovery status activity.
cpfiles_status	Indicates the **cpfiles** command status.
cpfiles_status_data	Indicates the **cpfiles** command activity.
snap_status	Indicates the **snap** command status.
installsoftware_status	Indicates the **installsoftware** command status.

Table 130. **lscmdstatus** output (continued)	
Attribute	**Description**
supportupload_status	Indicates the upload activity status. The values are: • `active` indicates that the upload is ongoing. • `complete` indicates that the upload completed successfully. • `failed` indicates that the upload failed. • `abort` indicates that the upload is aborted. • `wait` indicates that the upload is underway but not completed.
supportupload_status_data	Indicates upload activity information (for example `uploading`).
supportupload_progress_percent	Indicates the upload progress percentage. The value is a number in the range 0 - 100.
supportupload_throughout_KBps	Indicates the upload speed in kilobytes per second (KBps). The value is a number (integer).
supportupload_filename	Indicates the file name. The standard length is 256 characters.
downloadsoftware_status	Indicates the download activity status. The values are: • `active` indicates that the download is ongoing. • `complete` indicates that the download completed successfully. • `failed` indicates that the download failed. • `abort` indicates that the download is aborted. .
downloadsoftware_status_data	Indicates the download activity information (for example `Downloading the bundle`).
downloadsoftware_progress_percent	Indicates the download progress percentage. The value is a number in the range 0 - 100.
downloadsoftware_throughput_KBps	Indicates the download speed in kilobytes per second (KBps). The value is a number (integer).
downloadsoftware_size	Indicates the total size of the selected bundle to download. The value must be a number (decimal) and must be in the format TiB, GiB, MiB, or KiB.

sainfo lsfcportfeature

Use the **lsfcportfeature** command to list the status of Fibre Channel port features.

Syntax

sainfo lsfcportfeature

Description

This command lists the status of the Fibre Channel port features. The Fibre Channel port supports various features like fabric device management interface (FDMI) registration, discovery, and port state. The port state is represented in a 64-bit bitmap, where each bit represents a physical adapter port. The rightmost bit indicating the first Fibre Channel I/O port ID. You can use **lsfcportfeature** command to list the

status of various Fibre Channel port features. For example, use **lsfcportfeature** command to check that FDMI registration and discovery is enabled or disabled.

This table provides the possible values that are applicable to the attributes that are displayed as data in the output views.

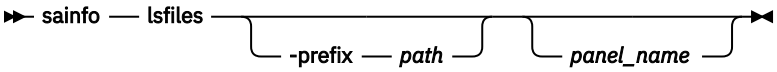
*Table 131. **lsfcportfeature** attribute values*	
Attribute	**Value**
feature	Indicates the name of the port feature.
status	Indicates the status of the port feature. The value is `enabled` or `disabled`.

Note: Some switches do not support all these parameters. For these switches, the respective parameters might be missing in switch CLI outputs.

lsfiles

Use the **lsfiles** command to display the files on the node that you want to retrieve with the **satask cpfiles** command.

Syntax



Parameters

panel_name
(Optional) The name of the panel. The command fails if the *panel_name*

ID is not in the list that is returned by the **lsservicenodes** command.

Note: If `panel_name` is not supplied, this parameter applies to the node on which the command is running.

-prefix*path*
(Optional) The path must exist in a permitted directory that supports lists. You can use the following **-prefix** paths:

- `/dumps` (the default if **-prefix** is not set)
- `/dumps/audit`
- `/dumps/cimom`
- `/dumps/cloud`
- `/dumps/configs`
- `/dumps/drive`
- `/dumps/easytier`
- `/dumps/elogs`
- `/dumps/enclosure`
- `/dumps/feature`
- `/dumps/iostats`
- `/dumps/iotrace`
- `/dumps/mdisk`
- `/dumps/syslogs`
- `/home/admin/update`

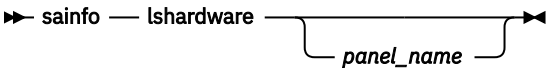
Description

This command displays a list of the files on the node that you want to retrieve by using the **satask cpfiles** command.

lshardware

Use the **lshardware** command to view the configured and actual hardware configuration of a node in the system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

panel_name
(Optional) The node panel name.

Note: If `panel_name` is not supplied, this parameter applies to the node on which the command is running.

Description

When the node is in a service state, use this command to view the current hardware configuration. This table provides the possible values that are applicable to the attributes that are displayed as data in the output views.

*Table 132. **lshardware** attribute values*

Attribute	Value
panel_name	Indicates the node panel name.
node_id	Indicates the node unique ID. The value is a number or blank if not in a system.
node_name	Indicates the node name. The value is an alphanumeric string or blank if not in a system.
node_status	Indicates the node status.
hardware	Indicates the hardware model, such as DH8.
actual_different	Indicates whether the node hardware is different than the configured hardware.
actual_valid	Indicates whether the node hardware is valid.
memory_configured	Indicates the configured amount of memory (in GB).
memory_actual	Indicates the currently installed amount of memory (in GB).
memory_valid	Indicates whether the actual memory is a valid configuration.
cpu_count	Indicates the maximum number of CPUs for the node.
cpu_socket	Indicates the ID of socket to which the CPU fields refer.

Table 132. **lshardware** attribute values (continued)	
Attribute	**Value**
cpu_configured	Indicates the configured CPU for this socket.
cpu_actual	Indicates the currently installed CPU in this socket.
cpu_valid	Indicates whether the currently installed CPU is a valid configuration.
adapter_count	Indicates the maximum number of adapters for the node (differs by node type).
adapter_location	Indicates the location of this adapter.
adapter_configured	Indicates the configured adapter for this location.
adapter_actual	Indicates the currently installed adapter for this location.
adapter_valid	Indicates whether the adapter in this location is valid.
ports_different	Indicates whether adapter ports can support more functions.

lsnodediscoverysubnet

Use the **lsnodediscoverysubnet** command to list the IPv4 subnets that are used for local node discovery.

Syntax

sainfo lsnodediscoverysubnet

Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

Description

This command returns a list of the 32-bit subnet nodes that used for local node discovery. Note: IPv6 is not supported. This table provides the possible output values.

Table 133. **lsnodediscoverysubnet** attribute values	
Attribute	**Value**
id	Displays the position of the subnet set in the subnet table.
subnet	Displays the subnet (32 bits) on which the target node is located.

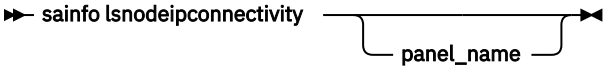
lsnodeip (Deprecated)

The **satask lsnodeip** command has been deprecated.

lsnodeipconnectivity

Use the **sainfo lsnodeipconnectivity** command to list the RDMA sessions.

Syntax



Parameters

panel_name
(Optional) Specifies the system or node on which the command is invoked. The value must be an alphanumeric string.

Description

This command lists the RDMA sessions.

This table provides the attribute values that can be displayed as output view data.

*Table 134. **sainfo lsnodeipconnectivity** output*

Attribute	Description
status	Indicates the status of the connection. A value of `Discovered` indicates that ports are configured, but the connection could not be established. The reason for the connection issue is described in the **error_data** field. A value of `Connected:protocol` indicates a successful connection.
local_port_id	Indicates the local port identifier for the node Ethernet port number that is displayed in the **lsservicestatus** output.
local_vlan	Indicates the local port Virtual LAN Identifier for the node Ethernet port number that is displayed in the **lsservicestatus** output.
local_rdma_type	Indicates the local port RDMA type.
local_ip_addr	Indicates the local IP address.
remote_port_id	Indicates the peer or target port identifier for the node Ethernet port number that is displayed in the **lsservicestatus** output.
remote_vlan	Indicates the remote port Virtual LAN Identifier for the node Ethernet port number that is displayed in the **lsservicestatus** output.
remote_rdma_type	Indicates the remote port RDMA type.
remote_ip_addr	Indicates the remote IP address.
remote_wwnn	Indicates the remote World Wide Node Name (WWNN).
remote_panel_name	Indicates the remote node unique panel name.
cluster_id	Indicates the Cluster ID of the peer system.
error_data	Displays information that details any problems with the connection.

The following table lists the values for the **error_data** field and possible corrective actions.

*Table 135. **sainfo lsnodeipconnectivity error_data***		
error_data value	**Description**	**Possible Corrective Action**
Protocol mismatch	Source port RDMA adapter and destination Port RDMA adapter are not compatible with each other.	Make sure that both of the RDMA adapters are the same type (Ex: RoCE to RoCE, iWARP to iWARP).
Unreachable	The source IP address and destination IP address are not reachable.	Make sure that both of the IPs are reachable. You can use the `satask ping` command to check the network connectivity of the local and remote ports.
Duplicate IP	A duplicate IP address is set.	Make sure that the IP addresses that are set are unique on the same network.
Degraded	Negotiated speed does not match. Degraded is set when a slower speed is selected than the adapter supported maximum speed.	Make sure that the switch adapters on both nodes support the maximum speed.
VLAN ID Mismatch	The local and remote port Virtual LAN Identifiers do not match.	Make sure that the VLAN IDs match.

lsservicenodes

Use the **lsservicenodes** command to display a list of all the nodes that can be serviced by using the service assistant CLI.

Syntax



Parameters

None

Description

This list includes nodes that are at a code level of at least 6.2.0, are visible on the fabric, and are one of the following candidates:

- The partner node in a control enclosure to the node that is running the command.
- In the same system as the node that is running the command.
- In candidate state.
- Not in a system and in service state.
- Not in an enclosure with a stored system ID (which is not the system ID of the local node).

Nodes participating in a different system are not be shown unless they are the partner node. This table shows possible outputs.

*Table 136. **lsservicenodes** outputs*	
Attribute	**Value**
panel_name	Indicates the front panel name, enclosure IDs, or canister IDs that identify the node. **Note:** You can find out the `panel_name` for each of your nodes by issuing **`lsservicenodes`**.
cluster_id	Indicates the system ID. Blank if node is a candidate; otherwise, the value is determined from vpd_cluster.
cluster_name	Indicates the system name. The value is blank if node is a candidate; otherwise, the value is determined from `vpd_cluster`.
node_id	Indicates the node ID. The value is blank if node is a candidate; otherwise, the value is determined from `vpd_cluster`.
node_name	Indicates the node name. The value is blank if node is a candidate; otherwise, the value is determined from `vpd_cluster`.
relation	Indicates the relationship. The values are: • `local` indicates that the node the CLI command was issued from. • `partner` indicates that the node is in the same enclosure as the local node. • `cluster` indicates that nodes other than the partner that is in the same system as the local node. • `candidate` indicates that the node is not part of the system.
node_status	Indicates the node status. • `Active` indicates that the node is part of a system and can perform I/O. • `Service` indicates that the node is in service, standby, or node rescue. • `Candidate` indicates that the node is not part of a system. • `Starting` indicates that the node is part of a system and is attempting to join the system, and cannot perform I/O.
error_data	Indicates the outstanding error and error data, by priority.

lservicerecommendation

Use the **`lservicerecommendation`** command to determine what actions must be completed when you service a node.

Syntax



Parameters

panel_name
(Optional) If no panel ID is provided, the service recommendation for the local node is returned. If a *panel_name* from the list that is returned by **`lsservicenodes`** is specified, then the service recommendation for that node is returned. The command fails if the *panel_name* is not in the list that is returned by **`lsservicenodes`**.

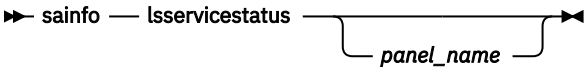
Note: If `panel_name` is not supplied, this applies to the node on which the command is running.

Description
This command determines what actions must be completed when you service a node.

lsservicestatus

Use the **lsservicestatus** command to display the status of a node.

Syntax



Parameters

panel_name
(Optional) If a *panel_name* is not provided, the service recommendation for the local node is returned. If a *panel_name* from the list that is returned by **lsservicenodes** is specified, then the service recommendation for that node is returned. The command fails if the *panel_name* ID is not in the list that is returned by **lsservicenodes**. This output is returned as the node status on all Universal Serial Bus (USB) flash drive commands.

Description

Use this command to display the status of a node. This command provides all the information that can be obtained by using the front panel of a system node.

You can run this command on any node (even one that is not part of a system) to obtain the vital product data (VPD) and error status.

This table shows the possible outputs.

*Table 137. **lsservicestatus** output*

Attribute	Value
panel_name	Indicates the front panel name, enclosure IDs, or canister IDs that identify the node.
cluster_id	Indicates the ID of a clustered system.
cluster_name	Indicates the name of a system. When you use this parameter, the detailed view of the specific system is displayed, and any value that you specified by the **-filtervalue** parameter is ignored. If you do not specify the *cluster_name* parameter, the concise view of all systems that match the filtering requirements that are specified by the **-filtervalue** parameter are displayed.
cluster_status	Indicates that the error code is the same as the one displayed on the front panel.
cluster_ip_count	Indicates the maximum number of management addresses that you can configure.
cluster_port	Indicates the system IP port. This field is repeated for each management address.
cluster_ip	Indicates the Internet Protocol Version 4 (IPv4) management IP address. This field is repeated for each management address.
cluster_gw	Indicates the IPv4 management IP gateway. This field is repeated for each management address.

*Table 137. **lsservicestatus** output (continued)*	
Attribute	**Value**
cluster_mask	Indicates the IPv4 management IP mask. This field is repeated for each management address.
cluster_ip_6	Indicates the Internet Protocol Version 6 (IPv6) management IP address. This field is repeated for each management address.
cluster_gw_6	Indicates the IPv6 management IP gateway. This field is repeated for each management address.
cluster_prefix_6	Indicates the IPv6 management IP prefix. This field is repeated for each management address.
node_id	Indicates the ID of the node that is being configured. The value is a number.
node_name	Indicates the name of the node that is being configured.
node_status	Indicates the node status. The values are: • `active` indicates that the node is part of a system and can perform I/O. • `service` indicates that the node is in service mode, standby mode, or node rescue mode. • `candidate` indicates that the node is not part of a system. • `starting` indicates that the node is part of a system and is attempting to join the system. • `spare` indicates that the node is a spare node.
config_node	Indicates whether it is a configuration node. The values are `yes` or `no`.
hardware	Indicates the hardware type.
service_IP_address	Indicates the IPv4 service address for the node.
service_gateway	Indicates the IPv4 service gateway for the node.
service_subnet_mask	Indicates the IPv4 service mask for the node.
service_IP_address_6	Indicates the IPv6 service address for the node.
service_gateway_6	Indicates the IPv6 service gateway for the node.
service_prefix_6	Indicates the IPv6 service prefix for the node.
service_IP_mode	Indicates the current mode of IPv4 service address. Modes include: • `blank` • `static` • `dhcp` - Indicates that the address was obtained from a DHCP server. • `dhcpfallback` - Indicates that DHCP failed and the default static address is in use.
service_IP_mode_6	Indicates the current mode of IPv6 service address. Modes include: • `blank` • `static` • `dhcp` - Indicates that the address was obtained from a DHCP server.
machine_part_number	Indicates the machine part number of the enclosure.

*Table 137. **lsservicestatus** output (continued)*	
Attribute	**Value**
node_machine_part_number_copy	Indicates the machine part number of the home enclosure of the node.
node_code_version	Indicates the version of system code for the node.
node_code_build	Indicates the build string for code on the node.
cluster_code_build	Indicates the CSM build that the system is running.
node_error_count	Indicates the number of node errors.
error_code	Indicates the error code.
error_data	Indicates the type of node errors.
fC_ports	Indicates the number of FC ports. This field is repeated for each management address.
port_id	Indicates the port ID. This field is repeated for each management address.
port_status	Indicates the port status. This value must match the port on the front panel, enclosure, or canister.
port_speed	Indicates the port speed. This value must match the port speed on the front panel, enclosure, or canister.
port_WWPN	Indicates the worldwide port number of the port.
SFP_type	Indicates the SFP type. The values are `long-wave` or `short-wave`.
rdma_type	A semicolon-separated list that displays the RDMA protocol capabilities of the port. **Note:** This field is blank for non-RDMA capable Ethernet ports.
ethernet_ports	Indicates the number of detected Ethernet ports.
ethernet_port_id	Indicates the ID of an Ethernet port.
port_status	Indicates the port status. The values are: • `online` • `offline` • `not configured` • `disabled`
port_speed	Indicates the port speed. The values are: • `10 Mbps` • `100 Mbps` • `1 Gbps` • `10 Gbps` • `25 Gbps` • `full` • `half`
MAC	Indicates a single MAC address.

*Table 137. **lsservicestatus** output (continued)*	
Attribute	**Value**
node_IP_address	Indicates the IPv4 management node IP address, which the system uses for node discovery and IP clustering.
node_gateway	Indicates the IPv4 management node IP gateway.
node_subnet_mask	Indicates the IPv4 management node IP mask.
product_mtm	Indicates the machine type and model.
product_serial	Indicates the node serial number.
disk_WWNN_prefix	The most recently used WWNN prefix.
dump_name	Indicates the dump name of the node.
node_WWNN	Indicates the last active WWNN stored in the node; blank if no system data.
disk_WWNN_suffix	The most recently used WWNN suffix.
panel_WWNN_suffix	Indicates the WWNN suffix stored in the front panel.
UPS_serial_number	Indicates the serial number of attached UPS.
UPS_status	Indicates the UPS status of the nodes.
enclosure_WWNN_1	Indicates the first canister WWNN from the enclosure VPD.
enclosure_WWNN_2	Indicates the second canister WWNN from the enclosure VPD.
node_part_identity	Indicates the 11S string from the hardware VPD.
node_FRU_part	Indicates the node field-replaceable unit (FRU) part number, if stored in node VPD
enclosure_identity	Indicates the S11 data.
PSU_count	Indicates the number of expected PSUs (two).
PSU_id	Indicates the ID of the slot the PSU is in. The value is a number.
PSU_status	Indicates the PSU status. The status is `missing`, `failed`, or `active`.
time_to_charge	Indicates the estimated time (in minutes) needed for 50% of the battery to be charged.
Battery_charging	Indicates the percentage of charge of the batteries.
Battery_count	Indicates the number of expected batteries (two).
Battery_id	Indicates the ID of the slot the battery is in.
Battery_status	Indicates that the status is `missing`, `failed`, `charging`, or `active`.
Battery_FRU_part	Indicates the FRU part number of the battery.
Battery_part_identity	Indicates the 11S FRU identity of the battery (includes the serial number).
Battery_fault_led	Indicates thee fault light-emitting diode (LED) status.
Battery_charging_status	Indicates the battery charge status.
Battery_cycle_count	Indicates the number of charge or discharge cycles that are performed by the battery.

*Table 137. **lsservicestatus** output (continued)*	
Attribute	**Value**
Battery_powered_on_hours	Indicates the number of hours the battery is in powered mode.
Battery_last_recondition	Indicates the system time stamp of the last successful gas gauge calibration.
Battery_midplane_FRU_part	Indicates the FRU part number of the battery midplane.
Battery_midplane_part_identity	Indicates the 11S FRU identity of the battery midplane (which includes the serial number).
Battery_midplane_FW_version	Indicates the firmware version that is running on the battery midplane.
Battery_power_cable_FRU_part	Indicates the FRU part number of the battery power cable.
Battery_power_sense_cable_FRU_part	Indicates the FRU part number of the battery midplane power sense cable.
Battery_comms_cable_FRU_part	Indicates the FRU part number of the battery midplane communication cable.
Battery_EPOW_cable_FRU_part	Indicates the FRU part number of the battery midplane EPOW cable.
node_location_copy	Indicates the node location copy status. This value is equivalent to the panel name and is blank if a node is removed from a system.
node_product_mtm_copy	Indicates the node product machine type copy status. This value is equivalent to panel product_mtm; blank if a node is removed from a system.
node_product_serial_copy	Indicates the node product serial copy status. This value is equivalent to product_serial; blank if a node is removed from a system.
node_WWNN_1_copy	Indicates the node WWNN first copy status. This value is equivalent to enclosure_WWNN_1; blank if a node is removed from a system.
node_WWNN_2_copy	Indicates the node WWNN second copy status. This value is equivalent to enclosure_WWNN_2; blank if a node is removed from a system.
latest_cluster_id	Indicates the system ID running on the current enclosure; blank if a node is removed from a system.
next_cluster_id	Indicates the system ID running on the current enclosure; blank if a node is removed from a system.
console_ip	Indicates an Internet Protocol (IP) Version 4 or 6 address. **Note:** This field might be blank if the node is not present in a system.
has_nas_key	Indicates the has_nas_key field value. The values are yes or no. **Note:** This field might be blank if the node is not present in a system.
fc_io_ports	Indicates the number of configured Fibre Channel (FC) input/output (I/O) ports on the node. The value is a number. This field is repeated for each management address.
fc_io_port_id	Indicates the ID of the FC I/O port. The value is a number. This field is repeated for each management address.

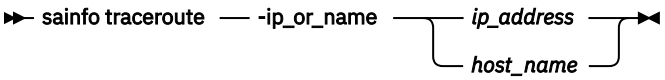
*Table 137. **lsservicestatus** output (continued)*	
Attribute	**Value**
fc_io_port_WWPN	Indicates the Worldwide port name (WWPN) of the FC I/O port. The value is 16 hexadecimal characters (lowercase). This field is repeated for each management address.
fc_io_port_switch_WWPN	Indicates the WWPN of the fabric port name. The value is 16 hexadecimal characters (lowercase) or all zeros if never active. This field is repeated for each management address.
fc_io_port_state	Indicates whether the port is working - `active` or `inactive`. This field is repeated for each management address.
fc_io_port_FCF_MAC	Indicates the FCF MAC address for active Ethernet. The value is blank for inactive or broken Ethernet, or N/A for FC type. This field is repeated for each management address.
fc_io_port_vlanid	Indicates the virtual local area network (VLAN) ID being used. The value is three hexadecimal digits for previously active or inactive Ethernet, blank for Ethernet that is never active, or N/A for FC type. This field is repeated for each management address.
fc_io_port_type	Indicates the type of the platform port for the FC I/O port. This field is repeated for each management address.
fc_io_port_type_port_id	Indicates the platform port type index of the platform port supplies the FC I/O port. Use both `fc_io_port_type_port_id` and `fc_io_port_type` to find the platform port. This field is repeated for each management address.
local_fc_port_mask	Indicates the FC I/O ports that a system can use for node-to-node communications on a local system if those FC I/O ports exist on a node. The value is 64 binary bits.
partner_fc_port_mask	Indicates the FC I/O ports that a system can use for system-to-system communications on a partnered system if those FC I/O ports exist on a node. The value is 64 binary bits.
topology	Indicates the system topology (set by using the **chsystem** command).
site_id	Indicates the site node value.
site_name	Indicates the site name.
identify_LED	Indicates that the node or node canister identify LED state (on, `off`, or blank).
password_reset_enabled	Indicates whether the superuser password reset is enabled (`yes` or no).
techport	Indicates the technician port (for emergency node maintenance): • `permanent` indicates a dedicated technician port. • `enabled` indicates an enabled technician port. • `disabled` indicates a disabled technician port. • none indicates that there is no technician port.
fcportmap	Indicates a mapping string of FC ports.
node_usb	Indicates whether the USB ports are enabled for use (**on**) or disabled (**off**).
superuser_locked	Specifies whether the superuser is locked. The value is either yes or no.

*Table 137. **lsservicestatus** output (continued)*	
Attribute	**Value**
superuser_multi_factor	Specifies whether multi-factor authentication is enabled for the superuser. Value can be *yes* or *no*. **Note:** The field attribute `superuser_multi_factor` indicates whether the command `chsecurity -superusermultifactor yes` has been used to indicate that superuser should be prompted for second factor authentication.
superuser_password_sshkey_required	Specifies whether superuser should provide both password and SSH public key during authentication. Value can be *yes* or *no*. **Note:** The field attribute `superuser_password_sshkey_required` indicates whether the command `chsecurity -superuserpasswordkeyrequired yes` has been used to indicate that superuser should have both a password and an SSH key configured for authentication.
superuser_gui_disabled	Specifies whether GUI access is disabled for superuser. Value can be *yes* or *no*. **Note:** The field attribute `superuser_gui_disabled` indicates whether the command `chsecurity -disablesuperusergui yes` has been used to indicate that superuser has disabled GUI access.
superuser_rest_disabled	Specifies whether REST-API access is disabled for superuser. Value can be *yes* or *no*. **Note:** The field attribute `superuser_rest_disabled` indicates whether the command `chsecurity -disablesuperuserrest yes` has been used to indicate that superuser has disabled REST API access.
superuser_cim_disabled	(Deprecated) CIMOM is disabled for all users and cannot be re-enabled.
system_cos	Indicates the system-wide priority tag that is used by system traffic. The value is in the range `0` to `7`, where `0` is the default value.
node_patch_count	Indicates the number of patches that are installed, obsolete, or in an error state. If the patch count is greater than `0`, patch information like patch name, patch description, and patch state are shown for each patch.
patch_name	Indicates the name of patch file.
patch_description	Indicates the short description of the patch file.
patch_status	Indicates the installation state of the patch: • `Installed` indicates that the patch is installed on the system. • `Obsolete` indicates that the patch was installed in the previous update and the new update does not require it. • `Error` indicates that the patch failed to install.
recovery_key_name	Indicates the name of the committed recovery key.

traceroute

Use the **traceroute** command to print wide area network (WAN) route packets to a specified host.

Syntax



Parameters

-ip_or_name *ip_address* | *host_name*
(Required) Specifies the host system IP address or host system name that you want to print WAN route packets to. The value for the IP address must be a standard IPv4 or IPv6 address. The value for the host name must be an alphanumeric string.

Description

This command prints WAN print route packets to a specified host.

An invocation example

```
sainfo traceroute -ip_or_name compass.ssd.hursley.ibm.com
```

The resulting output:

```
traceroute to compass.ssd.hursley.ibm.com (9.71.44.59), 30 hops max, 60 byte packets
  1  9.71.45.4 (9.71.45.4)  0.283 ms  0.337 ms  0.397 ms
  2  compass.ssd.hursley.ibm.com (9.71.44.59)  0.124 ms  0.124 ms  0.113 ms
```

Chapter 33. Service mode commands (Discontinued)

The service mode commands are discontinued.

applysoftware (Discontinued)

Attention: The **svcservicemodetask applysoftware** command is discontinued. Use the **satask installsoftware** command instead.

Discontinued.

cleardumps (Discontinued)

The **svcservicemodetask cleardumps** command is discontinued. Use the **cleardumps** command instead.

Discontinued.

dumperrlog (Discontinued)

The **svcservicemodetask dumperrlog** command is discontinued. Use the **dumperrlog** command instead.

Discontinued.

exit (Discontinued)

The **svcservicemodetask exit** command is discontinued. Use the **satask stopservice** command instead.

© Copyright IBM Corp. 2003, 2025

Chapter 34. Service mode information commands (Discontinued)

The service mode information commands are discontinued.

ls2145dumps (Discontinued)

The **svcservicemodeinfo ls2145dumps** command is discontinued. Use the **lsdumps** command to display a list of files in a particular dumps directory.

lscimomdumps (Discontinued)

The **svcservicemodeinfo lscimomdumps** command is discontinued. Use the **lsdumps** command to display a list of files in a particular dumps directory.

lsclustervpd (Discontinued)

Attention: The **svcservicemodeinfo lsclustervpd** command is discontinued. Use the **sainfo lsservicestatus** command instead.

lserrlogdumps (Discontinued)

The **svcservicemodeinfo lserrlogdumps** command is discontinued. Use the **lsdumps** command to display a list of files in a particular dumps directory.

lsfeaturedumps (Discontinued)

The **svcservicemodeinfo lsfeaturedumps** command is discontinued. Use the **lsdumps** command to display a list of files in a particular dumps directory.

lsiostatsdumps (Discontinued)

The **svcservicemodeinfo lsiostatsdumps** command is discontinued. Use the **lsdumps** command to display a list of files in a particular dumps directory.

lsiotracedumps (Discontinued)

The **svcservicemodeinfo lsiotracedumps** command is discontinued. Use the **lsdumps** command to display a list of files in a particular dumps directory.

lsmdiskdumps (Discontinued)

The **svcservicemodeinfo lsmdiskdumps** command is discontinued. Use the **lsdumps** command to display a list of files in a particular dumps directory.

lssoftwaredumps (Discontinued)

The **svcservicemodeinfo lssoftwaredumps** command is discontinued. Use the **lsdumps** command to display a list of files in a particular dumps directory.

© Copyright IBM Corp. 2003, 2025

Chapter 35. Service task commands

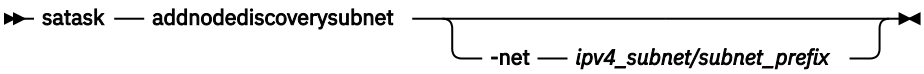
Use the service task commands to service the node hardware (such as IBM Storage Virtualize).

Note: You can find out the *panel_name* for each of your nodes by issuing `satask lsservicenodes`.

addnodediscoverysubnet

Use the **addnodediscoverysubnet** command to add a node discovery subnet where candidate nodes can be discovered.

Syntax



Parameters

-net
(Optional) Specify the subnet and network prefix notation on which the target node is located.

Description

This command adds a node discovery subnet where candidate nodes can be discovered. There are a total of 16 slots (16 records) in the subnets table. You do not need to specify which slot in the table to store the record. This command finds an available slot automatically for you.

If you add a subnet that has an overlapping IP address with other subnets that are configured on the node or add the same subnet twice, the command fails with an error message.

An invocation example

```
satask addnodediscoverysubnet -net 9.7.8.0/28
```

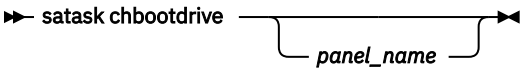
The resulting output:

```
No feedback
```

chbootdrive

Use the **chbootdrive** command to synchronize a broken drive or field-replaceable unit (FRU) replacement drive or acknowledge an intentional swap over of boot drives into a replacement canister.

Syntax



Parameters

-sync
(Required) Specifies synchronization for:

- New drives

© Copyright IBM Corp. 2003, 2025

- Drives marked out of synchronization
- Drives from another node

panel_name
(Optional) Specifies the panel name of the node on which to operate.

Note: If `panel_name` is not supplied, this command applies to the node on which the command is running.

Description
The command synchronizes a broken or FRU replacement drive. Specifying the **-sync** causes a node restart on the specified node.

An invocation example

```
satask chbootdrive -sync
```

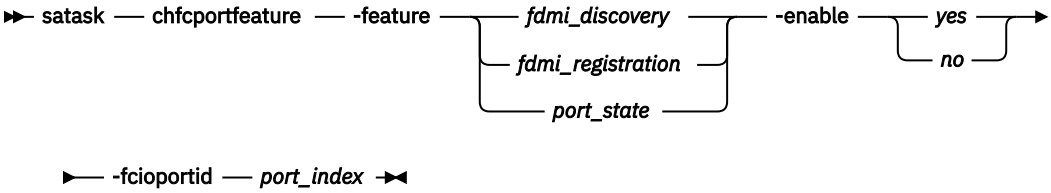
The following output is displayed:

```
No feedback.
```

satask chfcportfeature

Use the **chfcportfeature** command to change the status of the Fibre Channel port features.

Syntax



Parameters

-feature ***fdmi_discovery | fdmi_registration| port_state***
(Required) Specifies the name of the Fibre Channel port feature.

-enable ***yes | no***
(Required) Specifies whether the Fibre Channel port feature is enabled.

-fcioportid
Specifies the ID of the Fibre Channel I/O port, only when the feature is port_state.

Description

This command can enable or disable the Fibre Channel port features. The Fibre Channel port supports features like fabric device management interface (FDMI) registration, discovery, and port state. You can use **chfcportfeature** command to enable or disable the Fibre Channel port features.

An invocation example

```
satask chfcportfeature -feature fdmi_discovery -enable yes
```

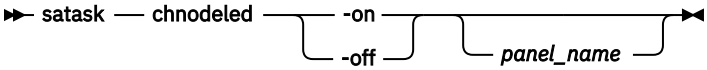
The resulting output:

```
No feedback
```

chnodeled

Use the **chnodeled** command to turn the identification light-emitting diode (LED) on or off for the specified node or control canister.

Syntax



Parameters

-on | -off
(Required) Turns the identification LED for the specified node or control canister on or `off`.

panel_name
(Optional) Specifies a unique panel name used to apply the command to a valid node on the fabric.

Note: If `panel_name` is not supplied, this applies to the node on which the command is running.

Description

This command turns the canister identification LED on or off.

Note: The identification LED is mapped onto the physical LEDs using different methods, depending on your hardware. Refer to the documentation for your hardware platform for more information.

An invocation example to turn on the LED to identify the battery in slot 2 on node KP2812

```
satask chnodeled -on -battery 2 KP2812
```

The resulting output:

```
No feedback
```

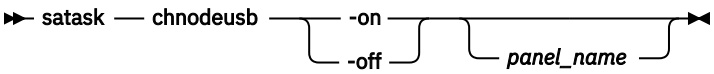
chnodeip (Discontinued)

The **satask chnodeip** command has been discontinued.

chnodeusb

Use the **chnodeusb** command to enable or disable USB ports on the node.

Syntax



Parameters

-on | -off
(Required) Setting this parameter to on enables USB ports on the node. Setting it to `off` disables USB ports on the node.

panel_name
(Optional) Specifies a unique panel name that is used to apply the command to a valid node on the system.

Note: If `panel_name` is not supplied, this command applies to the node on which the command is running.

Description

This command allows a superuser to change security features on a node. This command cannot be run on active nodes.

An invocation example for `-on/-off`

```
satask chnodeusb -on/-off
```

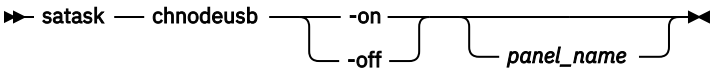
The resulting output:

```
No feedback.
```

chnodeusb

Use the **chnodeusb** command to enable or disable USB ports on the node.

Syntax



Parameters

`-on | -off`
(Required) Setting this parameter to on enables USB ports on the node. Setting it to `off` disables USB ports on the node.

panel_name
(Optional) Specifies a unique panel name that is used to apply the command to a valid node on the system.

Note: If `panel_name` is not supplied, this command applies to the node on which the command is running.

Description

This command allows a superuser to change security features on a node. This command cannot be run on active nodes.

An invocation example for `-on/-off`

```
satask chnodeusb -on/-off
```

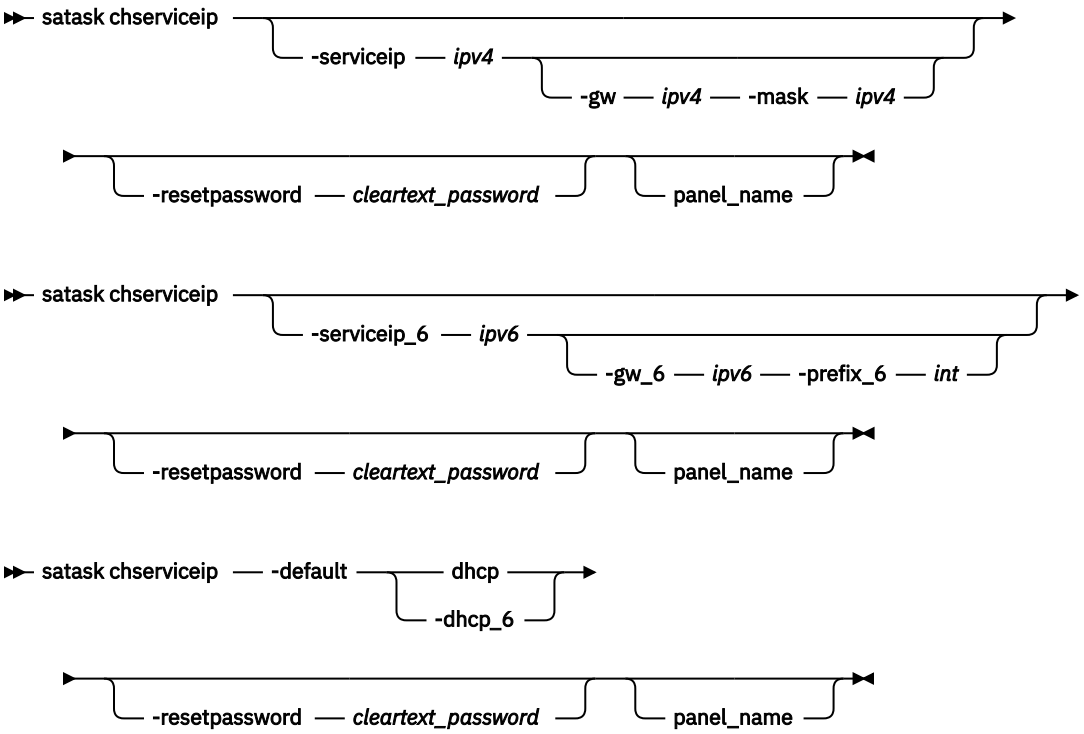
The resulting output:

```
No feedback.
```

satask chserviceip

Use the **satask chserviceip** command to set the service address for a specific node.

Syntax



Parameters

-satask
(Required) Specifies a system administrator task, such as issuing service commands that are only used in specific circumstances.

-default
(Optional) Resets the Internet Protocol Version 4 (IPv4) service address to the default address.

panel_name
(Optional) Specifies the node being serviced.

-serviceip
(Required) Specifies the IPv4 address for the service assistant.

Note: The IPv4 service address can be unconfigured by setting the address to `0.0.0.0`.

-gw
(Optional) Specifies the IPv4 gateway for the service assistant. If **-gw** is specified, **-mask** must be specified.

-mask
(Optional) Specifies the IPv4 subnet for the service assistant. If **-mask** is specified, **-gw** must be specified.

-serviceip_6 ***yes | no***
(Required) Specifies the Internet Protocol Version 6 (IPv6) address for the service assistant.

Note: The IPv6 service address can be unconfigured by setting the address to `0:0:0:0:0:0:0:0`.

-gw_6
(Optional) Specifies the IPv6 gateway for the service assistant. If **-gw_6** is specified, **-prefix_6** must be specified.

-prefix_6
(Optional) Specifies the IPv6 prefix for the service assistant. If **-prefix_6** is specified, **-gw_6** must be specified.

-default
(Optional) Resets to the default IPv4 address.

-resetpassword ***cleartext_password***
(Optional) Sets the service assistant password to **cleartext_password**. The password must be in the range of 6 - 64 characters.

-dhcp
(Optional) Attempts to obtain an IPv4 address from Dynamic Host Configuration Protocol (DHCP).

-dhcp_6
(Optional) Attempts to obtain an IPv6 address from DHCP.

Description

This command sets the service assistant IP address for a specific node. If the node is part of a system, then the system gateway, subnet, and prefix are used unless specified otherwise. If the node is a candidate node, the subnet, prefix, and gateway must be specified. If you specify an IPV4 or IPV6 address, but do not provide a gateway, mask, or prefix, then the existing gateway, mask, and prefix values are preserved.

When **-dhcpfallback** is specified, the current service interface is restarted and the new service IPv4 address is established using DHCP. If the DHCP request fails, the service IP address is set statically based on the node's physical location.

The location is based on the:

- Enclosure location within the chassis
- Node slot within the enclosure

Do not use the **-dhcpfallback** parameter for IPv6. These flags allocate a new address because the command causes the service interface to restart.

Therefore, you can configure both an IPv4 and an IPv6 address concurrently. Setting the IPv4 address does not change the IPv6 setting, and setting the IPv6 address does not change the IPv4 setting.

It is possible to clear any values set by setting the IPv4 address to 0.0.0.0 or leaving the IPv6 value empty.

If the **-resetpassword** parameter is provided, the superuser password for the system is reset to the value provided. The superuser password cannot be reset back to the default password. If the node or canister is active in a system, the superuser password for the system is reset; otherwise, the superuser password is reset on the individual node canister. If the node canister later becomes active in a system, the superuser password is reset to that of the system. You can configure the system to disable resetting the superuser password, see “setpwdreset” on page 216. If you disable that function, this action fails.

Use the **chserviceip** command to:

- Clear the IPv4 service IP address:

```
satask chserviceip -serviceip 0.0.0.0 -gw 0.0.0.0 -mask 0.0.0.0
```

- Clear the IPv6 service IP address:

```
satask chserviceip -serviceip_6 0::0 -gw_6 0::0 -prefix_6 64
```

Remember:

– If **-gw** is specified, **-mask** must also be specified.

– If **-gw_6** is specified, **-prefix_6** must also be specified.

An invocation example using specific -serviceip, -gw, and -mask parameters

```
satask chserviceip -serviceip 1.2.3.4 -gw 1.2.3.1 -mask 255.255.255.0
```

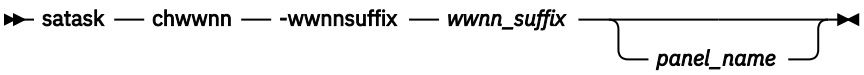
The resulting output:

```
No feedback
```

chwwnn

Use the **chwwnn** command to modify the node worldwide node name (WWNN). (This command applies to SAN Volume Controller 2145-CG8 and older nodes.)

Syntax



Parameters

-wwnnsuffix *wwnn_suffix*
(Required) Specifies the suffix to be used for node *wwnn*.

panel_name
(Optional) Specifies the node that you are servicing.

Note: If `panel_name` is not supplied, this parameter value applies to the node on which the command is running.

Description

This command modifies the WWNN. Use the **lsservicestatus** command to view suggested WWNNs.

An invocation example

```
chwwnn -wwnnsuffix 000cc
```

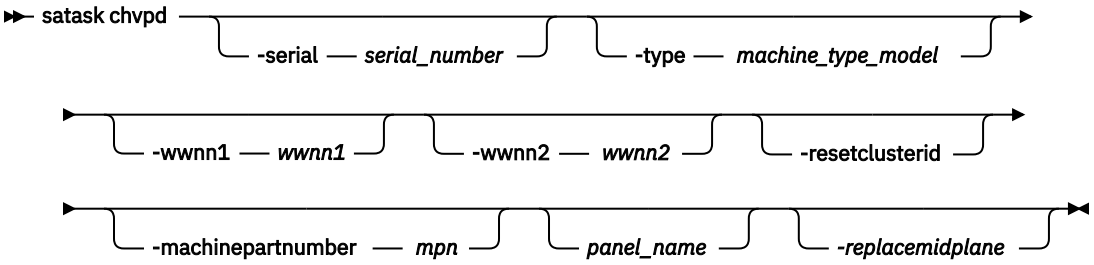
The resulting output

```
No feedback
```

chvpd

Use the **chvpd** command to set vital product data (VPD) such as serial number and machine type.

Syntax



Parameters

-serial *serial_number*
(Optional) Specifies the serial number for system board or enclosure.

-type *machine_type_model*
(Optional) Specifies the machine model type.

-wwnn1 *wwnn1*
(Optional) Specifies the WWNN of canister 1.

Note: **wwnn1** will no longer work on node 2, when the **panel_name** is not present.

-wwnn2 *wwnn2*
(Optional) Specifies the WWNN of canister 2.

Note: **wwnn2** will no longer work with node 1, when the **panel_name** is not present.

-machinepartnumber *mpn*
(Optional) Specifies the machine part number.

-resetclusterid
(Optional) Resets the system ID to 0.

Note:

- **resetclusterid** will only work on the current node it is on or the one specified by **panel_name**.
- The **resetclusterid** flag will no longer warm start the partner node.
- The partner cluster ID is ignored when clustering.

panel_name
(Optional) Identifies the node that is used.

-replacemidplane
(Optional) Fills in the details for the FRU part midplane replacement. Using this parameter fills in the following fields:

- The MTM in the midplane VPD, if empty.
- The serial number in the midplane VPD, if empty.
- The personality ID in the midplane VPD, if emtpy.

Note: There must be a location node error active on the node for this command to be successful.

Description

The command sets vital product data (VPD), including the serial number, WWPN, machine type, and system ID.

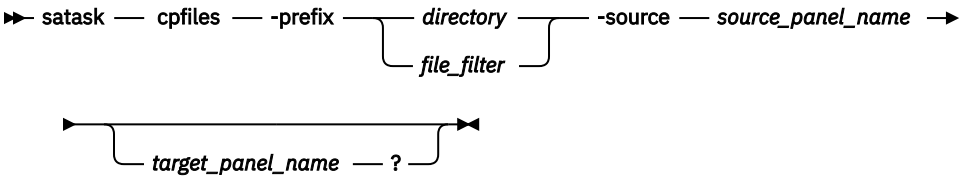
Remember: When you specify **chvpd**, one or more nodes might experience a reset or a restart. For example, these parameters might cause a restart:

- **-fcportmap** resets a node or one node canister
- **-idfile** resets a node
- **-serial** restarts both nodes or node canisters
- **-type** restarts both nodes or node canisters
- **-wwnn** resets a node or one node canister

cpfiles

Use the **cpfiles** command to copy files from another node.

Syntax



Parameters

satask
System Administrator task; service commands that are only used in specific circumstances.

-prefix ***directory | file_filter***
(Required) Specifies the directory, files, or directory and files to be retrieved. The path must exist in a permitted listable directory. You can use the following `-prefix` filters:

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

Note:

- You can also specify a file filter. For example, if you specify `/dumps/elogs/*.txt`, all files in the `/dumps/elogs` directory that end in *.txt* are copied.
- If you use a wildcard, the following rules apply:
  1. The wildcard character is an asterisk (*).
  2. The command can contain a maximum of one wildcard.
  3. When you use a wildcard, you must surround the filter entry with double quotation marks ("*x*"). For example: `satask cpfiles -prefix "/dumps/elogs/*.txt"`

-source ***source_panel_name***
(Required) Identifies the source node that the files are copied from.

target_panel_name
(Optional) Identifies the node that files are copied to. If no panel name is provided, the files are copied to the local node.

Note: If `target_panel_name` is not supplied, this applies to the node on which the command is running.

Description

This command copies files from another node. You can monitor the progress of the copy using the **`sainfo lscmdstatus`** command.

An invocation example to copy configuration information from canister 1 to enclosure 2

```
satask cpfiles -prefix /dumps/configs -source 02-1
```

The resulting output:

```
No feedback
```

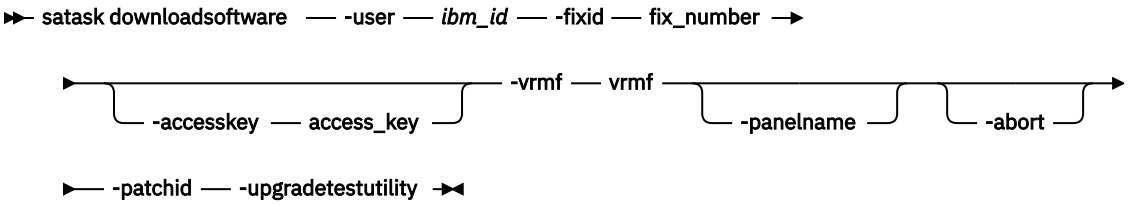
satask downloadsoftware

Use the **`satask downloadsoftware`** command to download the selected code bundles either through Cloud Call Home (and optionally via an HTTP proxy) or using FTP over SSH (SFTP). Use this command to download the software patch files through Call Home Using the Restful API. This command can also be used to abort a download.

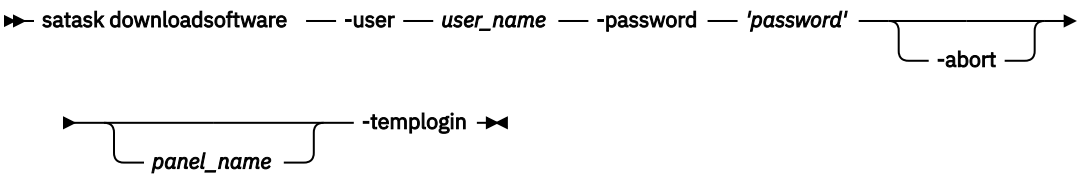
Syntax

The system supports two ways of downloading software packages from Fix Central. New for 8.4.2, you can download software packages from Fix Central using Cloud Call Home through `esupport.ibm.com`, which uses an IBMid user and password login to download software. The second method downloads packages using SFTP from Fix Central. The previous method of downloading software upgrade packages from Fix Central over SFTP is still supported, but is deprecated on systems running 8.4.2 or later.
The first syntax diagram shows downloading software from Fix Central using Cloud Call Home through `esupport.ibm.com`. The second syntax diagram shows downloading a software package from Fix Central using SFTP.

Using Cloud Call Home (new for 8.4.2)



Using SFTP from Fix Central (deprecated in 8.4.2)



Parameters

Using Cloud Call Home

`-user` ***ibm_id***
(Mandatory) Specifies the IBMid of the person who is downloading the software.

-fixid *fix_number*
(Mandatory) Specifies the fix to be downloaded from the Fix Central server. The value should be any string containing printable characters such as StorageDisk-2077-8.2.1.2 or Storage_Disk-V5000E-8.3.1.2-ifix7.

Note: You can use either **-patchid** or **-upgradetestutility** or **-fixid** or **-vrmf**. All of these are mutually exclusive parameters.

-vrmf *vrmf*
(Mandatory) Specifies the Version, Release, Modification, Fix of Storage Virtualize software to be downloaded from the Fix Central server. Four numeric values separated by three dots.

-patchid
(Mandatory) Specifies the name of the patch file that is to be downloaded through Call Home using Restful API.

-upgradetestutility
(Mandatory) Specifies that the upgrade test utility is to be downloaded.

-accesskey *access_key*
(Optional) Specifies the access key needed when downloading an ifix from Fix Central. The value must be an alphanumeric string that is 6 - 64 printable characters in length.

Note: This parameter can be used only with **fixid**.

-panelname
(Optional) Specifies the node that is being serviced. The value must be an alphanumeric string.

Note: If not specified, this variable applies to the node on which the command is running.

-abort
(Optional) Specifies that the download is canceled.

Using SFTP from Fix Central (deprecated in 8.4.2)

-user *user_name*
(Optional) Specifies a temporary user that is created by the Fix Central server to download an existing code bundle. The value must be an alphanumeric string that is 6 - 64 printable characters in length.

-password *password*
(Optional) Specifies a temporary password that is created by the Fix Central server for the download that is being performed by the temporary user. The value must be an alphanumeric string that is 6 - 64 printable characters in length. You must enclose the password in single quotation marks.

-abort
(Optional) Specifies that the download is canceled.

panel_name
(Optional) Specifies the node that is being serviced. The value must be an alphanumeric string.

Note: If not specified, this variable applies to the node on which the command is running.

-temploginn
(Mandatory) Indicates that the user and password parameters are generated by Fix Central in order to download a code bundle.

Note: This parameter is mandatory unless **-abort** parameter is used.

Description

This command downloads selected code bundles from the Fix Central server using either Cloud Call Home or SFTP. Use this command to download the software patch files through Call Home Using the Restful API. This command can also be used to abort an ongoing download.

Remember: Before you specify this command, you must:

1. Have internet access for all nodes
2. Configure service IP addresses on all nodes in your system.

3. Configure a DNS server on your system by specifying **mkdnsserver** to define a Domain Name System (DNS) server and **lsdnsserver** to display its values.
4. Depending on whether you are using an internal HTTP proxy or a firewall to manage remote access, additional network configuration might be required. For more information, see the topic, "Configuring support assistance" in the IBM Documentation for your product.
5. If you are using SFTP to download software packages from Fix Central, log on to the Fix Central server and configure the bundle that you want to download for your product. Fix Central servers prepare your download bundle and create a temporary username and password for you. These login credentials are valid for up to 72 hours (during this time you must finish downloading the bundle on to your node or system). As a part of the download bundle, the Fix Central server generates `md5sum` output for each file that this command validates the download against. If the `md5sum` output is not generated in any file, all files downloaded at that point are deleted and the download is aborted.

Various software packages might be available depending on your product. This includes new builds for upgrading or downgrading your system, upgrade checking software, remote support proxy server, ifixes, and other software packages.

Specify **lscmdstatus** to display detailed results. For example, specify `sainfo lscmdstatus`:

If you are using Cloud Call Home to download packages, the following results display:

```
last_command satask downloadsoftware -user ##### -vrmf x.x.x.x
last_command_status CMMVC8044E Command completed successfully.
T3_status
T3_status_data
cpfiles_status
cpfiles_status_data
snap_status
snap_filename
installcanistersoftware_status
supportupload_status
supportupload_status_data
supportupload_progress_percent 0
supportupload_throughput_KBps 0
supportupload_filename
downloadsoftware_status Active
downloadsoftware_status_data Downloading the bundle
downloadsoftware_progress_percent 38
downloadsoftware_throughput_KBps 321
downloadsoftware_size 467.6 MiB
```

If you are using SFTP through Fix Central to download packages, the following results display:

```
last_command satask downloadsoftware -user ##### -password #####
last_command_status CMMVC8044E Command completed successfully.
T3_status
T3_status_data
cpfiles_status
cpfiles_status_data
snap_status
snap_filename
installcanistersoftware_status
supportupload_status
supportupload_status_data
supportupload_progress_percent 0
supportupload_throughput_KBps 0
supportupload_filename
downloadsoftware_status Active
downloadsoftware_status_data Downloading the bundle
downloadsoftware_progress_percent 38
downloadsoftware_throughput_KBps 321
downloadsoftware_size 467.6 MiB
```

Invocation examples

The following example downloads software packages from Fix Central using Cloud Call Home with a **-fixid** and **-accesskey** specified.

```
satask downloadsoftware -user bob.example@ibm.com -accesskey 131315948042609040 -fixid
StorageSoftware-2145-IBM_INSTALL_patch_assert_1.00
```

The resulting output:

```
No feedback (use sainfo lscmdstatus to display software download information)
```

You can use the following CLI commands to obtain the package:

- To download the software patch files through the Call Home by using the RESTful API, enter the following command:

  ```
  satask downloadsoftware -patchid Storage-FlashSystem9100-PATCH-12345678
  ```

  where **-patchid** specifies the name of the patch file.
- To download the upgrade test utility by using the command-line interface, enter the following command:

  ```
  satask downloadsoftware -upgradetestutility
  ```

  where **-upgradetestutility** downloads the upgrade test utility.

The following example downloads software packages from Fix Central using Cloud Call Home with a **-vrmf** specified.

```
satask downloadsoftware -user bob.example@ibm.com -vrmf x.x.x.x
```

The resulting output:

```
No feedback (use sainfo lscmdstatus to display software download information)
```

The following example shows downloading a software package from Fix Central using SFTP with temporary username and password.

```
satask downloadsoftware -user rOLrhyPf -password 'E4yrr6WZM' -templogin
```

The resulting output:

```
No feedback (use sainfo lscmdstatus to display software download information)
```

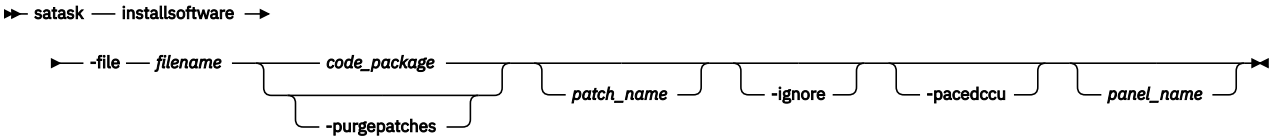
dumpinternallog (Discontinued)

The **dumpinternallog** command has been discontinued.

installsoftware

Use the command to install a specific system code package and a specific patch on a single node.

Syntax



Parameters

-file *filename*
(Required) The file name of code installation package.

Note: The argument to **-file** must be present on the local node; the argument is automatically copied to the target *panel_name*.

patch_name
Specifies the name of the patch file.

code_package

Specifies the name of the code package uploaded on the configuration node.

-purgepatches
When a code package is installed on a system, all the existing patches are reapplied. When you specify the **-purgepatches** parameter, all the existing patches on the node are removed.

-ignore
(Optional) Overrides prerequisite checking and forces installation of the code.

-pacedccu
(Optional) Causes the node to initiate a paced Concurrent Code Update (in which you define when the node begins its update) instead of a normal Concurrent Code Update (in which each node in the clustered system automatically updates in sequence).

panel_name
(Optional) Identifies the node that is being serviced.

Note: If `panel_name` is not supplied, this applies to the node on which the command is running.

Description

This command installs a specific code package and a specific patch on a single node. You can also use this command to install a code package without installing a patch.

Important: Use this command only under the direction of your support team.

An invocation example

```
satask installsoftware -file install_pkg.gpg nodeB_panel_name
```

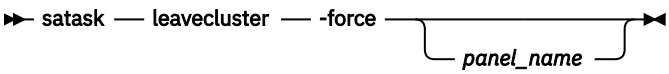
The resulting output:

```
No feedback.
```

leavecluster

Use the **leavecluster** command to remove system state data, location information, and other history from a node.

Syntax



Parameters

-force
(Required) The **-force** parameter is required because this service action can cause temporary or permanent loss of access to data.

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

panel_name
(Optional) Identifies the node that is being serviced.

Note: If `panel_name` is not supplied, this command applies to the node on which the command is running.

Description
Use this command to remove system state data, location information, and other history from a node.

An invocation example

```
satask leavecluster -force 78G00F3-2 /* this forces the node with panel_name=78G00F3-2 out of the clustered system */
```

The resulting output:

```
No feedback
```

Note: Using this command causes both nodes to become unmanaged, and the system warm starts them.

locksuperuser

Use the **locksuperuser** command to lock the superuser so that it cannot be used to log in to the system or perform service actions. The user of this task is limited to the IBM Remote Support user.

Syntax



Description
This command locks the superuser account.

Note:

- For nodes in candidate mode, this command is not allowed and fails with an error message.
- For nodes in spare mode, this command is not allowed and fails with an error message.
- For nodes in service mode, this command is not allowed and fails with an error message.
- For nodes in active mode, this command can run on any node in the cluster and pushes the setting to all active and spare nodes.
- Only the **sra_privileged** users can run this command in order to configure the remote support.

An invocation example

```
satask locksuperuser
```

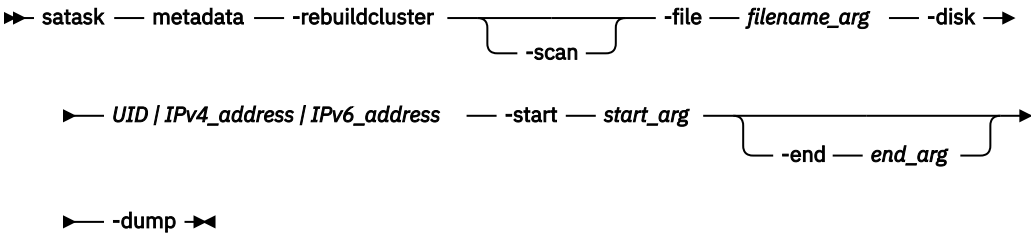
The resulting output

```
No feedback
```

metadata

Use the **metadata** command to recover the critical metadata describing arrays and volumes.

Syntax



Parameters

satask
(Required) Specifies a system administrator task, such as issuing service commands that are only used in specific circumstances.

-rebuildcluster
(Required) Creates a cluster from the metadata found in `/dumps/t3_recovery.bin` created by the -dump process.

-scan
(Optional) Scans the specified MDisk or drive for system metadata.

-file ***filename_arg***
(Required) Specifies the file in which you want the results of a scan operation. The file is placed into the node in the directory `/dumps`, and can be retrieved using Secure Copy (**scp**). The file can be subsequently cleaned using the **cleardumps** command.

-disk ***UID | IPv4_address | IPv6_address***
(Required) Specifies the UID, IPv4 address, or IPv6 address of the MDisk or drive that you want to scan or remove a dump from. If an IPv4 or IPv6 address is specified, the **-start** and **-end** parameters are ignored.

-start ***start_arg***
(Required) Specifies the following conditions:

- When used with **-scan**: The first LBA in which to look for metadata on the disk.
- When used with **-dump**: The first LBA at which the metadata resides (as reported in the scan file).

If the *start_arg* variable is not specified, *0* is used as the default.

-end ***end_arg***
(Optional) Specifies the last LBA in which to look for metadata on the disk.

-dump
(Required) Dumps metadata from the specified MDisk or drive to file /dumps/t3_recovery.bin

Description

Use this command to recover the critical metadata describing arrays and volumes.

Use the **lscmdstatus** command to see the status of the command.

An invocation example

```
satask metadata -scan -file scan.0.xml -disk 044234043539t5305304534534953454303 45 -start 0
```

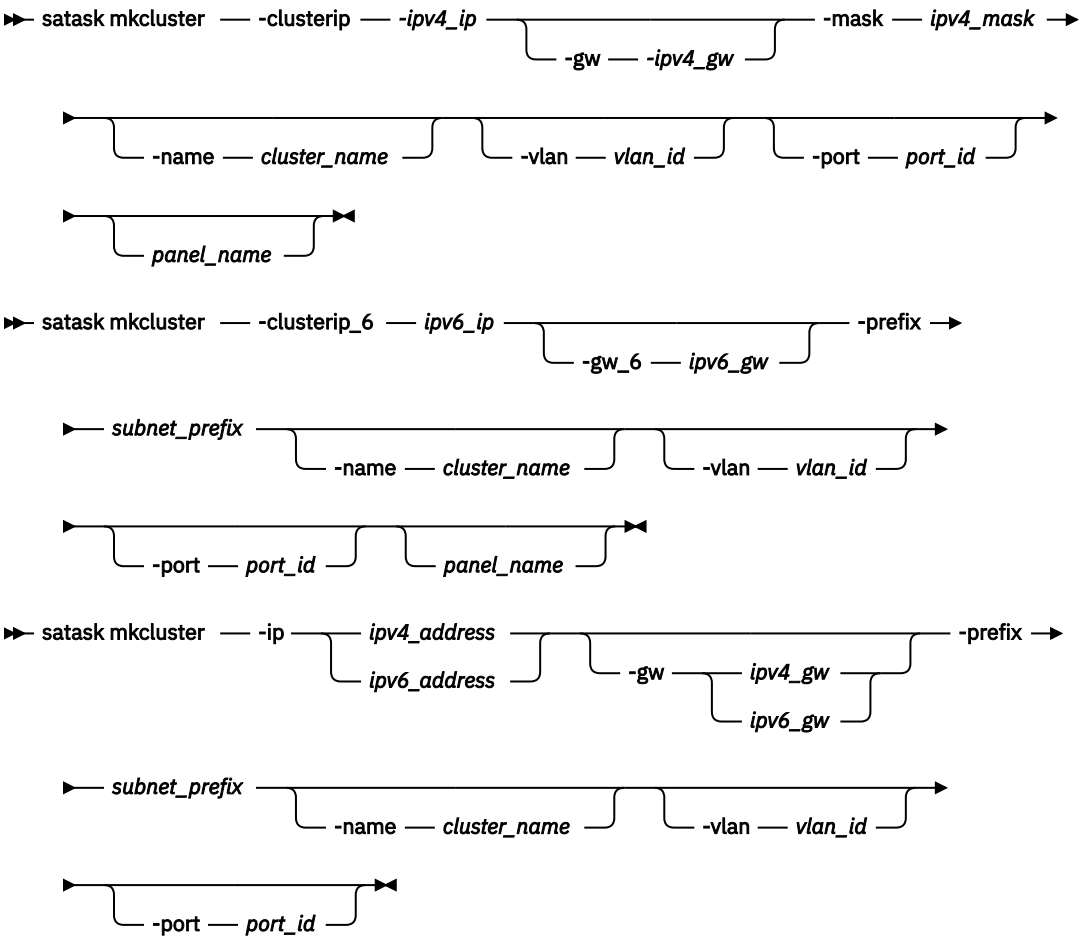
The resulting output:

```
No feedback
```

mkcluster

Use the **mkcluster** command to create a new system.

Syntax



Parameters

-clusterip *ipv4_ip*
(Optional) Specifies the Internet Protocol Version 4 (IPv4) address for the specific port ID.

-mask *ipv4_mask*
(Optional) Specifies the IPv4 subnet for the specific port ID.

-clusterip_6 *ipv6_ip*
(Optional) Specifies the Internet Protocol Version 6 (IPv6) address for the specific port ID.

-gw_6 *ipv6_gw*
(Optional) Specifies the IPv6 gateway for the specific port ID.

-prefix *ipv6_subnet*
(Optional) Specifies the IPv6 prefix for the specific port ID.

-name *cluster_name*
(Optional) Specifies the name of the new system.

panel_name
(Optional) Specifies the node being serviced.

Note: If `panel_name` is not supplied, this command applies to the node on which the command is running.

-ip ***ipv4_address|ipv6_address***
(Optional) Specifies the VLAN ID for a IPv4 or IPv6 address.

-prefix ***subnet_prefix***
(Required) Specifies the prefix of the subnet. The value for IPv4 addresses and IPv6 addresses can range from 1 to 31 and 1 to 127 respectively.

-gw ***gateway***
(Optional) Specifies the IPv4 or IPv6 gateway address. The value consists in the standard format for IPv4 address, or IPv6 address of a string of 40 characters.

-vlan ***vlan_id***
(Optional) Specifies the VLAN ID for a IPv4 or IPv6 address.

-port ***port_id***
(Optional) Specifies a new IP address to the Ethernet Port ID.

Note: If **-port** parameter is not specified the default port ID is 1.

Remember: You must specify one of the following:

- IPv4 system IP, gateway, and subnet
- IPv6 system IP, gateway, and prefix

Description
This command creates a new system.

An invocation example

The following example uses specific **-clusterip**, **-gw**, and **-mask** parameters.

```
satask mkcluster -clusterip x.xxx.xxx.xxx -gw x.xxx.xxx.xxx -mask x.xxx.xxx.xxx
```

The resulting output:

```
No feedback
```

mkcluster (Deprecated)

The **mkcluster** system command is deprecated. Use the **satask mkcluster** command to create a new system.

overridequorum

Use the **overridequorum** command to start the manual override command.

Syntax



Parameters

-force
(Required) Overrides any quorum decision that is made by the system.

Important: This command must only be used with assistance from IBM Support. Using this command incorrectly can cause data corruption, by creating two separate clusters which both write to the same disks. To avoid this, the command must only be used in a disaster recovery scenario, when it has been ensured that nodes at the other site can never come back online.

Description

This command starts the manual override command. This command is valid on nodes that are in a starting state with either of the following node errors:

- 551
- 921

An invocation example

```
satask overridequorum -force
```

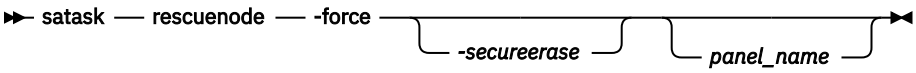
The resulting possible outputs:

```
No feedback.
```

satask rescuenode

Use the **satask rescuenode** command to start automatic recovery for a specific node. Use this command only when the node reports corrupted code. This command is not supported for IBM Storage Virtualize for Public Cloud on Amazon Web Services (AWS) and Azure.

Syntax



Parameters

panel_name
(Optional) Identifies the node that is being serviced.

Note: If panel_name is not supplied, it applies to the node on which the command is running.

-secureerase
(Optional) This parameter securely erases all boot devices and returns the system to the manufactured state.

Note: This parameter is not supported on cloud systems.

-force
(Required) The **-force** parameter is required because this service action can cause temporary or permanent loss of access to data. Use this command only when the node reports corrupted system code.

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

Description

This command starts automatic recovery for a specific node. Use this command only when the node reports corrupted code.

An invocation example

```
satask rescuenode -force 112233
```

The resulting output:

```
No feedback
```

resetpassword

Use the **resetpassword** command to reset the system superuser password to a user-defined value.

Syntax



Parameters

satask
System administrator task; service commands that are only used in specific circumstances.

-password ***'cleartext_password'***
(Required) Sets the system superuser password to **'cleartext_password'**. The password must be in the range of 6 - 64 characters. You must enclose the password in single quotation marks.

Description

This command resets the system superuser password to a user-defined value. This command also unlocks the superuser account if it is locked, removes the SSH key for the superuser and disables the *password and key required* feature for the superuser. The superuser password cannot be reset back to the default password.

An invocation example

```
satask resetpassword -password 'new_password'
```

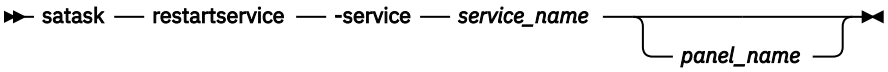
The resulting output

```
No feedback
```

satask restartservice

Use the **satask restartservice** command to restart a named service.

Syntax



Parameters

satask
Specifies a system administrator's task (such as service commands that are only used in specific circumstances).

-service ***service_name***
(Required) Specifies the name of the service that you want to restart. The following services are supported:

sshd
Specifies a Secure Shell (SSH) Daemon.

slpd
Specifies a Service Location Protocol Daemon.

slv_dpadmpd
Specifies Easy Tier service.

tomcat
Specifies a web server.

cimserver
Specifies cimserver.

reclaim_se
Specifies the reclamation daemon for thin-provisioned volumes.

panel_name
(Optional) Identifies the node that is being serviced.

Note: If `panel_name` is not supplied, this command applies to the node on which the command is running.

Description

Important: When directed to do so by your product support team, use this command to restart a named service.

An invocation example

```
satask restartservice -service cimserver
```

The following output is displayed:

```
No feedback
```

rmnodediscoverysubnet

Use the **rmnodediscoverysubnet** command to remove a configured subnet and stop it from being used for node discovery.

Syntax



Parameters

-id ***subnet_id***
(Required) Specify the subnet ID that you want to remove.

Description

This command removes a node discovery subnet from the subnets table.

An invocation example

```
satask rmnodediscoverysubnet -id 1
```

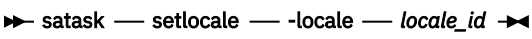
The resulting output:

```
No feedback
```

setlocale

Use the **setlocale** command to change the **satask** and **sainfo** command output to the chosen language on the current node canister.

Syntax



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).

Issue the **satask setlocale** command to change the locale setting for the system; all interface output is changed to the chosen language. For example, to change the language to Japanese, type the following:

```
satask setlocale -locale 3
```

where 3 is the value for Japanese. The following values are supported:

- 0 US English (default)
- 1 Simplified Chinese
- 2 Traditional Chinese
- 3 Japanese
- 4 French
- 5 German
- 6 Italian
- 7 Spanish
- 8 Korean
- 9 Portuguese (Brazilian)
- 10 Russian

An invocation example (where 3 is Japanese)

```
satask setlocale -locale 3
```

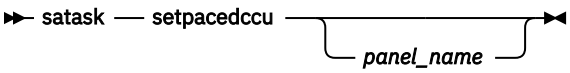
The resulting output:

```
No feedback
```

setpacedccu

Use the **setpacedccu** command to flag a node to participate in a user-paced system update.

Syntax



Parameters

panel_name
(Optional) Identifies the node being serviced.

Note: If `panel_name` is not supplied, this applies to the node on which the command is running.

Description

Use this command to flag a node to participate in a user-paced system update.

This command can only be used when the node is:

- In a service state
- Error-free
- Not part of a clustered system when the node is out of a service state

An invocation example

```
satask setpacedccu
```

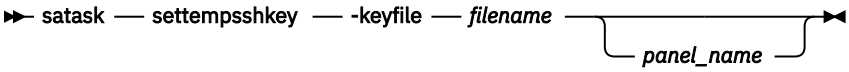
The resulting output:

```
No feedback
```

settempsshkey

Use the **settempsshkey** command to install a temporary Secure Shell (SSH) key for a superuser ID to run commands in the service assistant CLI.

Syntax



Parameters

-keyfile ***filename***
(Required) Specifies the name of the file that contains the Secure Shell (SSH) public key. The file identified by *filename* must be on the local node (or on the USB flash drive, if you execute the command from there).

panel_name
(Optional) Identifies the node being serviced.

Note: If `panel_name` is not supplied, this applies to the node on which the command is running.

Description

This command installs a temporary SSH key for a superuser ID to run commands in the service assistant CLI (for example, to copy files to or from the node).

You can only use this command when completing a service action. Installing a temporary key will replace any available existing keys. The key will be deleted when the node joins a cluster or is rebooted or power cycled.

An invocation example

```
satask settempsshkey -keyfile jvardy12
```

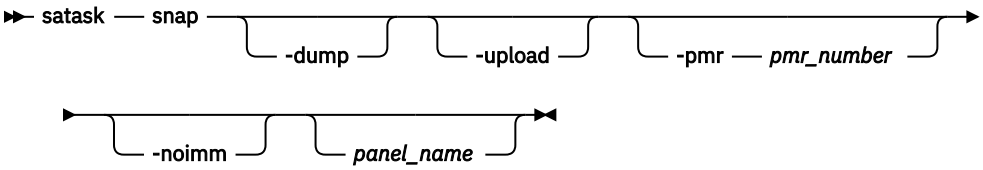
The resulting output:

```
No feedback
```

snap

Use the **satask snap** command to collect diagnostic information from the node canister.

Syntax



Parameters

-dump
(Optional) Specifies that the most recent dump file on the node should be included in the snap file.

-upload
(Optional) Specifies that the snap file be uploaded after it is generated.

-pmr *pmr_number*
(Optional) Specifies the PMR number to use to upload the snap file. The format for a PMR is a 13-character alphanumeric string or letters TS followed by 9 digits. If the specified PMR is invalid or unknown, it is uploaded to a generic location on the server with the prefix:

```
unknown_pmr_pmr_number_
```

If this option is not supplied, the snap file is uploaded using the machine type and serial number attributes.

-noimm
(Optional) Specifies that diagnostic data should not be collected from the Integrated Management Module (IMM).

panel_name
(Optional) Indicates the node on which to execute the **snap** command.

Description

This command completes immediately and the snap process runs as a background task. The user can use the **sainfo lscmdstatus** view to monitor the status of the task.

This command creates a new snap file to collect diagnostic information from the node. If this command is run from the command line, the snap file will be created in the /dumps directory on the specified node. Alternatively, if this command is run from a USB drive, the snap file will be created on the USB drive.

You can use the **sainfo lsfiles** command to list the files on the node.

Note: You cannot specify a *panel_name* that belongs to an SAS enclosure.

An invocation example

```
satask snap
```

The resulting output:

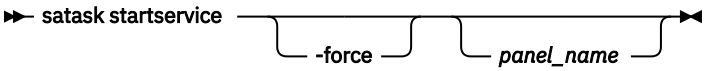
```
No feedback
```

Important: The name of the output file (placed on the specified node) is `snap.single.`*`nodeid`*`.`*`date`*`.`*`time`*`.tgz`.

startservice

Use the **startservice** command to enter a service state.

Syntax



Parameters

satask
(Required) Specifies a system administrator task, such as issuing service commands that are only used in specific circumstances.

-force
(Optional) Overrides checking of system membership.

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

panel_name
(Optional) Identifies the node being serviced.

Note: If `panel_name` is not supplied, this applies to the node on which the command is running.

Description

This command causes a node to go into service state. For example, a user might put a system node into service state to remove it from candidate state or to prevent it from automatically being added to a system again.

The **-force** flag is required if the action could interrupt I/O (last node in cluster or IO group). This command holds the node in service state until it is cleared using the **satask stopservice** command, or until the I/O process is restarted.

An invocation example

```
satask startservice
```

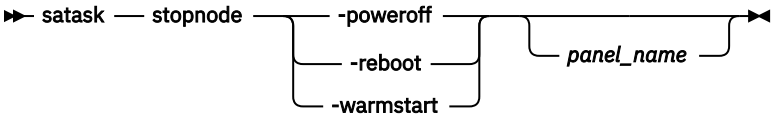
The resulting output:

```
No feedback
```

stopnode

Use the **stopnode** command to power off, reboot, or warmstart a node. Use the **stopnode** command to reboot or warmstart a node.

Syntax



Parameters

-poweroff
(Required if **-reboot** and **-warmstart** are not specified) Powers off the node.

-reboot
(Required if **-poweroff** and **-warmstart** are not specified)Required if **-warmstart** is not specified Reboots the node.

-warmstart
(Required if **-poweroff** and **-reboot** are not specified) Restarts the I/O process and issues a diagnostic dump (also known as a full dump).

Important: You can also specify `stopsystem -node -reset` to restart the I/O process (but in a more controlled manner).

(Required if **-reboot** is not specified) Restarts the I/O process and issues a diagnostic dump.

panel_name
(Optional) Identifies the node being serviced.

Note: If `panel_name` is not supplied, this applies to the node on which the command is running.

Description

Use the **stopnode** command to power off, reboot, or warmstart a node. Use the **stopnode** command to reboot or warmstart a node.

An example for powering off canister 1 in enclosure 2

```
satask stopnode -poweroff 02-1
```

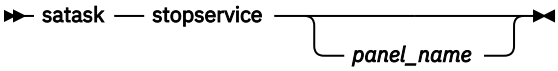
The resulting output:

```
No feedback
```

stopservice

Use the **stopservice** command to exit a service state.

Syntax



Parameters

panel_name
Identifies the node that is being serviced.

Note: If `panel_name` is not supplied, this command applies to the node on which the command is running.

Description

This command exits the service state that was entered with the **`startservice`** command, then exits the service state on the local node.

An invocation example

```
satask stopservice
```

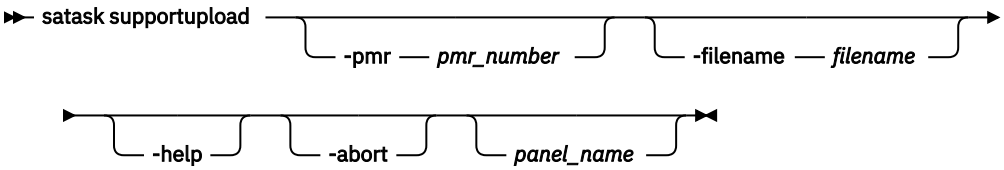
The resulting output

```
No feedback
```

supportupload

Use the **`supportupload`** command to upload files to a node or system.

Syntax



Parameters

`-pmr` ***pmr_number***
(Optional) Specifies the PMR (software service request) number that is related to a service report or request. If you do not specify this parameter, the system uses the machine type (MTM) and serial number during the upload. The format for a PMR is a 13-character alphanumeric string or letters TS followed by 9 digits. The PMR value is in the format `ppppp,bbb,ccc`, where:

- p indicates PMR.
- b indicates branch.
- c indicates country.

`-filename` ***filename***
(Optional) Specifies the file name to be uploaded. If you do not specify this parameter, the system uses the most recent snap file that is available in the `/dumps` directory to upload. The value that is specified must be a 256-character alphanumeric string.

`-help`
(Optional) Specifies that usage information is displayed.

`-abort`
(Optional) Specifies upload cancellation.

panel_name
(Optional) Specifies the node that is being serviced. The value must be an alphanumeric string.

Note: If not specified, this variable applies to the node on which the command is running.

Description

This command starts or stops node or system file uploads.

Note: If you do not specify a PMR number, files are uploaded to the server (MTSM specific location) under machine type and serial number attributes. Any user can specify this CLI command.

If you specify an incorrect, invalid, or unknown PMR, the files are uploaded to the system with the prefix `unknown_pmr_pmrnumber_filename`.

Remember: Before you specify this command:

1. Make sure that you have internet access for all nodes.
2. Configure service IP addresses on all nodes in your system.
3. Request your network administration to let the firewall allow connections to the following Internet Protocol (IP) addresses on port 443:
   - `129.42.56.189`
   - `129.42.54.189`
   - `129.42.60.189`
4. Configure a DNS server by specifying **`mkdnsserver`** to define a Domain Name System (DNS) server and **`lsdnsserver`** to display its values.

If you specify this command on a remote node, the upload completes by using the service IP address and port of that remote node.

Note: Only one file upload at one time is supported per node.

Do not specify the **`supportupload`** and **`snap`** commands simultaneously to upload files, or a message that indicates an upload is in progress is displayed. One upload can be completed at one time per node by using `supportupload/snap/svc_snap`.

Remember: Specify **`lscmdstatus`** to display detailed results. For example, specify **`sainfo lscmdstatus`**:

```
last_command satask supportupload -pmr 79556,019,866
last_command_status CMMVC8044E Command completed successfully.
T3_status
T3_status_data
cpfiles_status
cpfiles_status_data
snap_status
snap_filename
installcanistersoftware_status
supportupload_status Active
supportupload_status_data Uploading
supportupload_progress_percent 56
supportupload_throughput_KBps 99939
supportupload_filename /dumps/snap.single.7830619-1.161219.161359.tgz
downloadsoftware_status
downloadsoftware_status_data
downloadsoftware_progress_percent 0
downloadsoftware_throughput_KBps 0
downloadsoftware_size
```

An invocation example

```
satask supportupload -abort
```

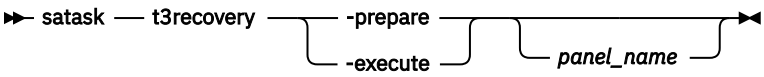
The resulting output:

```
No feedback (use lscmdstatus to display support upload information)
```

t3recovery

Use the **t3recovery** command to prepare and start a T3 recovery.

Syntax



Parameters

panel_name
Identifies the node that is being serviced.

Note: If `panel_name` is not supplied, this recovery applies to the node on which the command is running.

-prepare
Search for T3 recovery data. This parameter locates the necessary backup file by date and the quorum disk.

-execute
Start T3 recovery with recovered data.

Description

This command prepares and starts a T3 recovery.

Important: Progress of a T3 recovery can be displayed with the **sainfo lscmdstatus** command.

An invocation example

```
satask t3recovery -prepare
```

The resulting output

```
No feedback
```

unlocksuperuser

Use the **unlocksuperuser** command to unlock the superuser so that it can be used to log in to the system or perform service actions.

Syntax



Description

This command unlocks the superuser account.

Note:

- For nodes in candidate mode, this command has no effect. The superuser is always unlocked in candidate state.
- For nodes in spare mode, this command is not allowed and fails with an error message.
- For nodes in service mode, this command unlocks the superuser on the local node only. Rejoining the cluster discards any local setting.

- For nodes in active mode, this command can run on any node in the cluster and pushes the setting to all active and spare nodes.

An invocation example

```
satask unlocksuperuser
```

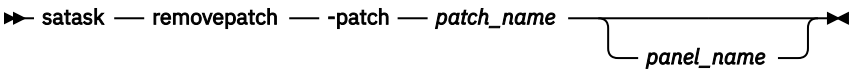
The resulting output

```
No feedback
```

satask removepatch

Use the **satask removepatch** command to remove a specified patch from the node.

Syntax



Parameters

-patch *patch_name*
(Required) Specifies the name of the patch file.

panel_name
(Optional) Specifies the panel name of the node to remove the patches from.

An Invocation example

```
satask removepatch -patch patch_name panel_name
```

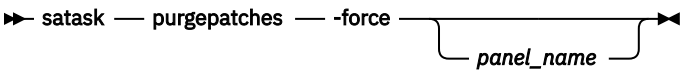
The resulting output:

```
No feedback.
```

satask purgepatches

Use the **satask purgepatches** command to remove all patches from a node.

Syntax



Parameters

-force
(Required) Specifies that the service action removes all installed patches from a node.

panel_name
(Optional) Specifies the panel name of the node to remove the patches from.

An invocation example

```
satask purgepatches -force panel_name
```

The resulting output:

```
No feedback.
```

Chapter 36. Service node task commands

Service node task (**sntask**) commands are used during the initial installation of IBM Storage Virtualize for Public Cloud software onto an IBM Cloud bare metal server or Amazon Web Services (AWS) container and afterward to manage the IBM Storage Virtualize for Public Cloud node on the bare metal server or AWS container. For example, you might start the node with **sntask startnode** and stop the node with **sntask stopnode**. On IBM Cloud, the **sntask** commands and service node information (**sninfo**) commands reside and run as terminal commands on the Red Hat Enterprise Linux 7 operating system of the bare metal server. On AWS, the **sntask** commands and service node information (**sninfo**) commands reside and run as terminal commands on the CentOS operating system of the EC2 instances.

The **sntask** and **sninfo** commands are installed as part of the package that is installed during the automatic or manual installation. The installation also creates the user (`sv_cloud`) for running the commands. The task commands gather information that is required to load the software and configure and manage the IBM Storage Virtualize for Public Cloud node. The **sninfo** commands output information about the nonce and the status of the installed IBM Storage Virtualize for Public Cloud node.

To use the **sntask** and **sninfo** commands, **ssh** to the bare metal server as the sv_cloud user. Log in with an SSH key, if possible, or with the sv_cloud password if you did not yet create and install an SSH key pair. Set the initial password for the sv_cloud user from the root user if the password is not yet set.

```
# passwd sv_cloud
```

To create the SSH key pair, refer to the RHEL man pages for a description of how to use the **ssh-keygen -t rsa** terminal command. To create the key pair from a Microsoft Windows system, use the `puttygen.exe` utility, as described in the PuTTY documentation.

To bypass the default session time limit for the SSH session, use **ssh** with the **TCPKeepAlive** option:

```
$ ssh server_name -l sv_cloud -o TCPKeepAlive=yes
```

You are prompted for the password if you are not using an SSH key. To change the sv_cloud password when a password exists, use the sv_cloud user to run the following terminal command:

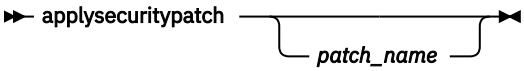
```
$ passwd sv_cloud
```

You can now run **sntask** and **sninfo** commands.

sntask applysecuritypatch

Use the **sntask applysecuritypatch** command to apply available security patches to IBM Storage Virtualize for Public Cloud node that is deployed in a supported cloud services provider.

Syntax



Parameters

patch_name
(Optional) Specifies the name of the patch. You can specify one or several patches, which are separated by a blank space. If you do not specify a patch name, the command tries to apply all available patches.

© Copyright IBM Corp. 2003, 2025

Description

This command applies available security patches on the host system.

Note: The IBM Storage Virtualize for Public Cloud node must be placed in service mode by using the **startservice** command before you can apply security patches.

An invocation example

```
sntask applysecuritypatch firefox
```

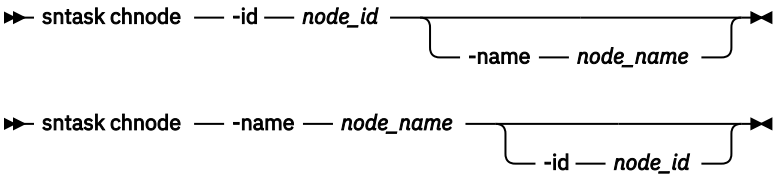
The following output is displayed:

```
Info: firefox is not installed on your system, it does not need to apply patch for it.
Security patches have been applied!
```

sntask chnode

Use the **sntask chnode** command to change IBM Storage Virtualize for Public Cloud node information.

Syntax



Parameters

-id *node_id*
(Optional) Specifies the node ID to change. The value must be a number.

-name *node_name*
(Optional) Specifies the node name to change. The value must be an alphanumeric string.

Description

This command changes IBM Storage Virtualize for Public Cloud node information. You can use this command to set specific values after you complete the initialization of the IBM Storage Virtualize for Public Cloud node. All parameters that are associated with this command are optional. However, you must specify one or more parameters.

An invocation example

```
chnode -name baremetal1
```

The following output is displayed:

```
No feedback
```

sntask cleansnap

Use the **sntask cleansnap** command to delete log files from the system. Both the **satask snap** command and the **sntask snap** command create the log files.

Syntax



Parameters

No parameters

Description

This command deletes most of the log files in the `/var/log/SpectrumVirtualize/` directory on the host.

Running the **satask snap** command automatically calls the **sntask snap** command to create the log files and a .tar file of the logs on the host computer. You can also run the **sntask snap** command directly to create the log files on the host computer.

The **sntask cleansnap** command erases the .tar file and all log files on the host computer except for an operating system log for debugging purposes.

Running the **satask snap** command creates log files, memory dumps, and the log files from the host computer in a .tar file in the `/dumps/` directory within the system. To clean up the `/dumps/` directory if space is unavailable or for any other reason, use the **cleardumps** command.

An invocation example of sntask cleansnap

```
sntask cleansnap
```

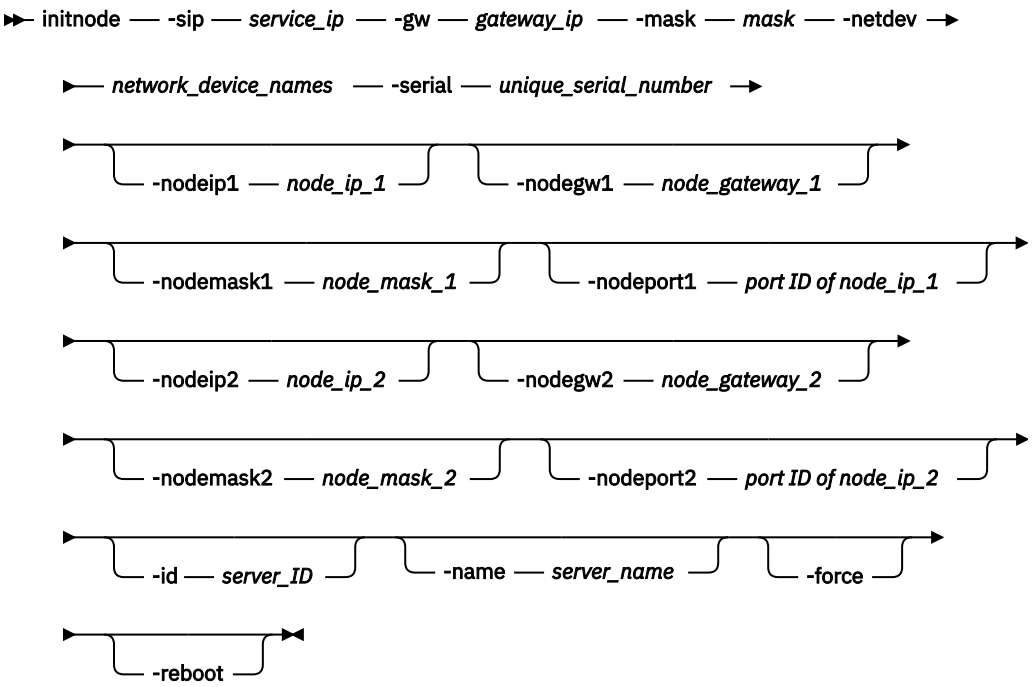
The following output is displayed:

```
snap tarballs under /var/log/SpectrumVirtualize have been cleaned up.
```

sntask initnode

Use the **initnode** command to initialize the IBM Storage Virtualize for Public Cloud software in a supported cloud environment.

Syntax



Parameters

-sip
(Required) Specifies the service IP address of the node. The format of this IP address is IPv4.

-gw
(Required) Specifies the gateway IP address of the node. The format of this IP address is IPv4.

-mask
(Required) Specifies the subnet mask of the service IP address. The format of this IP address is IPv4.

-netdev
(Required) Specifies the network interface names on the host. Two network interface names that are separated by a colon are required (example: `eth0:eth1`).

Note: This parameter is only supported for AWS and Microsoft Azure.

-serial
(Required) Specifies a unique serial number for the node.

Note: This parameter is only supported for IBM Cloud.

-nodeip1
(Optional) Specifies the node IP address for port 1.

-nodegw1
(Optional) Specifies the gateway address for port 1.

-nodemask1
(Optional) Specifies the subnet mask for port 1.

-nodeport1
(Optional) Specifies the port ID for port 1.

-nodeip2
(Optional) Specifies the node IP address for port 2.

-nodegw2
(Optional) Specifies the gateway address for port 2.

-nodemask2
(Optional) Specifies the subnet mask for port 2.

-nodeport2
(Optional) Specifies the port ID for port 2.

-id ***ID***
(Optional) Specifies the ID from the supported cloud service.

Note: This parameter is only supported for IBM Cloud.

-name ***server_name***
(Optional) Specifies the server name from supported cloud service.

Note: This parameter is only supported for IBM Cloud.

-force or -f
(Optional) Specifies to install the IBM Storage Virtualize for Public Cloud software without the partition template.

Note: This parameter is only supported on previous versions of IBM Storage Virtualize for Public Cloud software.

Note: This parameter is only supported for IBM Cloud.

-reboot
(Optional) Specifies to restart the server if necessary after initialization.

Note: This parameter is only supported for IBM Cloud.

Description

This command initializes the IBM Storage Virtualize for Public Cloud software.

An invocation example for IBM Cloud

```
sntask initnode -sip service_ip -gw gateway_ip -mask mask -serial serial_number -force
```

The following output is displayed:

```
Storage-virtualize node will be installed on /dev/sdb
  Downloading, it may take a few minutes.
  Installing, it may take a few minutes
  Storage-virtualize node is successfully installed
```

An invocation example for AWS

```
sntask initnode -sip 172.31.47.37 -gw 172.31.32.1 -mask 255.255.240.0 -netdev eth0:eth1
                -nodeip1 172.31.47.128 -nodegw1 172.31.32.1 -nodemask1 255.255.240.0 -nodeport1
1
                -nodeip2 172.31.34.222 -nodegw2 172.31.32.1 -nodemask2 255.255.240.0 -nodeport2
2
```

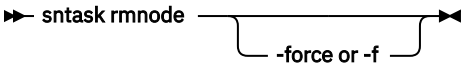
The following output is displayed:

```
CMMSV1036I Download image succeeded.
CMMSV1023I The node has successfully initialized. Reboot the node.
```

sntask rmnode

Use the **sntask rmnode** command to erase the virtual machine image when you uninstall the IBM Storage Virtualize for Public Cloud software from a supported cloud services provider.

Syntax



Parameters

-force or -f
(Optional) Specifies to force the erasure of the virtual machine image when you uninstall the IBM Storage Virtualize for Public Cloud software from a supported cloud services provider.

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

Description

This command erases the image from the cloud provider when you uninstall the IBM Storage Virtualize for Public Cloud software. This step is the first step of uninstalling the product. You must remove other files and objects that are used by the software from your instance in the cloud services provider.

An invocation example

```
sntask rmnode
```

The following output is displayed:

```
Remove the node on /dev/sdb: [y/n]  y
 Storage-virtualize node is removed.
```

sntask snap

Use the **sntask snap** command to collect the debugging logs. The **satask snap** command automatically calls the **sntask snap** command to also include the host machine logs with its system logs.

Syntax

Parameters

No parameters

Description

This command collects the host machine debugging logs. The log files are placed in `/var/log/SpectrumVirtualize/`. A tgz file is copied to the node under the `/dumps/` directory.

Although the **sntask snap** command is automatically called when you run the **satask snap** command, sometimes you might need to create the host machine logs for host debugging.

An invocation example

```
sntask snap
```

The following output is displayed:

```
Log saved to /var/log/SpectrumVirtualize/debug.host03.20170818.005917.tgz
```

sntask startnode

Use the **sntask startnode** command to power on an IBM Storage Virtualize for Public Cloud node.

Syntax



Parameters

No parameters

Description

This command powers on a IBM Storage Virtualize for Public Cloud node in a supported cloud provider.

An invocation example

```
sntask startnode
```

The following output is displayed:

```
The node started successfully.
```

sntask startservice

Use the **sntask startservice** command to enter a service mode state for IBM Storage Virtualize for Public Cloud.

Syntax



Parameters

No parameters.

Description

This command causes a node to go into a service mode state. For example, an administrator can put a system node into a service state to apply security patches.

An invocation example

```
sntask startservice
```

The resulting output:

```
No feedback
```

sntask stopnode

Use the **sntask stopnode** command to power off an IBM Storage Virtualize for Public Cloud node.

Syntax



Parameters

-poweroff
(Optional) If specified and the IBM Storage Virtualize for Public Cloud node is not in service mode, this flag powers off the host after the container is stopped.

-reboot
(Optional) If specified, the IBM Storage Virtualize for Public Cloud node restarts if it is in service mode.

Description

This command powers off a IBM Storage Virtualize for Public Cloud node that is deployed on a supported cloud services provider.

An invocation example

```
sntask stopnode
```

The following output is displayed:

```
No feedback
```

sntask stopservice

Use the **sntask stopservice** command to exit a service mode state for IBM Storage Virtualize for Public Cloud node.

Syntax



Parameters

No parameters.

Description

This command causes a node to exit service mode state.

An invocation example

```
sntask stopservice
```

The resulting output:

```
No feedback
```

Chapter 37. Single sign-on commands

Use the following set of commands to configure the system-wide single sign-on authentication settings.

lsauthsinglesignon

Use the **lsauthsinglesignon** command to list the system-wide single sign-on authentication configuration.

Syntax



Parameters

None.

Description

Use this command to display a detailed view to list the system-wide single sign-on authentication configuration.

This table provides the attribute values that can be displayed as output view data.

*Table 138. **lsauthsinglesignon** output*

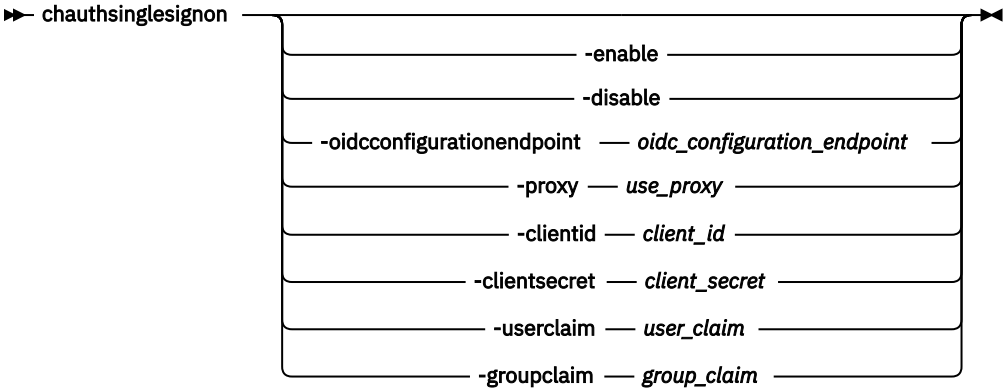
Attribute	Description
status	Indicates the SSO enablement status. Value can be *enabled* or *disabled*.
oidc_configuration_endpoint	Indicates the full path to the openid-configuration endpoint of the SSO authentication server. It must consist of a string up to 255 printable ASCII characters.
use_proxy	Indicates whether the SSO authentication server uses a configured proxy server to connect. Value can be *yes* or *no*.
client_id	Indicates OpenID Connect client ID for the system. It must consist of a string up to 64 printable ASCII characters.
client_secret_set	Indicates whether the OpenID Connect client secret has been set or not. Value can be *yes* or *no*.
user_claim	Indicates the claim in the OpenID Connect ID token used to identify the user's name. It must consist of a string up to 64 printable ASCII characters.
group_claim	Indicates the claim in the OpenID Connect ID token used to identify the user's group. It must consist of a string up to 64 printable ASCII characters.

© Copyright IBM Corp. 2003, 2025

chauthsinglesignon

Use the **chauthsinglesignon** command to change the system-wide single sign-on authentication configuration.

Syntax



Parameters

-enable
(Optional) Specifies the enablement of the single sign-on service.

-disable
(Optional) Specifies the disablement of the single sign-on service. This parameter is mutually exclusive with `-enable`, `-oidcconfigurationendpoint`, `-proxy`, `-clientid`, `-clientsecret`, `-userclaim`, `-groupclaim`.

-oidcconfigurationendpoint ***oidc_configuration_endpoint***
(Optional) Specifies the full path to the openid-configuration endpoint of the single sign-on provider. It must consist of a string up to 255 printable ASCII characters.

-proxy ***use_proxy***
(Optional) Specifies whether to use the configured proxy to connect to the single sign-on provider. Value can be *yes* or *no*.

-clientid ***client_id***
(Optional) Specifies the OpenID Connect client ID for the system, required to authenticate with the single sign-on provider. It must consist of a string up to 64 printable ASCII characters.

-clientsecret ***client_secret***
(Optional) Specifies the OpenID Connect client secret for the system, required to authenticate with the single sign-on provider. It must consist of a string up to 64 printable ASCII characters.

-userclaim ***user_claim***
(Optional) Specifies the claim in the OpenID Connect ID token used to identify the user's name. It must consist of a string up to 64 printable ASCII characters.

-groupclaim ***group_claim***
(Optional) Specifies the claim in the OpenID Connect ID token used to identify the user's group. It must consist of a string up to 64 printable ASCII characters.

An invocation example

```
chauthsinglesignon -oidcconfigurationendpoint https://wildduks.verify.ibm.com/oidc/endpoint/
default/.well-known/openid-configuration
-clientid a3514ba2-aa5b-492b-983d-d34e2ba39f61 -clientsecret mtU81MuMQO -groupclaim marvel
-userclaim spiderman -enable
```

The resulting output:

```
No feedback.
```

Chapter 38. System certificates commands

Use the System Certificates commands to manage the system certificate. The system certificate is used to enable secure connections between systems or applications across networks.

sainfo lsrootcertificate

Use the **sainfo lsrootcertificate** command to view the system's root certificate. The certificate can be copied into a trust store or into a file using a text editor.

Syntax

Note:

The **sainfo lsrootcertificate** command can be used on nodes in service state before the cluster is created or when performing maintenance. The root certificate should be obtained from the node that will be used to initialize the cluster.

Note: The **sainfo lsrootcertificate** command will not run on the spare nodes.

Note: The **sainfo lsrootcertificate** command can be executed by any user.

Parameters

None.

satask exportrootcertificate

Use the **satask exportrootcertificate** command to export the system's root certificate to `/dumps/root_certificate.pem`. The root certificate can be downloaded using scp and installed in web browsers and truststores.

Syntax

Note: The **satask exportrootcertificate** command can be used on nodes before the system is initialized or when the node is part of a system. The root certificate should be obtained from the node that will be used to initialize the system.

Note:

- The **satask exportrootcertificate** command will not run on spare nodes.
- The **satask exportrootcertificate** command can be executed by any user.

Parameters

None.

© Copyright IBM Corp. 2003, 2025

Chapter 39. Snapshot commands

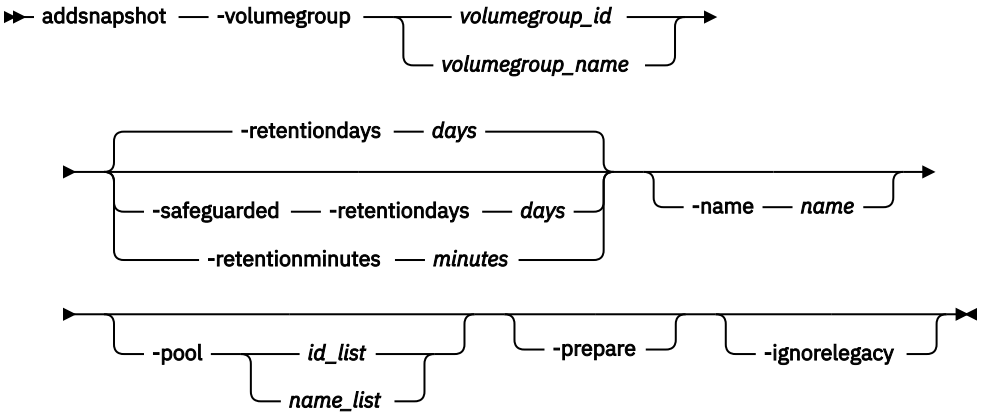
Use the following set of commands to create a management model that allows a storage administrator or automation engine to handle the defined use-cases without having to manage the internal implementation.

addsnapshot

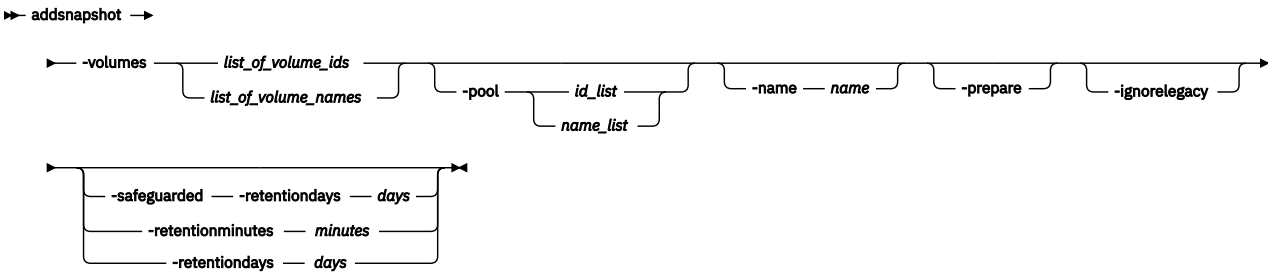
Use the **addsnapshot** command to save a version of a set of volumes.

Syntax

Add snapshot to a volumegroup



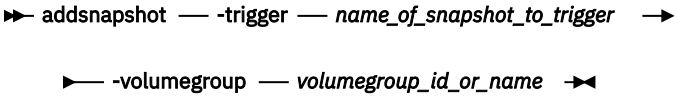
Add snapshot to a list of volumes



Trigger a snapshot using the snapshot name and parentuid



Trigger a snapshot of a volumegroup



Trigger a snapshot using a snapshot ID



Parameters

-pool *id_list | name_list*
(Optional) Specifies the name or ID for the maximum of two storage pools for the new snapshots.

© Copyright IBM Corp. 2003, 2025

-name *name*
(Optional) Specifies the name given to the new snapshot.

-prepare
(Optional) Prepare the snapshot by creating all the objects. The snapshot will not be populated until it is triggered.

-ignorelegacy
(Optional) Specifies the addition of the volume snapshots although there are already legacy FlashCopy mappings using the volume as a source.

-retentionminutes
(Optional) Specifies the number of minutes to wait before removing the snapshot. This parameter cannot be specified together with the **-safeguarded** parameter.

Note: If a snapshot is created using the **-retentionminutes** parameter then the system will automatically remove the snapshot after the specified number of minutes has passed. A snapshot that is being used at that time to populate a clone or restore operation will not be immediately deleted, but will instead move to the `dependent_deleting` state. A snapshot that is not being used will be removed immediately.

-trigger *name*
(Optional) Specifies the name of a previously prepared snapshot. This parameter must be specified together with either the volume group ID or the parent UID.

-volumegroup *volumegroup_id_or_name*
(Optional) Specifies the name or ID of the volume group, the volumes within which the snapshots are to be added or triggered.

-volumes *list_of_volume_ids_or_names*
(Optional) Specifies the names or IDs of the volumes to which snapshots are added.

-parentuid *id*
(Optional) Specifies the UID identifying the parent volumes of the snapshot to be triggered. This parameter enables the user to identify a prepared snapshot which does not have a volume group.

-triggerid *snapshot_id*
(Optional) Specifies the snapshot ID of a previously prepared snapshot.

-safeguarded
(Optional) Specifies a snapshot as Safeguarded. This parameter must be specified together with the **-retentiondays** parameter.

-retentiondays *days*
(Optional) Specifies a retention period in days. The value must be a number 1 - 365.

Invocation example: add snapshot to a volumegroup

```
addsnapshot -volumegroup volumegroup12
```

The resulting output:

```
Snapshot, id [1], successfully created or triggered
```

Invocation example: adding a snapshot to a list of volumes

```
addsnapshot -volumes volume0:volume1:volume2:volume3
```

The resulting output:

```
Snapshot, id [1], successfully created or triggered
```

Invocation example: add a snapshot called noon_snapshot to a volumegroup

```
addsnapshot -name noon_snapshot -volumegroup volumegroup12
```

The resulting output:

```
Snapshot, id [2], successfully created or triggered
```

Invocation example: create a prepared snapshot

```
addsnapshot -prepare -pool child -name noon_snapshot -volumegroup volumegroup12
```

The resulting output:

```
Snapshot, id [1519], successfully created or triggered
```

Invocation example: trigger the snapshot using the snapshot name

```
addsnapshot -trigger noon_snapshot -volumegroup volumegroup12
```

The resulting output:

```
Snapshot, id [1519], successfully created or triggered
```

Invocation example: create a prepared snapshot of a list of volumes

```
addsnapshot -prepare -name snap1 -volumes vol1:vol2:vol3
```

The resulting output:

```
Snapshot, id [1520], successfully created or triggered.
```

Invocation example: trigger the snapshot of a list of volumes

```
addsnapshot -trigger snap1 -parentuid 22000
```

The resulting output:

```
Snapshot, id [1520], successfully created or triggered.
```

Invocation example: trigger a snapshot using a snapshot id

```
addsnapshot -triggerid 1520
```

The resulting output:

```
Snapshot, id [1520], successfully created or triggered.
```

An invocation example to create a Safeguarded snapshot with an expiration of 10 days

```
addsnapshot -volumegroup 0  -safeguarded -retentiondays 10
```

The following output is displayed:

```
Snapshot, id [0], successfully created or triggered.
```

Invocation example: create a snapshot with a retention period of 5 minutes

```
svctask addsnapshot -retentionminutes 5 -volumegroup volumegroup12
```

The following output is displayed:

```
Snapshot, id [1], successfully created or triggered.
```

Invocation example: create a prepared snapshot of a list of volumes with a retention period of 3 minutes

```
svctask addsnapshot -prepare -retentionminutes 3 -volumes 1:4
```

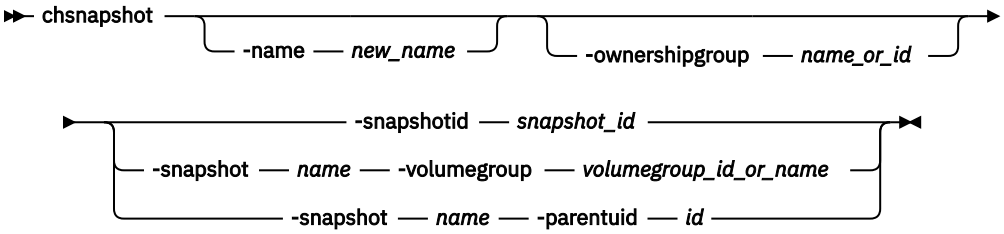
The following output is displayed:

```
Snapshot, id [8], successfully created or triggered.
```

chsnapshot

Use the **chsnapshot** command to change the name or the ownership group of an existing snapshot.

Syntax



Parameters

-name *new_name*
(Optional) Specifies the new name of the snapshot.

-snapshot *name*
(Optional) Specifies the name of the changed snapshot.

-ownershipgroup *name_or_id*
(Optional) Specifies the name or ID of the new ownership group.

-volumegroup *volumegroup_id_or_name*
(Optional) Specifies the name or ID of the volume group with the snapshot to be modified.

-parentuid *id*
(Optional) Specifies the UID identifying the parent volumes of the snapshot to be changed.

-snapshotid *snapshot_id*
(Optional) Specifies the ID of the snapshot to be changed.

An invocation example

```
chsnapshot -name midday -snapshot noon -volumegroup volumegroup12
```

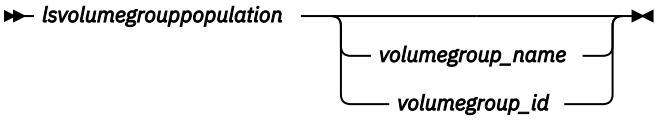
The resulting output:

```
No feedback.
```

lsvolumegrouppopulation

Use the **lsvolumegrouppopulation** command to list the volume group population information. This view includes volume groups of type clone or thin-clone.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

Description

This table provides the attribute values that can be displayed as output view data.

Table 139. ***lsvolumegrouppopulation*** *output*

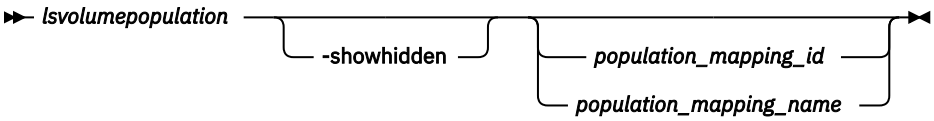
Attribute	Description
volume_group_id	Indicates the ID of the volume group.
volume_group_name	Indicates the name of the volume group.
volume_group_type	Indicates the volume groups of type clone or thin-clone. **Note:** A group can contain volumes that do not have a type as well as volumes of type thin-clone or clone. They cannot contain volumes that are of thin-clone and volumes of type clone. In this case, the type refers to the type for the volumes that are being populated.
start_time	Indicates the time when volume groups image was first populated. The valid time stamp format is YYMMDDHHMMSS.
source_volume_group_id	Indicates the ID of the source volume group.
source_volume_group_name	Indicates the name of the source volume group. Blank if no type is set and no restore process ongoing.
parent_uid	Indicates the ID identifying the source volume.
source_snapshot	Indicates the name of the snapshot that is used to populate the volume group.

Table 139. **lsvolumegrouppopulation** output (continued)	
Attribute	**Description**
estimated_completion_time	Indicates the estimated time when background populate of the clone volumes in the volume group completes. This estimate is calculated by using the used_capacity (used_capacity_before_reduction for DRP) of the volumes. The actual background copy rate is achieved to estimate how long each volume is going to take and then choosing the longest time remaining. The valid time stamp format is YYMMDDHHMMSS.
rate	Indicates the rate at which the volume is populated. By default this rate is 2 MBps. Possible values are 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024, 2048 MBps. If volumes have a mix of population rates, then this value is the lowest rate.
data_to_move	Indicates the amount of data to be copied. As the copy progresses, the value in this field decreases. An empty field indicates no more data to copy. If the user uses the optional **-bytes** parameter, capacity values are displayed in bytes.
restore_start_time	Indicates the time that the latest restore operation was started. It is displayed as blank or in YYMMDDHHMMSS format if no restore operation is ongoing.
restore_snapshot_name	Indicates the name of the snapshot that is used to populate the volume group. It is displayed as blank if no restore operation is ongoing.
restore_estimated_completion_time	Indicates an estimate time when the background restore of the volumes in the volume group completes. The estimated time is calculated by using the used_capacity (used_capacity_before_reduction for DRP) of the volumes and the actual background copy rate achieved. This calculation is used to estimate the duration that is required for each volume and then it select the longest remaining time. It is displayed as blank or in YYMMDDHHMMSS format if no restore operation is ongoing

lsvolumepopulation

Use the **lsvolumepopulation** command to list the population details of clone or thin-clone volumes.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

Description

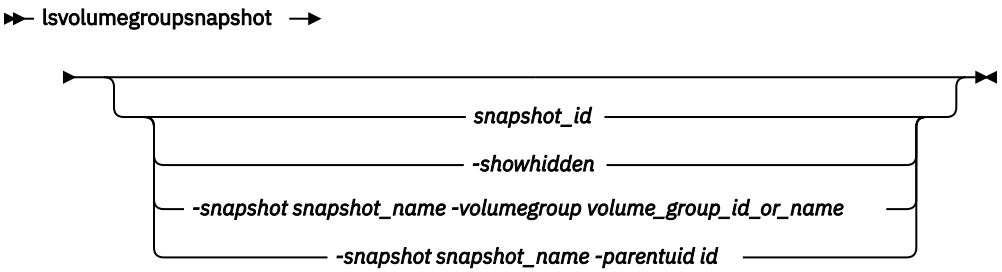
This table provides the attribute values that can be displayed as output view data.

Table 140. **lsvolumepopulation** output	
Attribute	**Description**
volume_id	Indicates the ID of the volume.
volume_name	Indicates the name of the volume.
volume_group_id	Indicates the ID of the volume group.
volume_group_name	Indicates the name of the volume group.
volume_type	Indicates the volume type.
start_time	Indicates the time volume image was first populated. The valid time stamp format is YYMMDDHHMMSS.
source_volume_id	Indicates the ID of the source volume.
source_volume_name	Indicates the name of the source volume. Blank if no type is set and no restore process ongoing.
source_snapshot	Indicates the name of the snapshot used to populate the volume.
estimated_completion_time	Indicates the estimated time when background populate of the clone will complete. The valid time stamp format is YYMMDDHHMMSS.
rate	Indicates the rate at which the volume is populated. By default this will be 2 MBps. Possible values are 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024, 2048 MBps.
data_to_move	Indicates the amount of data to be copied. As the copy progresses the value in this field will decrease. An empty field indicates no more data to copy. If the user uses the optional **-bytes** parameter, capacity values will be displayed in bytes.
populating_mapping_id	Indicates the ID of the flashcopy mapping that is driving the volume population.
populating_mapping_name	Indicates the name of the flashcopy mapping that is driving the volume population.

lsvolumegroupsnapshot

Use the **lsvolumegroupsnapshot** command to list a view snapshot objects in the system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

snapshot_id
(Optional) Specifies the ID of the snapshot.

-showhidden
(Optional) Specifies the snapshots that are in the deleting state.

-volumegroup *volume_group_id_or_name*
(Optional) Specifies the ID or name of the volume group.

-parentuid *id*
(Optional) Specifies the UID identifying the parent volume group.

-snapshot *snapshot_name*
(Optional) Specifies the snapshot name.

Description

This table provides the attribute values that can be displayed as output view data.

Table 141. ***lsvolumegroupsnapshot*** *output*

Attribute	Description
id	Indicates the ID of the volume group snapshot.
name	Indicates the name of the volume group snapshot
volume_group_id	Indicates the ID of the parent volume group. Blank if the group no longer exists.
volume_group_name	Indicates the name of the parent volume group. Blank if the group no longer exists.
time	Indicates the time at which the snapshot was triggered. Blank if the group no longer exists. The valid time stamp format is YYMMDDHHMMSS.
state	Indicates the state of the snapshot. Values can be • ready: If the snapshot is not triggered. • active: Maintain the snapshot image • deleting: Process of deleting • dependent_deleting: It is used to populate a thinclone or clone then it will go into the dependent_deleting state
matches_group	Indicates whether the snapshot is associated with all the volume in the group. It will show no if the snapshot was created from a list of volumes rather than a group. Valid values can be *yes* or *no*.
parent_uid	Indicates the UID identifying the parent volume group.
protection_provisioned_capacity	Indicates the amount of capacity provisioned to store all the volume snapshots in the volume group snapshot.
protection_written_capacity	Indicates the amount of capacity written to store all the volume snapshots in the volume group snapshot before any data reduction.
operation_start_time	Indicates the time when the last delete was started. It will be blank if the state is Active or blank. The valid time stamp format is YYMMDDHHMMSS.

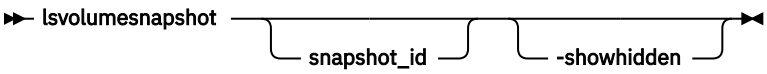
Table 141. **lsvolumegroupsnapshot** output (continued)	
Attribute	**Description**
operation_completion_estimate	Indicates the estimated time to complete the operation. The valid time stamp format is YYMMDDHHMMSS.
owner_id	Indicates the ID of the ownership group.
owner_name	Indicates the name of the ownership group.
expiration_time	Indicates the expiration time. The valid time stamp format is YYMMDDHHMMSS or blank.
auto_snapshot	Indicates the snapshot triggered by the snapshot scheduler. Value can be *yes* or *no*.
safeguarded	Indicates if snapshot is Safeguarded. Values can be *yes* or *no*.
last_restored_from	Indicates the time that the volumegroup snapshot was last used to restore a volumegroup. The timestamp is in the format YYMMDDhhmmss. It will be blank if the volumegroup snapshot has never been used for a restore.

Note: The **expiration_time** field value is assigned at the time that a volume snapshot is created based on the source volume being part of a volume group with a snapshot policy. This **expiration_time** remains unchanged until the volume snapshot is removed manually or automatically unless the system clock changes including during a power-off situation. In that case, the expiration times of volume snapshots are adjusted to accommodate the difference. In the scenario of a power-off, the expiration times can be extended up to the length of the power-off.

lsvolumesnapshot

Use the **lsvolumesnapshot** command to display a list of all volume snapshots.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

snapshot_id
(Optional) Specifies the ID of the snapshot.

-showhidden
(Optional) Specifies that snapshots that are in the process of deleting are displayed. This parameter also displays additional attributes of the volume that contains snapshots such as *pool_1_id*, *pool_2_id*, *pool_1_name*, *pool_2_name*, *snapshot_mapping_id*, *snapshot_mapping_name*, *snapshot_vdisk_id*, and *snapshot_vdisk_name*. You must include the **-showhidden** parameter to display these attributes.

Description

The snapshot ID can be specified to show the specific volume snapshot objects for the specific ID.

Filter value can be used to restrict the objects that are returned to a specific volume group. To view these object, use the **-filtervalue** parameter.

This table provides the attribute values that can be displayed as output view data.

*Table 142. **lsvolumesnapshot** output*

Attribute	Description
snapshot_id	Indicates the ID of the volume snapshot.
snapshot_name	Indicates the name of the volume snapshot.
volume_id	Indicates the ID of the parent volume.
volume_name	Indicates the name of the parent volume.
volume_group_id	Indicates the ID of the volume group that is associated with the snapshot. Blank if the group no longer exists.
volume_group_name	Indicates the name of the volume group that is associated with the snapshot. Blank if the group no longer exists.
time	Indicates the snapshot trigger time. The valid timestamp format is YYMMDDHHMMSS.
parent_uid	Indicates the unique ID of the parent volumes.
state	Indicates the state of the snapshot. Values can be: • Ready: If the snapshot is not triggered. • Active: Maintain the snapshot image. • Deleting: Process of deleting. • Failed.
expiration_time	Indicates the expiration time of the snapshot. The valid timestamp format is YYMMDDHHMMSS.
safeguarded	Indicates if snapshot is Safeguarded. Values can be *yes* or *no*.
mirrored	Indicates whether the snapshot is mirrored. Values can be *yes* or *no*.
pool_1_id	Indicates the ID of the pool for copy 1.
pool_1_name	Indicates the name of the pool for copy 1.
pool_2_id	Indicates the ID of the pool for copy 2.
pool_2_name	Indicates the name of the pool for copy 2.
snapshot_mapping_id	Indicates the ID of the FlashCopy mapping that is used to maintain the snapshot.
snapshot_mapping_name	Indicates the name of the FlashCopy mapping that is used to maintain the snapshot.
snapshot_vdisk_id	Indicates the ID of the volume that is used to maintain the snapshot.
snapshot_vdisk_name	Indicates the name of the volume that is used to maintain the snapshot.

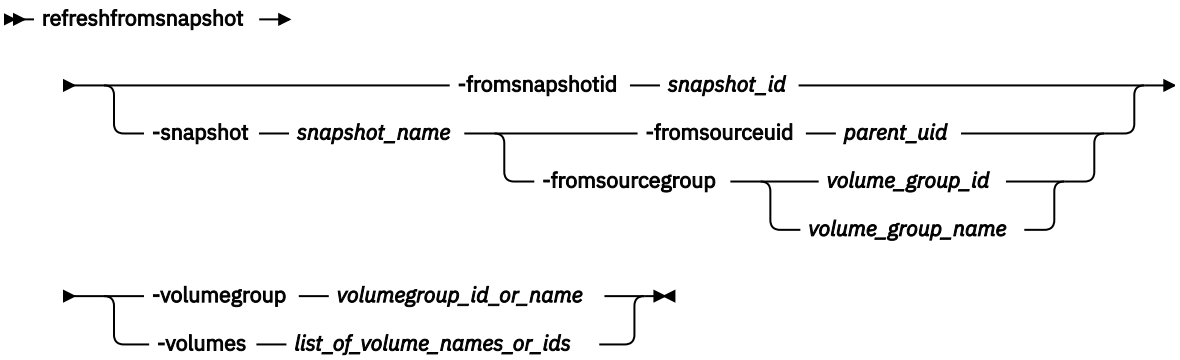
Table 142. **lsvolumesnapshot** output (continued)	
Attribute	**Description**
volume_size_mismatch	Indicates whether the provisioned capacity of the parent volume and the volume snapshot are different. Values are yes or no. **Note:** A size mismatch will prevent the volume snapshot being used in a restore operation.
written_capacity	Indicates the amount of data to be written.

Note: The **expiration_time** field value is assigned at the time that a volume snapshot is created based on the source volume being part of a volume group with a snapshot policy. This **expiration_time** remains unchanged until the volume snapshot is removed manually or automatically unless the system clock changes including during a power-off situation. In that case, the expiration times of volume snapshots are adjusted to accommodate the difference. In the scenario of a power-off, the expiration times can be extended up to the length of the power-off.

refreshfromsnapshot

Use the **refreshfromsnapshot** command to refresh the data on a set of volumes, whether the volumes are in a volume group or in a prepared list of volumes. The data used for the update comes from the specified snapshot. Any existing data on the set of volumes will be discarded.

Syntax



Parameters

-volumegroup ***volume_group_id_or_name***
(Optional) Specifies the volume group ID or name to be refreshed.

-volumes volume_ids_or_names
(Optional) Specifies a list of up to 512 volume ID or names. These are the volumes that will be refreshed.

-fromsourcegroup volume_group_id_or_name
(Optional) Specifies the parent volume group of the snapshot from which the volumes are to be refreshed.

-fromsourceuid parent_uid
(Optional) Specifies the UID used to identify the parent volumes of the snapshot from which the volumes are to be refreshed.

-snapshot snapshot_name
(Optional) Specifies the name of the source snapshot from which the volumes are to be refreshed.

-fromsnapshotid snapshot_id
(Optional) Specifies the ID of the source snapshot from which the volumes are to be refreshed.

Description

This command refreshes the data on a set of volumes, whether the volumes are in a volume group or in a prepared list of volumes.

The data used for the update comes from the specified snapshot. Any existing data on the set of volumes will be discarded. This command can only be used to refresh thinclone volumes.

An invocation example

```
refreshfromsnapshot -volumegroup volumegroup12 -snapshot snapshot2 -fromsourcegroup volumegroupA
```

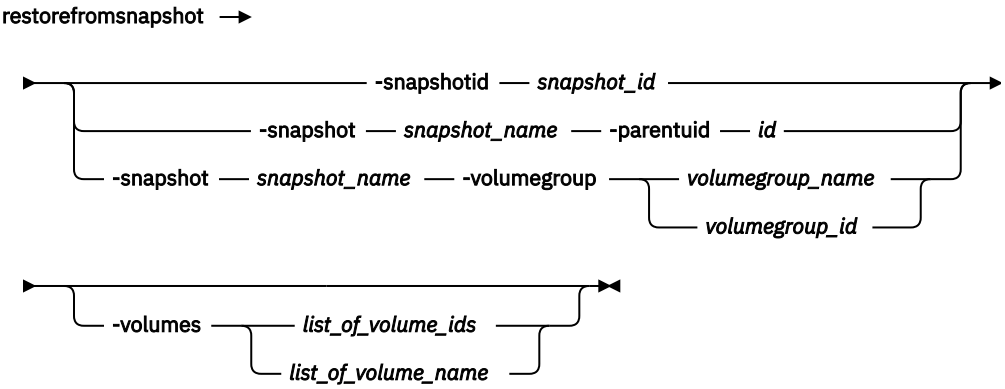
The detailed resulting output

```
No feedback.
```

restorefromsnapshot

Use the **restorefromsnapshot** command to restore a parent volume group to a snapshot taken at an earlier point in time.

Syntax



Parameters

-snapshot *name*
Specifies the name of the snapshot to be restored.

-volumegroup *volumegroup_id_or_name | -parentuid id*
Specifies the name or ID of the volume group with the snapshot to be restored.

-parentuid *id*
Specifies the UID identifying the parent volumes of the snapshot.

-snapshotid *snapshot_id*
Specifies the ID of the snapshot to be restored.

-volumes *list_of_volume_ids | list_of_volume_name*
Specifies the name or ID of the parent volumes.

Note: If **restorefromsnapshot** command is used on a High Availability (HA) volume group, HA is temporarily lost.

An invocation example

```
restorefromsnapshot -snapshotid 5
```

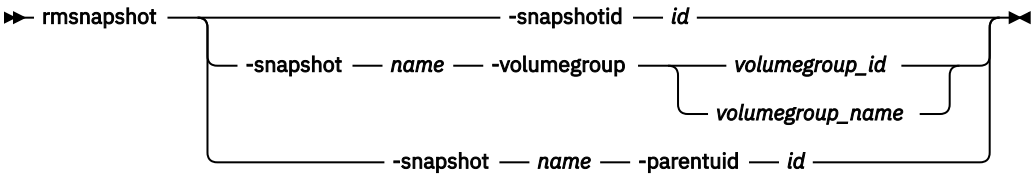
The resulting output:

```
No feedback.
```

rmsnapshot

Use the **rmsnapshot** command to remove the snapshot specified either by the snapshot name and the volume group or by the snapshot name and snapshot ID.

Syntax



Parameters

-snapshot *name*
(Optional) Specifies the name of the snapshot to be removed.

-volumegroup *volumegroup_id_or_name*
(Optional) Specifies the name or ID of the volume group with the snapshot to be removed.

-parentuid *id*
(Optional) Specifies the UID identifying the parent volumes of the snapshot.

-snapshotid *snapshot_id*
(Optional) Specifies the ID of snapshot to be removed.

Note: To remove a Safeguarded snapshot, you must have Security Administrator or superuser role. This access is only required if the snapshot is Safeguarded.

An invocation example

```
rmsnapshot -snapshot noon -volumegroup volumegroup12
```

The resulting output:

```
No feedback.
```

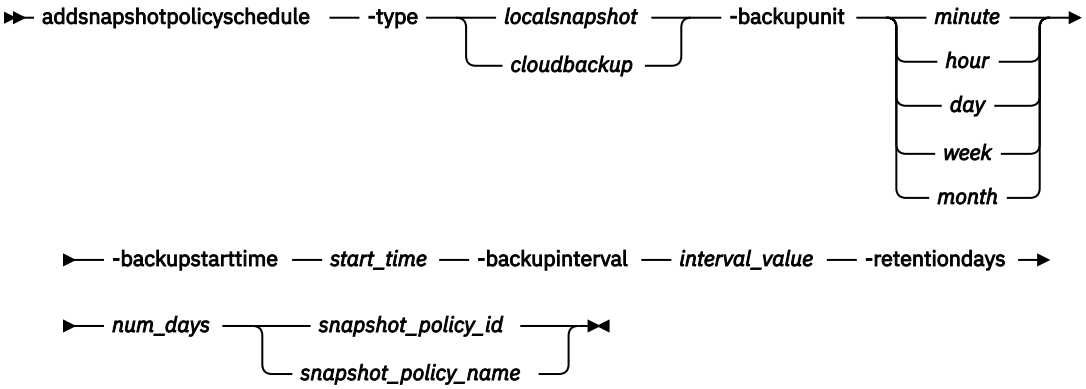
Chapter 40. Snapshot policy commands

Use the following set of commands to set a schedule for snapshot creation and retention.

addsnapshotpolicyschedule

Use the **addsnapshotpolicyschedule** command to add a new schedule in an existing snapshot policy. It is added as a second schedule in addition to the existing schedule in the policy.

Syntax



Parameters

-type ***localsnapshot | cloudbackup***
(Required) Specifies the type of schedule.

-backupunit minute | hour | day | week | month
(Required) Specifies the unit of time of the backup interval.

-backupstarttime start_time
(Required) Specifies when the first snapshot backup is to be taken. The **backup_start_time** parameter is displayed in the standard *YYMMDDHHMM* format, where: YY is year (The first) MM is month DD is day HH is hour (The second) MM is minute.

-backupinterval interval_value
(Required) Indicates the interval between snapshot backups. The value must be an integer (1 - 65535).

-retentiondays num_days
(Required) Specifies the number of days to keep the snapshot backup. The value must be an integer (1 - 365).

snapshot_policy_id
(Required) Specifies the snapshot policy ID. The value must be an integer (0-39).

snapshot_policy_name
(Required) Specifies the snapshot policy name.

Note: Backup-start-time supplied to **mksnapshotpolicy** will apply to the first schedule that it will create. Subsequent schedule is created using this CLI which gets its own backup-start-time. However, note that backup start time optionally supplied to **mkvolumegroup -snapshotpolicy** will apply to all the schedules in that policy.

© Copyright IBM Corp. 2003, 2025

An invocation example

The following example adds a new schedule to an existing snapshot policy that generates a `localsnapshot` every day and retains the local snapshots for 15 days.

```
addsnapshotpolicyschedule -type localsnapshot -backupunit day -backupinterval 1 -retentiondays
15  3
```

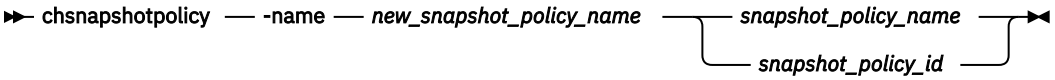
The resulting output:

```
No feedback.
```

chsnapshotpolicy

Use the **chsnapshotpolicy** command to change an existing snapshot policy.

Syntax



Parameters

-name new_snapshot_policy_name
(Required) Specifies a new name for the snapshot policy.

snapshot_policy_name | snapshot_policy_id
Specifies the ID or name of an existing snapshot policy.

Note: Only the name of an existing snapshot policy can be changed.

Note: The **-backupstarttime** parameter indicates the time of snapshot backups. It is displayed in the standard *YYMMDDHHMMSS* format. where: Y is *year* (The first) M is *month* D is *day* H is *hour* (The second) M is *minute* S is *second*.

An invocation example

```
chsnapshotpolicy -name mysgpolicy3 3
```

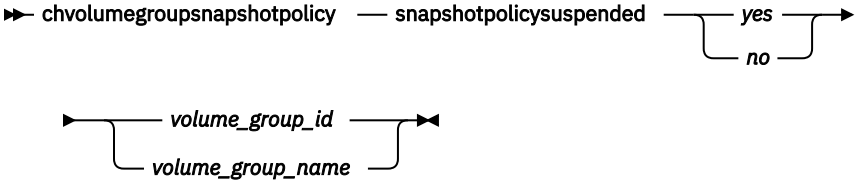
The resulting output:

```
No feedback
```

chvolumegroupsnapshotpolicy

Use the **chvolumegroupsnapshotpolicy** command to change the volume group properties related to snapshot policies.

Syntax



Parameters

-snapshotpolicysuspended
(Required) Specifies whether to suspend the snapshot policy that is associated with the volume group. When *yes* is specified, the policy stops creating new snapshots and stops deleting expired snapshots. The *no* value restarts a suspended policy. The value can be *yes* or *no*.

volume_group_id | volume_group_name
Specifies the ID or name of a volume group to change.

Note: Suspending the snapshot policy on a volume group stops the system from creating any new snapshots based on the snapshot policy and from deleting any expiring snapshots associated with this volume group. The parameter can be set to *yes* even when the volume group has no associated snapshot policy. This stops the system from deleting any previously created expiring snapshots associated with this volume group. The value of this parameter is not affected by associating or disassociating a snapshot policy to the volume group. The value of this parameter has no effect on the manual creation or deletion of snapshots.

An invocation example

```
chvolumegroupsnapshotpolicy -snapshotpolicysuspended yes 0
```

The resulting output:

```
No feedback
```

lssnapshotpolicy

Use the **lssnapshotpolicy** command to list the snapshot policies.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

snapshot_policy_id
(Required) Specifies the snapshot policy ID.

snapshot_policy_name
(Required) Specifies the snapshot policy name.

Description

Use this command to display a concise list or a detailed view of snapshot policies.

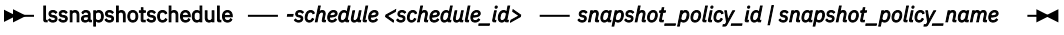
This table provides the attribute values that can be displayed as output view data.

Table 143. **lssnapshotpolicy** output	
Attribute	**Description**
id	Indicates the unique ID of the snapshot policy.
name	Indicates the unique name of the snapshot policy.
volume_group_count	Indicates the number of volume groups associated with the snapshot policy.
schedule_id	Indicates the schedule ID of the snapshot policy.
schedule_type	Indicates the type of schedule.
backup_unit	Indicates the unit of time of the backup_interval. The **backup_unit** parameter is displayed as minute, hour, day, week or month.
backup_interval	Indicates the interval between the snapshot backups. The value must be a number 1 - 65535.
backup_start_time	Indicates the time of snapshot backups. The **backup_start_time** parameter is displayed in the standard *YYMMDDHHMMSS* format. where: Y is *year* (The first) M is *month* D is *day* H is *hour* (The second) M is *minute* S is *second*.
retention_days	Indicates the number of days to keep the snapshot backups. The value must be a number 1 - 65535.

lssnapshotschedule

Use the **lssnapshotschedule** command to list the snapshot schedules.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

schedule_id
(Required) Specifies the snapshot schedule ID.

snapshot_policy_id
(Required) Specifies the snapshot policy ID

snapshot_policy_name
(Required) Specifies the snapshot policy name.

Description

This table provides the attribute values that can be displayed as output view data.

*Table 144. **lssnapshotschedule** output*	
Attribute	**Description**
policy_id	Indicates the unique ID of the snapshot policy.
policy_name	Indicates the unique name of the snapshot policy.
schedule_id	Indicates the scheduled ID of the snapshot policy.
backup_unit	Indicates the unit of time of the backup_interval. The **backup_interval** parameter is displayed in the minute, hour, day, week and month.
backup_interval	Indicates the interval between the snapshot backups. The value must be a number *1 - 65535.*
backup_start_time	• Indicates the time the first backup will be taken. If this **backup_start_time** is in the past, internally a new start time will be calculated by moving the original **backup_start_time** forward in increments of the **backup_interval** value until a current start time is found. **Note:** The **backup_start_time** value can be overridden at the time a volume group is associated with the policy. Both **mkvolumegroup** and **chvolumegroup** commands have an input option to add a policy start time. • The **backup_start_time** parameter is displayed in the standard YYMMDDHHMMSS format. where: YY is year, (the first) MM is month, DD is day, HH is hours, (the second) MM is minutes, and SS is seconds.
retention_days	Indicates the number of days to keep the snapshot backups. The value must be a number *1 - 365.*
schedule_type	Indicates the type of schedule.

lsvolumegroupsnapshotpolicy

Use the **lsvolumegroupsnapshotpolicy** command to list the snapshot policy attributes associated with a volume group.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

volume_group_id
(Required) Specifies the volume group ID.

volume_group_name
(Required) Specifies the volume group name.

Description

Use this command to display a concise list or a detailed view of snapshot policy.

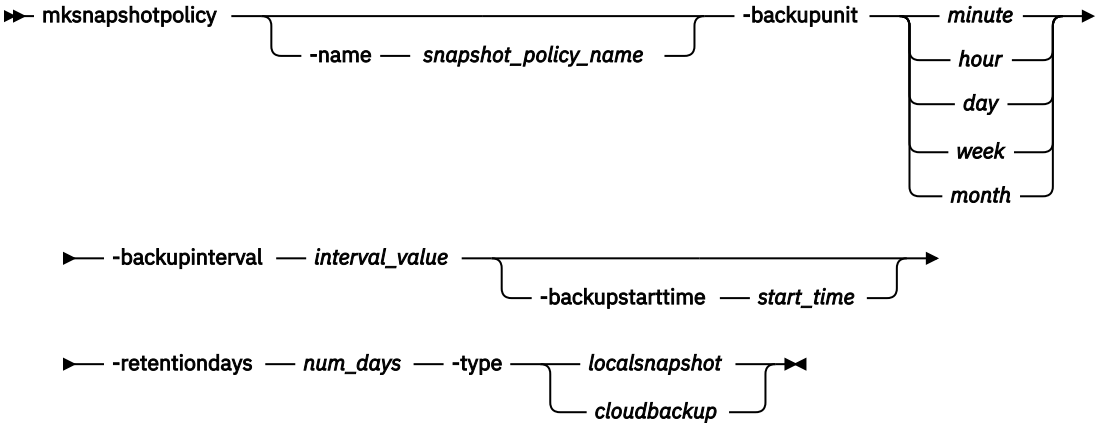
This table provides the attribute values that can be displayed as output view data.

*Table 145. **lsvolumegroupsnapshotpolicy** output*	
Attribute	**Description**
id	Indicates the ID of the volume group.
name	Indicates the name of the volume group.
snapshot_policy_id	Indicates the ID of the snapshot policy.
snapshot_policy_name	Indicates the name of the snapshot policy.
snapshot_policy_start_time	Indicates the value specified with the **-policystarttime** parameter. It is displayed in the standard *YYMMDDHHMMSS* format.
snapshot_policy_next_time	Indicates the next scheduled snapshot by snapshot policy. It is displayed in the standard *YYMMDDHHMMSS* format.
snapshot_policy_last_time	Indicates the last snapshot triggered by snapshot policy. It is displayed in the standard *YYMMDDHHMMSS* format.
snapshot_policy_suspended	Indicates the snapshot policy suspended for volume group. The value can be *yes* or *no*.
error_sequence_number	Indicates the snapshot policy error sequence number for the volume group.
snapshot_policy_safeguarded	Indicates whether snapshot policy is Safeguarded. Values can be *yes* or *no*.
volume_group_id	Indicates the volume group id to which the policy and schedules are associated.

mksnapshotpolicy

Use the **mksnapshotpolicy** command to create a new snapshot policy that can be associated with volume groups.

Syntax



Parameters

-name *snapshot_policy_name*
(Optional) Indicates a unique name for the snapshot policy. If name is not specified then a unique name will be generated and assigned to the snapshot policy object.

-backupunit minute | hour | day | month | year
(Required) Specifies the unit of time of the backup interval.

-backupinterval interval_value
(Required) Indicates the interval between snapshot backups. The value must be an integer (1 - 65535).

-backupstarttime start_time
(Optional) Specifies when the first snapshot backup is to be taken. The **backup_start_time** parameter is displayed in the standard *YYMMDDHHMM* format. where: Y is *year* (The first) M is *month* D is *day* H is *hour* (The second) M is *minute*.

-retentiondays num_days
(Required) Specifies the number of days to keep the snapshot backup. The value must be an integer (1 - 365).

-type localsnapshot | cloudbackup
(Optional) Specifies the type of schedule being created. By default, the type is *localsnapshot*.

An invocation example

The following example creates a snapshot policy that, by default, generates a local snapshot every day and retains the local snapshot for 15 days with start time 2102281800

```
mksnapshotpolicy -backupunit day -backupinterval 1 -backupstarttime 2102281800 -retentiondays 15
```

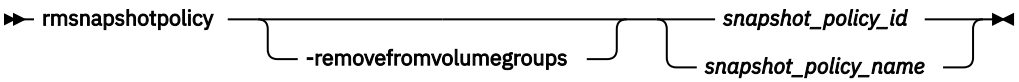
The resulting output:

```
snapshot Policy, id[3], successfully created
```

rmsnapshotpolicy

Use the **rmsnapshotpolicy** command to remove an existing snapshot policy by ID or by name.

Syntax



Parameters

-removefromvolumegroups
(Optional) Removes the volume group associations to this snapshot policy. This option is available to Security Administrator users only.

snapshot_policy_id | snapshot_policy_name
(Required) The ID or name of a snapshot policy to remove.

Note: The predefined snapshot policies cannot be removed.

An invocation example in which the security administrator remove existing snapshot policy that has volume group associations:

```
rmsnapshotpolicy 1
```

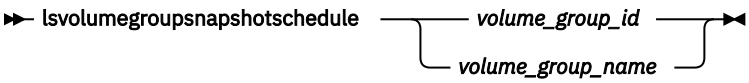
The resulting output:

```
CMMVC9695E The snapshot policy cannot be removed because of one or more associated volume
groups.
```

lsvolumegroupsnapshotschedule

Use the **lsvolumegroupsnapshotschedule** command to view the concise list of multiple schedule details of policies associated with volume groups.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

volume_group_id
(Required) Specifies the unique ID of the volume group.

volume_group_name
(Required) Specifies the name of the volume group.

Description

This table provides the attribute values that can be displayed as output view data.

Table 146. ***lsvolumegroupsnapshotschedule*** *output*

Attribute	Description
id	Indicates the sequence in which the snapshot schedule for volume groups are listed.
volume_group_id	Indicates the volume group id to which the policy and schedules are associated.
volume_group_name	Indicates the unique name of the volume group.
snapshot_policy_id	Indicates the unique ID of the snapshot policy.
snapshot_policy_name	Indicates the unique name of the snapshot policy.
schedule_id	Indicates the scheduled ID of the snapshot policy.
schedule_type	Indicates the type of schedule.
snapshot_schedule_start_time	Indicates the start time of the schedule. It is displayed in the standard *YYMMDDHHMMSS* format.

*Table 146. **lsvolumegroupsnapshotschedule** output (continued)*	
Attribute	**Description**
snapshot_schedule_next_time	Indicates the next scheduled snapshot by snapshot policy. It is displayed in the standard *YYMMDDHHMMSS* format.
snapshot_schedule_last_time	Indicates the last snapshot triggered by snapshot policy. It is displayed in the standard *YYMMDDHHMMSS* format.
error_sequence_number	Indicates the snapshot policy error sequence number for volume group.
snapshot_policy_suspended	Indicates the snapshot policy suspended for volume group.

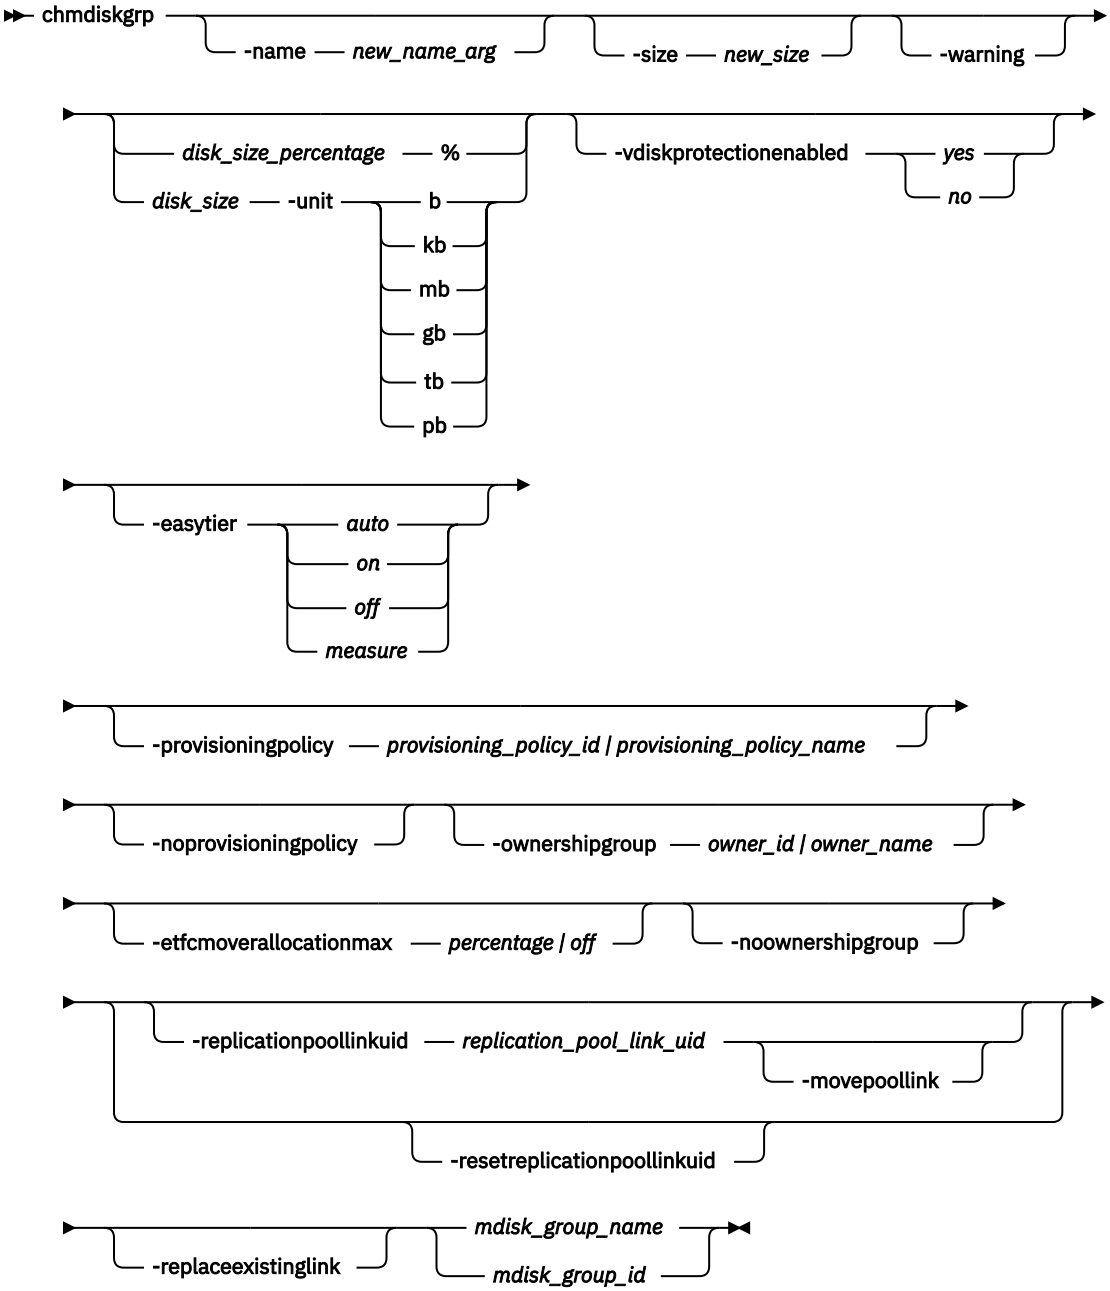
Chapter 41. Storage pool commands

Use the storage pool commands to work with storage pool options on the system.Use storage pool commands to work with managed disk group options.

chmdiskgrp

Use the **chmdiskgrp** command to modify properties of an existing storage pool.

Syntax



© Copyright IBM Corp. 2003, 2025

Parameters

-name *new_name_arg*
(Optional) Specifies the new name of the storage pool.

-warning *disk_size | disk_size_percentage %*
(Optional) Sets a threshold at which a warning is generated. The warning is generated the first time that the threshold is exceeded by the used-disk capacity in the storage pool. You can specify a *disk_size* integer, which defaults to megabytes (MB) unless the **-unit** parameter is specified; or you can specify a *disk_size_percentage* %, which is a percentage of the storage pool size. To disable warnings, specify **0** or **0%**. This parameter is not applicable for data reduction child pools.

-size *new_size*
(Optional) Specifies the new size of a child pool. This parameter is not applicable for data reduction child pools.

Note: This parameter cannot be used with parent pools. Use **addmdisk** or **rmmdisk** to change storage pool capacity.

-unit b | kb | mb | gb | tb | pb
(Optional) Specifies the data units for the **-warning** parameter.

-vdiskprotectionenabled *yes | no*
(Optional) Specifies whether volume protection in enabled for this storage pool.

Note: Volume protection can not be enabled for pools with the owner **vvol_child_pool**. Thus, the default value is *no* while creating a pool with this owner and the **chmdiskgrp** command fails if the argument **-vdiskprotectionenabled yes** is passed and the pool has the owner **vvol_child_pool**.

-easytier *auto | on | off | measure*
(Optional) Specifies whether the Easy Tier function is on or off for this storage pool, or if it is automatically determined. **-easytier** is `active` in storage pools with multiple tiers and is `balance` with single tiers.

Note: -easytier must be followed by one of the following:

- If **-easytier** is set to *auto*, system automatically enables Easy Tier functions when the storage pool contains MDisks from more than one tier, and enables automatic rebalancing when the storage pool contains MDisks from only one tier.
- If **-easytier** is set to *on*, then Easy Tier functions are active.
- If **-easytier** is set to *off*, then Easy Tier functions are inactive.
- If **-easytier** is set to *measure* Easy Tier statistics are collected but Easy Tier management is disabled. (No extents are moved by Easy Tier.)

auto equates to:

- *on* if Easy Tier is licensed or no license is required.
- *off* if Easy Tier is not licensed and a license is required.

Specifying `-easytier on` enables Easy Tier:

- Management of both single-tier and multi-tier pools
- Auto rebalance

Extents are moved to balance the I/O load on the MDisks in the pool.

-provisioningpolicy *policy_id | policy_name*
(Optional) The name or ID of the provisioning policy that the pool uses to manage capacity savings for volumes in the pool.

-noprovisioningpolicy
(Optional) If specified, the provisioning policy is removed from the pool.

-ownershipgroup ***owner_id | owner_name***
(Optional) The name or ID of the ownership group to which the object is being added. This parameter can be specified on child pools only.

-noownershipgroup
(Optional) If specified, the object is removed from the ownership group to which it belongs. This parameter can be specified on child pools only.

mdisk_group_id | mdisk_group_name
(Required) Specifies the ID or name of the storage pool to modify.

-etfcmoverallocationmax ***percentage | off***
(Optional) Sets the maximum over allocation which Easy Tier can migrate onto FlashCore Module arrays, when the array is used as the top tier in a multitier pool. The value acts as a multiplier of the physically available space. The allowed values are a percentage in the range of 100% (default) to 400% or **off**. Setting the value to **off** disables this feature.

-replicationpoollinkuid ***replication_pool_link_uid***
(Optional) Specifies the new replication pool link unique identifier (UID) value to assign to the pool. Pools are linked between partnered systems, for use with policy-based replication, when the same pool link UID value is assigned to the pools on each system.

-movepoollink
(Optional) Specifies that the replication pool link UID value, specified with the **-replicationpoollinkuid** parameter, is to be moved from an existing pool to the specified pool. This parameter is required to move an existing pool link to a different pool. The pool that is currently linked gets unlinked and the specified pool is linked.

-replaceexistinglink
(Optional) Allows the link UID to be changed even if the pool already has links used by a replication policy.

-resetreplicationpoollinkuid
(Optional) If specified, the existing replication pool link UID for the pool is reset and a new pool link UID value is generated. This action removes any links that exist between the local pool and pools on partnered systems.

Description

Table 147. Parameter differences for child pools and parent pools

Specifies the ID or name of the storage pool to Parameter	Child pool usage	Storage pool usage
-name	Optional	Optional
-easytier	Cannot be used with child pools	Optional
-size	Optional	Cannot be used with parent pools
-unit	Optional	Optional
-warning	Optional	Optional

This command modifies the name, or label, assigned to a given storage pool. You can use the new name to refer to the storage pool.

The command can also be used to set the warning threshold for the storage pool. The warning threshold is the threshold which a warning is generated when it is exceeded by the used-disk capacity in the storage pool.

You can also use this command to change other settings for parent pools and child pools.

Note: Only a user with the Security Admin or superuser role is allowed to use the **-size** parameter to shrink the space of a Safeguarded backup location. If the specified pool is a Safeguarded backup location, **-ownershipgroup** parameter is not allowed.

An invocation example

```
chmdiskgrp -name testmdiskgrp -easytier on Group0
```

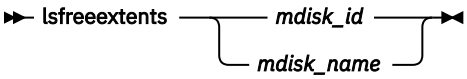
The resulting output:

```
No feedback
```

lsfreeextents

Use the **lsfreeextents** command to list the number of free extents that are available on a specified MDisk.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

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.

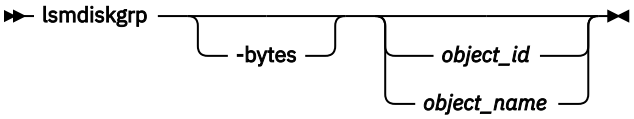
Description

This command displays a count of the number of free extents on the specified MDisk.

lsmdiskgrp

Use the **lsmdiskgrp** command to display a concise list or a detailed view of storage pools that are visible to the system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-bytes
(Optional) Specifies that you want the report to display 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 that match the filtering requirements that are specified by the **-filtervalue** parameter are displayed.

Description

This command returns a concise list or a detailed view of storage pools visible to the system.

Command output includes values for the following attributes:

status
The state of the MDisk with the highest-priority status in the group, excluding image mode MDisks.

vdisk_count
The number of volume copies that are in the storage pool.

capacity
The total amount of MDisk storage that is assigned to the storage pool. For quotaless child pools, this is the free capacity of the parent pool.

extent_size
The sizes of the extents for this group are the following values: 16, 32, 64, 128, 256, 512, 1024, 2048, 4096, or 8192 (MB).

free_capacity
The amount of MDisk storage that is immediately available. Additionally, `reclaimable_capacity` can eventually become available. For quotaless child pools, this is the free capacity of the parent pool.

real_capacity
The total MDisk storage capacity assigned to volume copies. Reports 0.00 MB for quotaless child pools.

Note: It includes `reclaimable_capacity`.

virtual_capacity
The total host mappable capacity of all volume copies in the storage pool.

used_capacity
The amount of data that is stored on MDisks. standard-provisioned volumes contribute their entire capacity. Reports 0.00 MB for quotaless child pools.

Note: It includes `reclaimable_capacity`.

overallocation
Expressed as a percentage, the ratio of the `virtual_capacity` value to the capacity. A storage pool overallocation of over 100 is only possible if you configure thin-provisioned volume copies. For data reduction parent pools, this calculation includes the virtual capacity of any quotaless child pools. No value is reported for quotaless child pools.

warning
This field is a percentage. A warning is generated when the amount of space in the storage pool that is assigned exceeds this level. For quotaless child pools, this is the parent pool's warning level.

easy_tier
This value is set by the user and determines whether Easy Tier is permitted to manage the pool.

Note: The values are:

1. on indicates that Easy Tier actively manages the extents (including single-tier storage pools), and the Easy Tier status must be `active` - unless a license is required.
2. `off` indicates that Easy Tier does not actively manage the extents, and the Easy Tier status must be `inactive` .
3. `auto` indicates that the value of Easy Tier status is determined by the number of tiers in a storage pool.

   Note: The following values apply to `auto`:

   - on if Easy Tier is licensed or no license is required.
   - `off` if Easy Tier requires a license and none exists.
4. `measure` indicates that Easy Tier s collects statistics on that storage pool but does not move any extents in the storage pool.

easy_tier_status
This field indicates whether the Easy Tier functions are active on a storage pool.

- `active` indicates that a pool is being managed by Easy Tier to provide tier management performance-based pool balancing (for example, extents can be moved).
- `inactive` indicates that Easy Tier is inactive.
- `measured` indicates that Easy Tier statistics are being collected but no Easy Tier management is detected.
- `balanced` indicates that a pool is being managed by Easy Tier to provide performance-based pool balancing (for example, extents can be moved).

tier
Indicates which tier information is being reported. The values are:

- `tier0_flash`
- `tier1_flash`
- `tier_enterprise`
- `tier_nearline`
- tier_scm

tier_mdisk_count
Indicates the number of MDisks in the tier.

tier_capacity
The amount of MDisk storage in this tier that is assigned to the storage pool.

tier_free_capacity
The amount of MDisk storage in this tier that has not been assigned.

compression_active
Indicates whether any compressed volume copies are in the storage pool. This field is blank for storage pools that are data reduction pools.

compression_virtual_capacity
Indicates the total provisioned (virtual) capacity for all compressed volume copies in regular storage pools.

compression_compressed_capacity
Indicates the total used capacity for all compressed volume copies in regular storage pools.

compression_uncompressed_capacity
Indicates the total uncompressed used capacity for all compressed volume copies in regular storage pools.

site_id
Indicates the site value for the storage pool group. This numeric value is 1, 2, 3, or blank.

site_name
Indicates the site name for the storage pool. This value is an alphanumeric value or is blank.

parent_mdisk_grp_id
Indicates the storage pool group ID. This value is a numeric string (in the range 0 - 127 characters) or blank.

parent_mdisk_grp_name
Indicates the storage pool group name. This value is an alphanumeric string (in the range 1 - 63 characters) or blank.

child_mdisk_grp_count
Indicates the number of child pools in the parent pools. This value is a numeric string (in the range 0 - 127 characters) or blank.

child_mdisk_grp_capacity
Indicates the total amount of space that is reserved for child pools.

type
Indicates the MDisk group type. The values are `parent`, `child_thick`, and `child_quotaless`.

A child pool that does not have a quota does not contribute any extents to the parent pool.

The following define the status fields, from lowest to highest priority:

Online
Indicates that the storage pool is online and available.

Offline
Indicates that all paths to the storage pool are lost.

owner_type
Indicates the type of owning object, such as a file system or application. This attribute is an alphanumeric string up to 20 characters in length.

data_reduction
Indicates that the storage pool is a data reduction pool. The values are `yes` or `no`.

physical_capacity
Indicates the total physical capacity of MDisks that belong to this storage pool. For any disks that do not display their physical capacity, displays the logical capacity value. The value must be a number (indicated in units) that is rounded to two decimal places.

physical_free_capacity
Indicates the total free physical capacity of MDisks that belong to this storage pool. For any disks that do not display their physical capacity, displays the logical capacity value. The value must be a number (indicated in units) that is rounded to two decimal places.

shared_resources
Indicates that one or more MDisks in this storage pool shares a resource provisioning group with an MDisk in another storage pool. It indicates cross contamination over-provisioning of physical resources between the storage pools. The value must be `yes` or `no`.

overhead_capacity
The MDisk capacity that is reserved for internal usage. Reports 0.00 MB for data reduction child pools.

vdisk_protection_enabled
Indicates if this pool has volume protection enabled. This field is combined with the system setting **vdisk_protection_enabled** to determine the pool **vdisk_protection_status**. The value is either yes or no.

vdisk_protection_status
Indicates whether volume protection is active for this pool. Volumes with one or more copies in a pool where **vdisk_protection_status** is active are protected. Volumes with all copies in pools with an inactive volume protection status are not protected volumes. This field depends on the system setting for **vdisk_protection_enabled** and the pool setting for **vdisk_protection_enabled**. It is active only when the storage pool has **vdisk_protection_enabled** set to `yes` and the system setting has **vdisk_protection_enabled** set to `yes`. Otherwise, it is inactive.

easy_tier_fcm_over_allocation_max
Displays the maximum over allocation which Easy Tier can migrate onto FlashCore Module arrays, when the array is used as the top tier in a multitier pool. The value acts as a multiplier of the physically available space. This value is a percentage or blank for unsupported products.

provisioning_policy_id
Displays the ID of the provisioning policy.

provisioning_policy_name
Displays the name of the provisioning policy.

protection_provisioned_capacity
The amount of capacity provisioned to store point-in-time copies for volumes in the storage pool. This includes snapshots and consistency protection for replication.

protection_written_capacity
The amount of capacity written to store point-in-time copies for volumes in the storage pool before any data reduction. This includes volume snapshots and consistency protection for replication.

protection_vdisk_count
The number of point-in-time copies for volumes in the storage pool that contribute to the protection capacity values. This includes snapshots and consistency protection for replication. These hidden volumes are included in the value in the vdisk_count field.

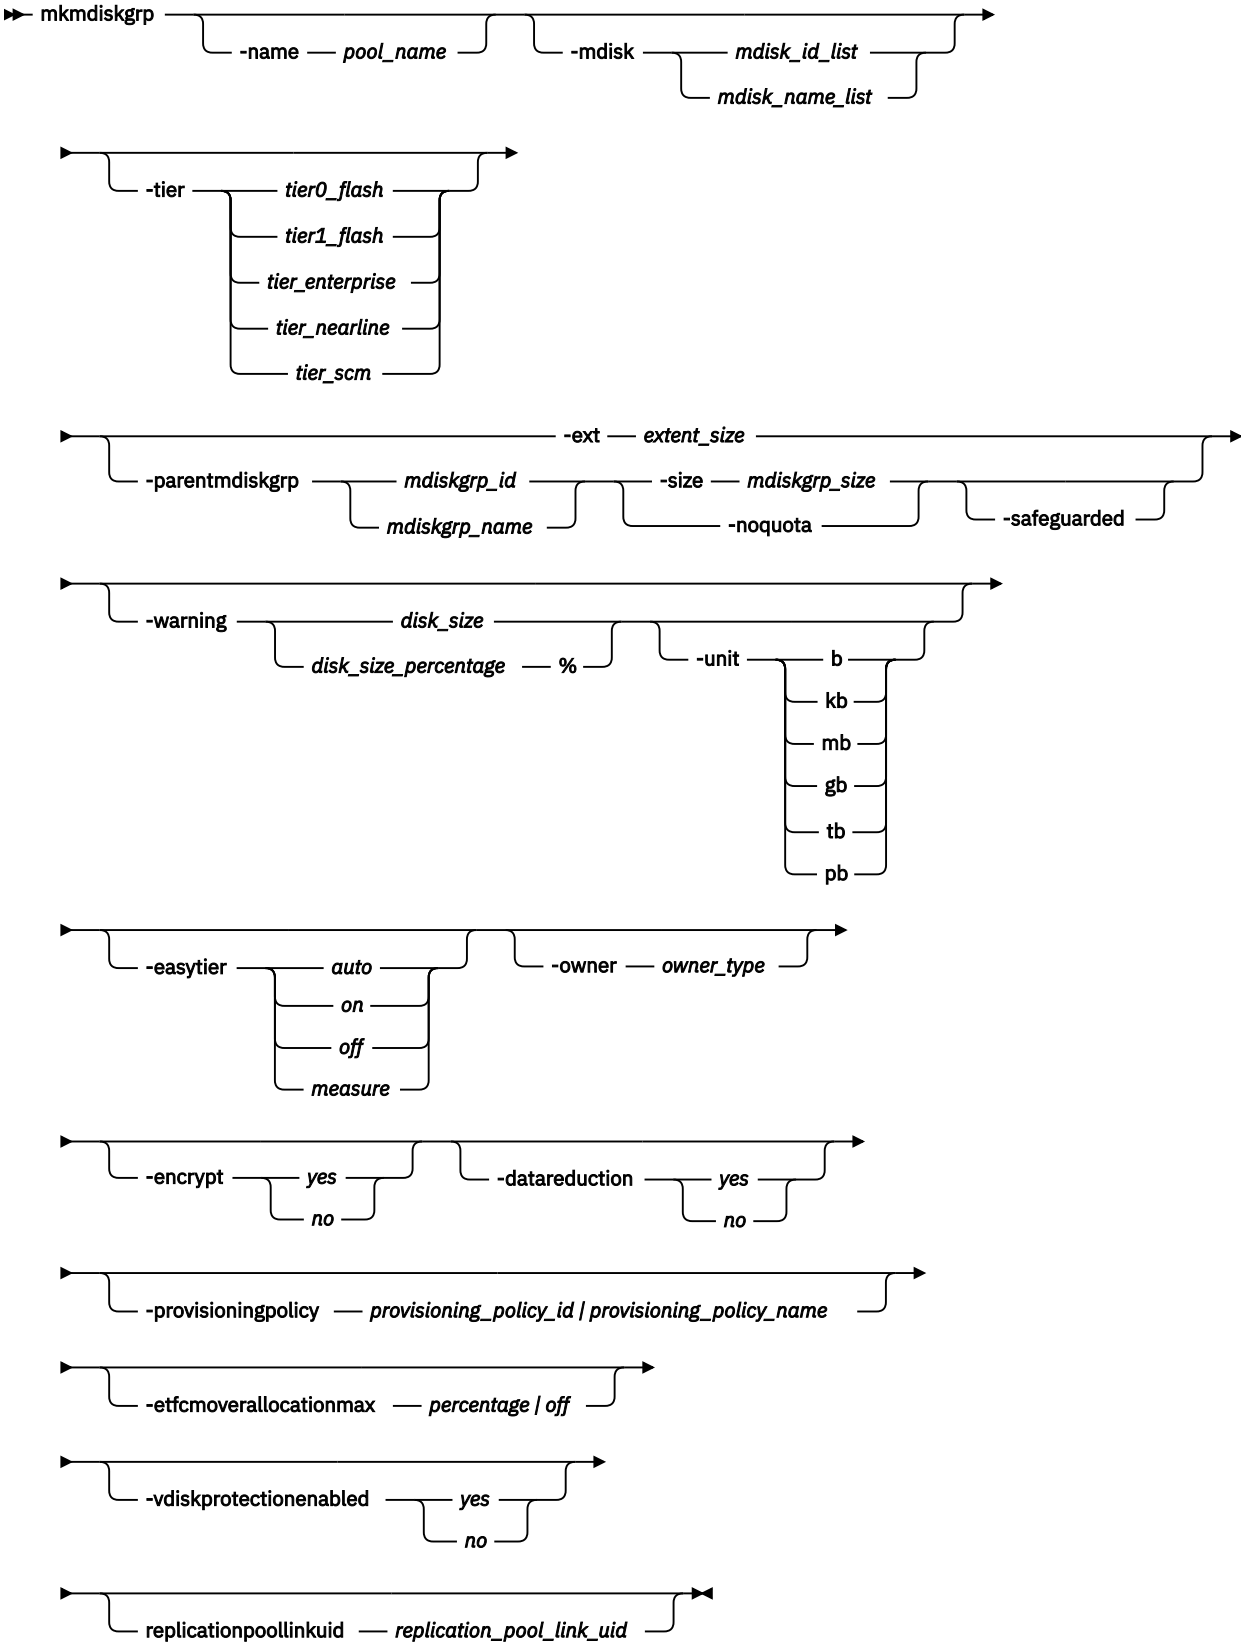
replication_pool_link_uid
Display the replication pool link unique identifier (UID) assigned to the pool. This value is used to link pools between systems for policy-based replication.

format_required
Indicate that freed extents belonging to the storage pool require formatting before they can be allocated to volumes. The formatting process is automatic.

mkmdiskgrp

Use the **mkmdiskgrp** command to create a new storage pool.

Syntax


mkmdiskgrp
-name
pool_name
-mdisk
mdisk_id_list
mdisk_name_list
-tier
tier0_flash
tier1_flash
tier_enterprise
tier_nearline
tier_scm
-ext
extent_size
-parentmdiskgrp
mdiskgrp_id
mdiskgrp_name
-size
mdiskgrp_size
-noquota
-safeguarded
-warning
disk_size
disk_size_percentage
%
-unit
b
kb
mb
gb
tb
pb
-easytier
auto
on
off
measure
-owner
owner_type
-encrypt
yes
no
-datareduction
yes
no
-provisioningpolicy
provisioning_policy_id | provisioning_policy_name
-etfcmoverallocationmax
percentage | off
-vdiskprotectionenabled
yes
no
replicationpoollinkuid
replication_pool_link_uid

Parameters

-name *pool_name*
(Optional) Specifies a name to assign to the new pool.

-mdisk *mdisk_id_list | mdisk_name_list*
(Optional) Specifies a colon-separated list of managed disk IDs or names to add to the storage pool. You can create an empty storage pool by not specifying the `-mdisk` parameter.

-tier *tier0_flash | tier1_flash | tier_enterprise | tier_nearline | tier_scm*
(Optional) Specifies the tier of the MDisk or MDisks being added.

tier0_flash
Specifies a `tier0_flash` hard disk drive or an external MDisk for the newly discovered or external volume.

tier1_flash
Specifies an `tier1_flash` (or flash drive) hard disk drive or an external MDisk for the newly discovered or external volume.

tier_enterprise
Specifies a `tier_enterprise` hard disk drive or an external MDisk for the newly discovered or external volume.

tier_nearline
Specifies a `tier_nearline` hard disk drive or an external MDisk for the newly discovered or external volume.

tier_scm
Specifies a `tier_scm` hard disk drive or an external MDisk for the newly discovered or external volume.

If you do not specify a tier, the current tier value of the MDisk is retained. The default value for an external MDisk is *enterprise*.

Note: Since Release 8.1.0, the tier of MDisks mapped from certain types of IBM System Storage Enterprise Flash is fixed to `tier0_flash`, and cannot be changed.

-ext *extent_size*
(Required) Specifies the size of the extents for this group in MiB.

-size *mdiskgrp_size*
(Optional) Specifies the child pool capacity. The value must be a numeric value (and an integer multiple of the extent size).

-parentmdiskgrp *mdiskgrp_id | mdiskgrp_name*
(Optional) Specifies the parent pool from which the volume extents of the child pool are allocated when you create a child pool. The value must be an *mdiskgrp_id* or *mdiskgrp_name*.

-warning *disk_size | disk_size_percentage%*
(Optional) Generates a warning when the used disk capacity in the storage pool first exceeds the specified threshold. You can specify a *disk_size* integer, which defaults to megabytes (MiB) unless the `-unit` parameter is specified; or you can specify a *disk_size*%, which is a percentage of the storage pool size. To disable warnings, specify `0` or `0%`. The default value is `0`.

-unit b | kb | mb | gb | tb | pb
(Optional) Specifies the data units for the `-warning` parameter.

-easytier *on | off | auto | measure*
(Optional) Specifies whether the Easy Tier function is active for this storage pool, or if it is automatically determined. *auto* is the default value. **-easytier** is `active` in storage pools with multiple tiers and is `balance` with single tiers.

Note:

- If **-easytier** is set to *auto*, the system automatically enables Easy Tier functions when the storage pool contains MDisk from more than one tier, and enables automatic rebalancing when the storage pool contains an MDisk from only one tier.

- If **-easytier** is set to *on*, then Easy Tier functions are active.
- If **-easytier** is set to *off*, then Easy Tier functions are inactive.
- If **-easytier** is set to *measure* Easy Tier statistics are collected but Easy Tier management is disabled. (No extents are moved by Easy Tier).

auto equates to:

- *on* if Easy Tier is licensed or no license is required.
- *off* if Easy Tier is not licensed and a license is required.

Specifying `-easytier on` enables Easy Tier:

- Management of both single-tier and multitier pools
- Auto rebalance

Extents are moved to balance the I/O load on the MDisks in the pool.

-owner ***owner_type***
(Optional) Specifies the owner type required for vVol child pools. The value must be vvol_child_pool for the system to use as a vVol storage container.

-encrypt ***yes | no***
(Optional) Specifies the encryption status for this storage pool. The values are `yes` or `no`.

Remember:

- If you do not specify **-encrypt** and encryption is enabled, the system defaults to `-encrypt yes` (the default encryption setting).
- When you create a child pool in an encrypted parent pool, the value must not be `no`. (All other permutations are permitted).
- When you create a data reduction child pool, you cannot specify `-encrypt`. The encryption status is inherited from the parent pool.

-datareduction ***yes | no***
(Optional) Specifies whether the storage pool is a data reduction storage pool. The values are `yes` or `no`. A value of `no` specifies that the storage pool is a standard storage pool.

-provisioningpolicy ***provisioning_policy_id | provisioning_policy_name***
(Optional) The name or ID of the provisioning policy for the pool.

-etfcmoverallocationmax ***percentage | off***
(Optional) Sets the maximum over allocation which Easy Tier can migrate onto FlashCore Module arrays, when the array is used as the top tier in a multitier pool. The value acts as a multiplier of the physically available space. The allowed values are a percentage in the range of 100% (default) to 400% or **off**. Setting the value to **off** disables this feature.

-vdiskprotectionenabled ***yes | no***
(Optional) Specifies whether volume protection is enabled for this storage pool. The default value is **yes**, except for child pools where the default value is inherited from the parent pool.

Note: Volume protection can not be enabled for pools with the owner **vvol_child_pool**. Thus, the default value is *no* while creating a pool with this owner and the **mkmdiskgrp** command fails if the argument **-vdiskprotectionenabled yes** is passed and the pool has the owner **vvol_child_pool**.

-noquota

(Optional) Creates a data reduction child pool without a size. The child pool must be within a parent pool that is also a data reduction storage pool.

-replicationpoollinkuid ***replication_pool_link_uid***
(Optional) Specifies the replication pool link unique identifier (UID) value to assign to the new pool. This value is used for linking pools between local and remote systems that use policy-based replication.

Description

The following items apply when you use this command:

- The **-noquota** and **-size** parameters are mutually exclusive.
- The **-noquota** parameter must be used with **-datareduction** set to *yes* to create a data reduction child pool.
- The **-noquota** parameter must be used with **-parentmdiskgrp** in a parent data reduction storage pool.

Table 148. Parameter differences for child pools and storage pools

Parameter	Child pool usage	Storage pool usage
-name	Optional	Optional for both parent pools and child pools.
-mdisk	Cannot be used with child pools.	Optional
-tier	Cannot be used with child pools.	Optional
-easytier	Cannot be used with child pools.	Optional
-size	Mandatory for traditional child pools. Cannot be used with child pools that are data reduction pools.	Cannot be used with parent pools.
-parentmdiskgrp	Mandatory	Cannot be used with parent pools.
-ext	Cannot be used for child pools.	Mandatory
-unit	Optional	Optional for both parent pools and child pools.
-warning	Optional for traditional child pools. Cannot be used with child pools that are data reduction pools.	Optional for both parent pools and child pools.
-encrypt	Optional. Cannot be used with child pools that are data reduction storage pools.	Optional for both parent pools and child pools.
-datareduction	Optional. Can be used to create a traditional child pool when the value is set to *no*. Can be used to create a data reduction child pool when the value is set to yes, and **-parentmdiskgrp** is a parent data reduction pool.	Data reduction pools must be a parent pool.
-noquota	Optional. Cannot be used with **-size**, **-warning**, or **-encrypt**.	Cannot be used with parent pools.

The **mkmdiskgrp** command creates a new storage pool and assigns the storage pool name if specified. The ID of the new storage pool is returned if the command is successful. Storage pools are collections of managed disks. Each storage pool is divided into chunks, called extents, which are used to create volumes.

Optionally, you can specify a list of managed disks that are added to this storage pool. These managed disks cannot belong to another storage pool, and they must have a mode of unmanaged. Use the

lsmdiskcandidate command to get a list of suitable candidates. If **-tier** is specified, it applies to all of the MDisks.

Each managed disk that is a member of this group is split into extents. The storage that is available on these disks is added to a pool of extents that is available in this group. When a volume is created from this group, free extents from the pool are used, in accordance with the policy used when the volume was first created.

Subsequently, all managed disks added to this group are split into extents of the same size as the size that is assigned to the group.

When you choose an extent size, be aware of the amount of storage you want to virtualize in this group. The system maintains a mapping of extents between volumes and managed disks. One system can virtualize the following number of extents:

- 64 TB - if all storage pools have extent sizes of 16 MiB.
- 2 PB - if all storage pools have extent sizes of 512 MiB.
- 32 PB - if all storage pools have extent sizes of 8192 MiB.

Important: The extent size for the storage pool can also limit volume size. Consider the maximum volume size that you want to use when you create storage pools. Refer to the information on creating storage pools for a comparison of the maximum volume capacity for each extent size. The maximum is different for thin-provisioned volumes.

Note: When an image mode volume is created, the storage pool increases in capacity by the size of the image mode volume (not the MDisk capacity) because the image mode volume might be smaller than the MDisk itself. If an extent is migrated from the image mode volume or MDisk to elsewhere in the group, the volume becomes a striped volume (no longer image mode). At this point, the available capacity might increase because the extra capacity available on the MDisk (for example, the capacity that was not part of the image mode volume) becomes available.

Remember: When you use the **-mdisk** parameter, the storage pool is not created if one of the MDisks is a SAS MDisk.

When you specify `-name` *pool_name* if you do not also specify **-parentmdiskgrp**, you create a parent pool where *pool_name* is the name of the new storage pool. When you specify `-name` *pool_name*, if you also specify **-parentmdiskgrp** and a size for it or **-noquota**, you create a child pool where *pool_name* is the name of the new storage pool.

Note:

A data reduction pool is created by using the **-datareduction** parameter set to `yes`. The pool can be used to create standard-provisioned, thin or compressed volumes, or volume copies.

There is a maximum number of four data reduction pools in a system. When this limit is reached, creating any further pools with **-datareduction** set to `yes` is not possible.

The **-datareduction** parameter can be used to create a child pool that is a data reduction pool. The parent pool must also be a data reduction pool.

In most cases volume protection is applied on a volume level, but a volume might have more than one copy, with copies in different storage pools, which in turn can have different protection settings. A volume is protected when it has one or more copies in a storage pool that has `vdisk_protection_status active`.

For **rmvdiskcopy**, volume protection is applied on a volume copy level, but only when you delete the last volume copy. When you delete one of two volume copies, volume protection is not applied. It means that a volume copy might be deleted from a protected storage pool and leaves a remaining copy in an unprotected storage pool that might then later be deleted without volume protection checks.

An invocation example

This example adds a list of MDisks to the storage pool.

```
mkmdiskgrp -mdisk mdisk0:mdisk1:mdisk2:mdisk3 -ext 4096
```

The resulting output:

```
MDisk Group, id [0], successfully created
```

rmmdisk

Use the **rmmdisk** command to delete a managed disk (MDisk) from a storage pool.

Syntax



Parameters

-mdisk *mdisk_id_list | mdisk_name_list*
(Required) Specifies one or more managed disk IDs or names to delete from the group.

-force
(Optional) Migrates data on the specified disks to other disks in the group. The command completes asynchronously if **-force** is specified.

mdisk_group_id | mdisk_group_name
(Required) Specifies the ID or name of the storage pool to delete the disks from. The warning threshold for a storage pool is automatically scaled when MDisks are deleted.

Note: The command cannot be used with image-mode volumes.

Description

This command attempts to remove the managed disk or disks from the group.

Remember: This command cannot be used for child pools.

Deleting a managed disk from a group can be done only if the managed disk does not contain any extents in use by a volume. If there are extents in use and you do not supply the force flag, the command fails.

Attention: If this disk being removed has already been powered down, removed, or is experiencing a power outage, the migration is pending and does not complete until the MDisk comes back online. The MDisk is not removed from the list of 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 delete them from the storage pool that they belong to.

The **rmmdisk** command fails if there are insufficient free extents on other disks in the storage pool for the duration of the command.

If you do specify the force flag, an attempt is made to migrate the extents that are in use onto other free extents within the storage pool. If there are not enough free extents in the storage pool, the command will fail even if the force flag is specified.

When an array MDisk is in a storage pool, five extents in the storage pool are reserved for internal use. If you attempt to remove an MDisk when an array MDisk is in the storage pool, the command fails (even if the **-force** flag is specified), if five free extents do not remain in the storage pool.

To delete the disks from the group, you have the following options:

- You can delete the volume that is using the extents that are 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 returns with a success code, notifying you that migration is in progress (if migration is required). An event is logged when the migration is complete and the disk is deleted from the group at this time. You can also check the progress of any active migrations by running the **lsmigrate** command.

If you specify **-force**, the **rmmdisk** command fails if there are offline MDisks. If there are no online quorum disks the migration fails.

Remember: When using the **-mdisk** parameter, MDisks are removed if there is one (or more) SAS MDisk specified in the list.

An invocation example

```
rmmdisk -mdisk mdisk12 -force Group3
```

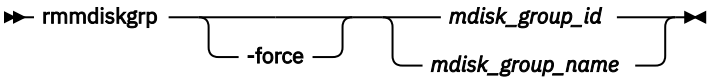
The resulting output:

```
No feedback
```

rmmdiskgrp

Use the **rmmdiskgrp** command to delete a storage pool without being able to recover it.

Syntax



Parameters

-force
(Optional) Specifies that all volumes and host mappings be deleted. When you use this parameter, all managed disks in the storage pool are removed and the storage pool itself is deleted.

Remember:

- You must specify **-force** to delete a standard child pool if it contains any volumes. The command fails if **-force** is used to delete a data reduction child pool if it contains any volumes.
- You cannot specify **-force** to delete a parent pool if it has child pools.

Note: The command fails if **-force** is used to delete an MDisk group if:

- Any of the VDisks in the MDisk group are mirrored across multiple MDisk groups (other than the one that is being deleted).
- AND any of the VDisk mirrors are out of sync.
- AND an attempt is made to delete the in-sync copy. Deleting the only in-sync copy requires **-force**. Otherwise, it isn't needed if the VDisk has another in-sync copy.
- AND the out-of-sync copy is a thin-provisioned or compressed copy in a data reduction pool.

Note: When you use the **-force** parameter to delete a storage pool that has volume protection inactive but includes mirrored volumes with out-of-sync copies in a storage pool with volume protection active, then volume protection applies to these volumes. The out-of-sync copies are deleted from the storage pool with volume protection active as a part of the forced delete. The command fails if host I/O is being sent to these mirrored volumes in the defined volume protection time period.

mdisk_group_id | mdisk_group_name
(Required) Specifies the ID or name of the storage pool that is to be deleted.

Note: You cannot delete a parent pool that has child pools. You must first delete the child pools.

Description

Important: Before you issue the command, ensure that you want to delete all mapping information. Data that is contained on the volume cannot be recovered after the storage pool is deleted.

The `rmmdiskgrp` command deletes the specified storage pool. The **-force** parameter is required if there are volumes that have been created from this storage pool or if there are managed disks in the storage pool. Otherwise, the command fails.

Note: This command also removes any associated storage pool throttling.

Deleting a storage pool is essentially the same as deleting a system or part of a system because the storage pool is the central point of control of virtualization. Because volumes are created by using available extents in the storage pool, mapping between volume extents and managed disk extents is controlled based on the storage pool.

The command deletes all volume copies in the specified storage pool. If the volume has no remaining synchronized copies in other storage pools, the volume is also deleted.

This command deletes the associated MDisk group (storage pool) throttle if that storage pool is removed.

Remember: This command is unsuccessful if:

- Volume protection is enabled (by using the **chsystem** command).
- The MDisk being removed is mapped to any volume that received I/O within the defined volume protection time period.

Remember: This command partially completes asynchronously. All volumes, host mappings, and Copy Services relationships are deleted before the command completes. The deletion of the storage pool then completes asynchronously.

In detail, if you specify the **-force** parameter and the volumes are still using extents in this storage pool, the following actions are initiated or occur:

- The mappings between that disk and any host objects and the associated Copy Services relationships are deleted.
- If the volume is a part of a FlashCopy mapping, the mapping is deleted.

  Note: If the mapping is not in the **idle_or_copied** or stopped states, the mapping is force-stopped and then deleted. Force-stopping the mapping might cause other FlashCopy mappings in the system to also be stopped. For more information, see the description for the **-force** parameter in the **stopfcmap** command.
- Any volume that is in the process of being migrated into or out of the storage pool is deleted. It frees up any extents that the volume was using in another storage pool.
- Volumes are deleted without first flushing the cache. Therefore, the storage controller LUNs that underlie any image mode MDisks might not contain the same data as the image mode volume before the deletion.
- If managed disks exist in the storage pool, all disks are deleted from the storage pool. They are returned to the unmanaged state.
- The storage pool is deleted.

Attention: If you use the **-force** parameter to delete all the storage pools in your system, you are returned to the processing state where you were after you added nodes to the system. All data that is contained on the volumes is lost and cannot be recovered.

Note: The **rmmdiskgrp** command fails if the pool contains volumes that are used in policy-based replication. To remove the pool, you must remove the replication policy from any volume groups with volumes in that pool. The **rmmdiskgrp** command succeeds if mirrored volume copies are synchronized and available. Synchronized copies ensures that if the pool is deleted, volumes do not get deleted from the volume group.

An invocation example

```
rmmdiskgrp -force Group3
```

The resulting output:

```
No feedback
```

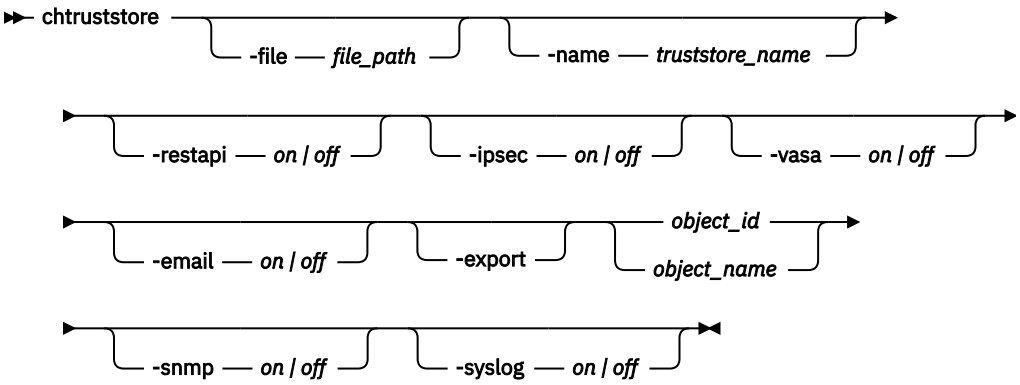
Chapter 42. Truststore management commands

Use the **truststore** management commands to configure repositories of certificates and tag stores for use with network services.

chtruststore

Use the **chtruststore** command to manage the certificate store and append additional certificates.

Syntax



Parameters

-file *file_path*
(Optional) Specifies the location of the file to import the certificates to append those already in the store.

Note: You cannot append the certificates on the truststore with object ID 0.

-name *truststore_name*
(Optional) Specifies a new name for the certificate store.

Note: You cannot change the name of the truststore with object ID 0.

-restapi *on | off*
(Optional) Specifies the certificates in the store are for use by the REST API. Values can be *on* or *off*.

Note: You cannot enable this parameter on the truststore with object ID 0.

-ipsec *on | off*
(Optional) Specifies the certificates in the store are for use by the IPsec service. Values can be *on* or *off*.

Note: You cannot enable this parameter on the truststore with object ID 0.

-vasa *on | off*
(Optional) Specifies the certificates in the store are for use by VASA Provider. Values can be *on* or *off*.

Note: You cannot enable this parameter on the truststore with object ID 0.

-email *on | off*
(Optional) Specifies the certificates in the store are used to validate the email server. Values can be *on* or *off*.

Note: You can enable this parameter on the truststore with object ID 0.

© Copyright IBM Corp. 2003, 2025

-snmp ***on | off***
(Optional) Specifies the certificates in the store are used to validate the SNMP servers. Value can be *on* or *off*. Default is set to *off*.

-export
(Optional) Exports the specified truststore to `/dumps/truststore.pem` on the configuration node.

object_id | object_name
(Required) Specifies the ID or name of the store to manage.

-syslog ***on | off***
(Optional) Specifies the certificates to be bundled and provided to rsyslog client for making TLS connections. Values can be *on* or *off*.

An invocation example

```
chtruststore -file /tmp/certificate-bundle.pem 1
```

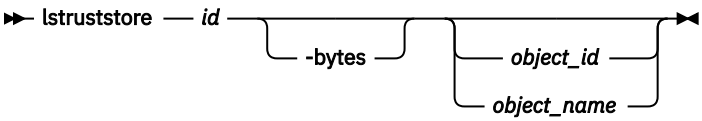
The resulting output:

```
No feedback.
```

lstruststore

Use the **lstruststore** command to display the current certificate stores.

Syntax



Parameters

-bytes
(Optional) Displays all capacity values in bytes.

object_id | object_name
(Optional) Specifies the object ID or name that you want to display.

Description

This command displays all trust stores that are available to the system.

This table provides the attribute values that can be displayed as output view data.

Table 149. lstruststore output

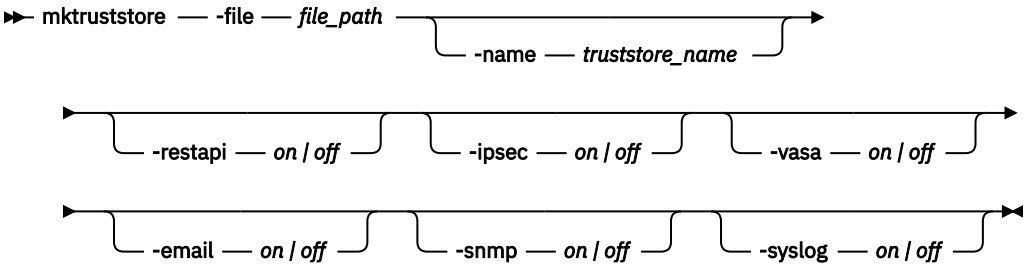
Attribute	Description
id	Indicates the store ID.
name	Indicates the name of the store.
percent_used	Indicates the used space in the store as a percentage of available space.
space	Indicates the remaining space available. Default is KiB up to 1 decimal place.
ipsec	Indicates if the certificates in the store are used for the IPsec service. Value can be *on* or *off*.

Table 149. lstruststore output (continued)	
Attribute	**Description**
vasa	Indicates if the certificates in the store are used for the VASA Provider. Value can be *on* or *off*.
restapi	Indicates if the certificates in the store are used for the REST API. Value can be *on* or *off*.
email	Indicates if the certificates in the store are used to validate the configured email servers. Value can be *on* or *off*.
snmp	Indicates if the certificates in the store are used to validate the configured SNMP servers.
syslog	Indicates if the certificates in the store are used to validate the configured syslog servers.
flash_grid_address	Indicates the IP address or FQDN of the FlashSystem grid member that this entry relates to.

mktruststore

Use the **mktruststore** command to create the certificate store that contains the certificate bundle.

Syntax



Parameters

-file ***file_path***
(Required) Specifies the location of the file to import the certificates.

-name ***truststore_name***
(Optional) Specifies the name for the new certificate store.

-restapi ***on | off***
(Optional) Specifies the certificates in the store are used for the REST API. Values can be *on* or *off*.

-ipsec ***on | off***
(Optional) Specifies the certificates in the store are used for the IPsec service. Values can be *on* or *off*.

-vasa ***on | off***
(Optional) Specifies the certificates in the store are used for the VASA Provider. Values can be *on* or *off*.

-email ***on | off***
(Optional) Specifies the certificates in the store are used to validate the email server. Values can be *on* or *off*.

-snmp ***on | off***
(Optional) Specifies the certificates in the store are used to validate the SNMP servers. Values can be *on* or *off*. Default is set to *off*.

-syslog *on | off*
(Optional) Specifies the certificates to be bundled and provided to rsyslog client for making TLS connections. Values can be *on* or *off*..

An invocation example

```
mktruststore -file /tmp/certificate-bundle.pem
```

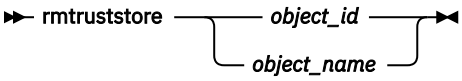
The resulting output:

```
No feedback.
```

rmtruststore

Use the **rmtruststore** command to remove a certificate store from the system.

Syntax



Parameters

object_id |object_name
(Required) Specifies the object ID or name of the store to remove.

Note: You cannot remove the truststore with object ID 0.

An invocation example

```
rmtruststore 1
```

The resulting output:

```
No feedback.
```

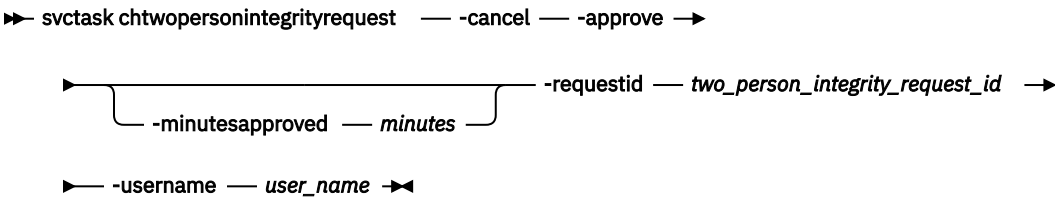
Chapter 43. Two person integrity commands

Use two person integrity (TPI) commands to create, view, approve, cancel, and update role elevation requests. To use these commands, you must have the Restricted Security Administrator role or the Security Administrator role. For example, if you have Security Administrator role you already hold a valid approved TPI request, and you can approve or cancel your own or other TPI request.

chtwopersonintegrityrequest

Use the **chtwopersonintegrityrequest** command to approve or cancel an existing role elevation request. You must have the Restricted Security Administrator role or the Security Administrator role to use this command.

Syntax



Parameters

-cancel
(Required) Specifies to cancel and remove the role elevation request.

-approve
(Required) Specifies to approve the role elevation request.

-minutesapproved <minutes>
(Optional) Specifies the duration (in minutes) for an approved role elevation request.

-requestid <request_id>
(Required) Specifies the ID of the role elevation request.

-username <user_name>
(Required) Specifies the username that is associated with the role elevation request.

Description

You must have the Restricted Security Administrator or the Security Administrator role to cancel or approve a role elevation request. The request can be approved or canceled by either specifying the ID or the user name that is associated with the request. The minimum duration for an elevation request is 10 minutes and the maximum duration is 24 hours. If **-minutesapproved** is not provided when doing an approve (-approve) operation, the duration that was originally specified when the request was created is used.

- The **-cancel** and **-approve** parameters are mutually exclusive, but one of them is required.
- The **-requestid** and **-username** parameters are mutually exclusive, but one of them is required.
- The **-cancel** and **-minutesapproved** parameters are mutually exclusive.
- When the **-approve** parameter is specified, you can also specify the **-minutesapproved** parameter.

© Copyright IBM Corp. 2003, 2025

An invocation example

The following example shows the approval of a role elevation request by username with a duration of 30 minutes, regardless of how many minutes were originally requested:

```
chtwopersonintegrityrequest -approve -minutesapproved 30 -username secadmin11
```

The resulting output:

```
No feedback.
```

lstwopersonintegrityrequest

Use the **lstwopersonintegrityrequest** command to list the role elevation requests active on the system.

Syntax



Parameters

request_id | user_name
(Required) Specifies the ID or username that is associated with the request.

Description

The following attribute values can be displayed in the output:

*Table 150. **lstwopersonintegrityrequest** output*

Attribute	Description
request_id	Indicates the unique ID of the role elevation request.
user_name	Indicates the user name that is associated with the role elevation request.
is_remote	Indicates whether the user named in the role elevation request is a remote user. The value can be yes or no.
minutes_requested	Indicates the number of minutes requested by the user.
minutes_approved	Indicates the number of minutes approved by the user.
request_time	Indicates the time (YYMMDDhhmmss) of request.
approval_time	Indicates the time (YYMMDDhhmmss) of approval.
expiration_time	Indicates the date and time (YYMMDDhhmmss) when the temporary role elevation expires.

Note: The **expiration_time** field value is adjusted as the role elevation request progresses.

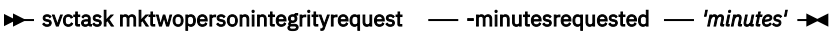
- It is calculated first at the time that the role elevation request happens (for example **mktwopersointegrityrequest**) based on the time of the request (*request_time* field value) plus 24 hours.
- When the role elevation request is approved (for example **chtwopersointegrityrequest -approve**), the expiration time is recalculated based on the time the request is approved (**approval_time** field value) plus the approved duration of the request (**minutes_approved** field value).

- Additionally, if the system clock shifts (for example, system time change) five minutes or more, the expiration times of all role elevation requests are updated to reflect that difference. That is the case after a system power-off, any role elevation requests present at that time, when the system restarts will have their expiration times adjusted resulting in those expiration times extending approximately by the length of the power-off.

mktwopersonintegrityrequest

Use the **mktwopersonintegrityrequest** command to request a role elevation for yourself. This command requests on behalf of the current user invoking this command.

Syntax



Parameters

-minutesrequested 'minutes'
(Required) Specifies the duration (in minutes) of the role elevation.

Description

To run the **mktwopersonintegrityrequest** command, the following conditions must be met:

- TPI must be enabled.
- You must not have any active role elevation requests.
- You must not be a superuser.

An invocation example

The following example shows a role elevation request:

```
mktwopersonintegrityrequest -minutesrequested 45
```

The resulting output:

```
Two Person Integrity request, id[3], successfully created
```

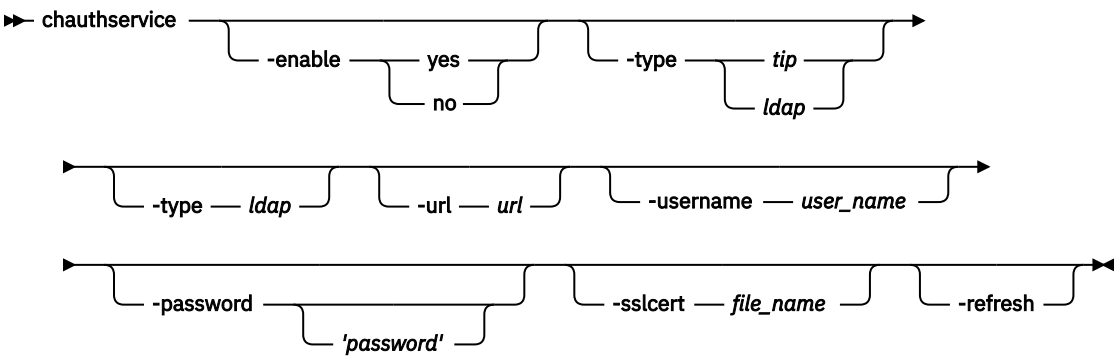
Chapter 44. User management and access control commands

Use the user management and access control commands to configure remote authentication service, access control, and manage users and groups on the system.

chauthservice

Use the **chauthservice** command to configure the remote authentication service of the system.

Syntax



Parameters

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

-type ldap
(Optional) Specifies the authentication service type (Tivoli Integrated Portal or native LDAP). An LDAP server must be configured. Before you change **-type**, ensure that the remote authentication type that is selected is properly configured.(Optional) Specifies the authentication service type (which must be LDAP). An LDAP server must be configured.

Remember:

- The remote authentication service must be enabled (**-enable yes**) for this setting to come into effect.
- Before you change **-type** from *ldap* to *tip*, ensure that all users who are configured for remote authentication have both an SSH key and password configured.

Remember: The remote authentication service must be enabled (**-enable yes**) for this setting to come into effect.

-url *url*
(Optional - IBM Security Services only) Specifies the website address (URL) of Security Services, which is referred to as TIP in the CLI. The host part of the URL must be a valid numeric IPv4 or IPv6 network address. You can use the following characters in the URL:

- a - z
- A - Z
- 0 - 9

© Copyright IBM Corp. 2003, 2025

- _
- ~
- :
- [
-]
- %
- /

The maximum length of the URL is 100 characters.

This option is no longer used.

-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 except for 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. You must enclose the password in single quotation marks. 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. Do not enclose the password in single quotation marks if you use the prompt.

-sslcert *file_name*
(Optional) Specifies the name of the file that contains the SSL certificate, in privacy enhanced mail (PEM) format, for the remote authentication server. The certificate file must be in valid PEM format and have a maximum length of 12 KB.

-refresh
(Optional) Causes the system to invalidate any remote user authorizations that are cached on the system. Use this option when you modify user groups on the authentication service and want the change to immediately take effect on the system.

Note: If you clear the cache, anyone who uses the system might have to log in again (for example, if credentials are provided to one of the defined LDAP servers).

Description

The system authenticates remote users by using Lightweight Directory Access Protocol (LDAP).

Before you enable remote authentication, ensure that the properties of the service are properly configured on the system. It is not necessary to disable the remote authentication service to change its properties. LDAP authentication can be configured by using the **chldap** command, and LDAP servers can be added to the system by using the **mkldapserver** command.

Remember: For the authentication type to be set to LDAP with authorization enabled (true), an LDAP server must be configured.

When the authentication service is enabled, the system does not test whether the remote authentication system is operating correctly.

- To establish whether the system is operating correctly, enter the **lscurrentuser** command for a remotely authenticated user. If the output lists the user roles that are obtained from the remote

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

- To establish whether LDAP is operating correctly, in addition to the **lscurrentuser** command, enter the **testldapserver** command. The **testldapserver** command can be entered whether or not remote authentication is enabled, and can be used to test the connection to LDAP servers, as well as user authorization and authentication.

To disable the remote authentication service in a controlled manner when it is not available, use the **enable** parameter with the no option.

Note: On a system where two person integrity (TPI) is enabled, you can't disable the remote authentication service if the minimum number of users that are required by TPI also includes remote users.

This command can be used to select and enable a remote authentication service for use with the system.

An invocation example

To fully configure and enable authentication with IBM Security Services:

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

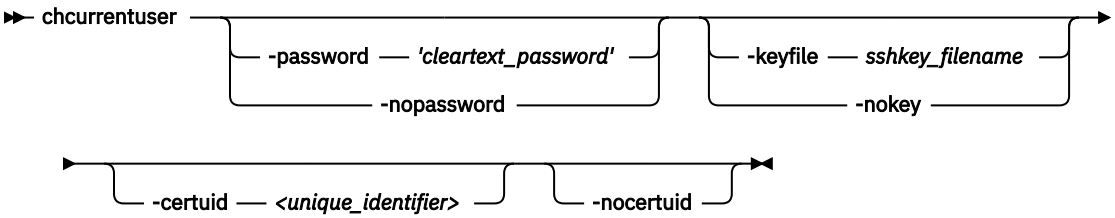
The following text is displayed when the command runs:

```
No feedback
```

chcurrentuser

Use the **chcurrentuser** command to change the attributes of the current user.

Syntax



Parameters

-password ***'cleartext_password'***
(Optional) Specifies the new password to be associated with the current user. The password cannot start or end with a blank. It must consist of a string of 6 - 64 printable ASCII characters. You must enclose the password in single quotation marks. 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. Do not enclose the password in single quotation marks if you use the prompt. 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.

-certuid *<unique_identifier>*
(Optional) Specifies a Unique ID that can be used by the authentication services that might run on the system. It must consist of a string up to *255* printable ASCII characters.

-nocertuid
(Optional) Deletes the **certuid** associated with the user.

Note: The **-certuid** parameter is mutually exclusive with **-keyfile** and **-nocertuid**.

Description

Use the **chcurrentuser** command to modify the attributes of the current user.

An invocation example

```
chcurrentuser -password 'secret' -nokey
```

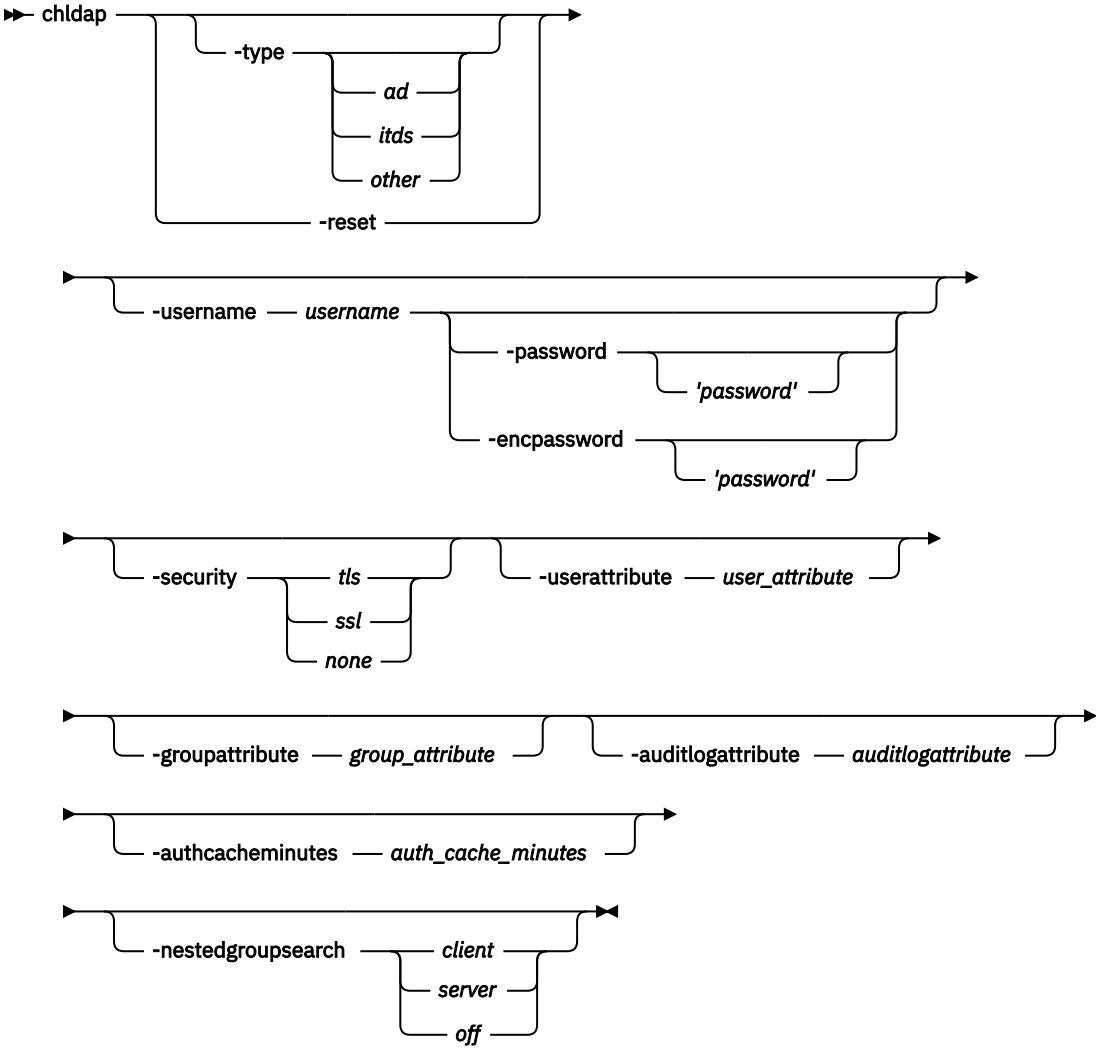
The resulting output:

```
No feedback
```

chldap

Use the **chldap** command to change system-wide Lightweight Directory Access Protocol (LDAP) configuration. This command can be used to configure remote authentication with LDAP. These settings apply when authenticating against any of the LDAP servers configured using the **mkldapserver** command.

Syntax



Parameters

-type ***ad*** **|** ***itds*** **|** ***other*** **| -reset**
(Optional) Specify the LDAP server type, or reset LDAP configuration to defaults for the current server type. Defaults for the configured server type:

- Active Directory (AD)
- IBM Security Directory Server (ISDS)
- Other

-username ***username***
(Optional) Specifies a username for administrative binding. This can be:

Note:

- A distinguished name (DN)
- A user principal name (UPN) or NT login name for Active Directory

-password ***password***
(Optional) Specifies the password for the administrative binding. You can optionally specify the password with this parameter. You must enclose the password in single quotation marks. 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. Do not enclose the password in single quotation marks if you use the prompt.

-encpassword *password*
(Optional) Specifies the password for the enclosure. You can optionally specify the password with this parameter. You must enclose the password in single quotation marks. 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. Do not enclose the password in single quotation marks if you use the prompt.

-security *tls | ssl | none*
(Optional) Specifies the type of security to use when communicating with LDAP servers. Specifying `tls` enables Transport Layer Security (TLS) security. Specifying `ssl` enables Secure Socket Layer (SSL) security. The default value is `none`.

-userattribute *user_attribute*
(Optional) Specifies the LDAP attribute used to determine the user name of remote users. The user attribute must exist in your LDAP schema and must be unique for each of your users.

-groupattribute *group_attribute*
(Optional) Specifies the LDAP attribute used to determine the group memberships of remote users. The attribute must contain either the DN of a group or a colon-separated list of group names.

-auditlogattribute *auditlogattribute*
(Optional) Specifies the LDAP attribute used to determine the identity of remote users. When a user performs an audited action, this information is recorded in the audit.

-authcacheminutes *auth_cache_minutes*
(Optional) Specifies the period for which to cache authentication details.

-nestedgroupsearch *client | server | off*
(Optional) Specifies whether nested groups are evaluated on the client (clustered system), server (authentication service), or are not evaluated not at all.

Description

At least one parameter must be specified.

The **chldap** command can be run whether or not LDAP authentication is enabled. Specifying **-reset** or **-type** populates the default values unless otherwise specified.

You can only specify **-password** or **-encpassword** if **-username** is specified.

The **-type** parameter values are only set to defaults for the specified type if the type is different from the existing type.

If the type is **itds**, **-nestedgroupsearch** cannot be executed (nested groups are evaluated by default). If the type is **ad**, **-nestedgroupsearch** can only be set to **client** or **off** because there is no server support. If the type is **other**, the **-nestedgroupsearch** parameter is fully configurable.

Use **-username** to specify a distinguished name (DN), user principal name (UPN), or NT login name. Distinguished names (DN) must be a sequence of attribute=value pairs separated by a comma (`,`), semi-colon(`;`), or plus sign (+). A backslash (`\,`) must be used to escape special characters, and can also be used to specify UTF-8 characters using their byte encoding. For example, c acute can be represented as `\C4\87`. NT logins are valid for only the Active Directory and must be in the `DOMAIN\user` format. These logins must not start or end with a period (`.`) and both the DOMAIN and the user must not use the following characters: `\/:?"<>|` UPN logins are valid for Active Directory only and must be in the format user@suffix. Both user and suffix can not use spaces or the following characters: `()<>,;:\"[]@`

Tip:

- Remember that **-userattribute**, **-groupattribute**, and **-auditlogattribute** accept values that:
  1. Must begin with a letter
  2. Only contain ASCII letters, digit characters, and hyphens
  3. Are case-insensitive

The following LDAP (first-time) configuration suggestions assist with LDAP server setup:

Important:

- Ensure that the system is configured appropriately according to your LDAP schema. Issue **chldap -type** to populate the system's LDAP configuration with the server type defaults. Issue **chldap -reset** to return to these defaults at any time.
  - (Advanced) For all server types, users are authenticated with a username configured in the LDAP attribute `user_attribute`. This attribute must exist in the LDAP schema and must be unique for each user. It is configurable by issuing **chldap -userattribute**. Active Directory users can also authenticate using their UPN or NT login names.
  - (Advanced) Authenticated users are assigned roles according to their LDAP group memberships. Each user's group memberships must be stored in the LDAP attribute `group_attribute`. This can be either an LDAP attribute containing the DN of the user's LDAP group, or an LDAP attribute containing a colon-separated list of user group names. It is configurable by issuing **chldap -groupattribute**.
  - (Advanced) When an LDAP authenticated user runs a command that is audited, the user's login name is placed in the audit log. The name is extracted from the LDAP attribute `audit_log_attribute`, which is configurable by issuing **chldap -auditlogattribute**.
- Ensure that the system is able to search within the user and group trees on LDAP servers. By default the system authenticates anonymously. Consequently, you must either permit anonymous searches of the LDAP directory, or create an LDAP user with the appropriate permissions and issue the **chldap -username** and **chldap -password** commands to instruct the system to search as this user.
- Ensure that the system is able to connect with the appropriate level of security. Passwords are sent to the LDAP server as clear text, so Transport Layer Security (TLS) encryption is recommended. Issue **chldap -security** to change the security level.
- (Advanced): On Active Directory and some other LDAP servers, the system (by default) identifies groups to which users belong directly. To assign users permissions according to a parent group, enable the nested group search on the client by issuing **chldap -nestedgroupsearch**. This setting has an additional performance overhead and supports up to 8 levels of nesting.

An invocation example

```
chldap -type
itds -username uid=joebloggs,cn=admins,dc=company,dc=com -password 'passw0rd'
-auditlogattribute descriptiveName
```

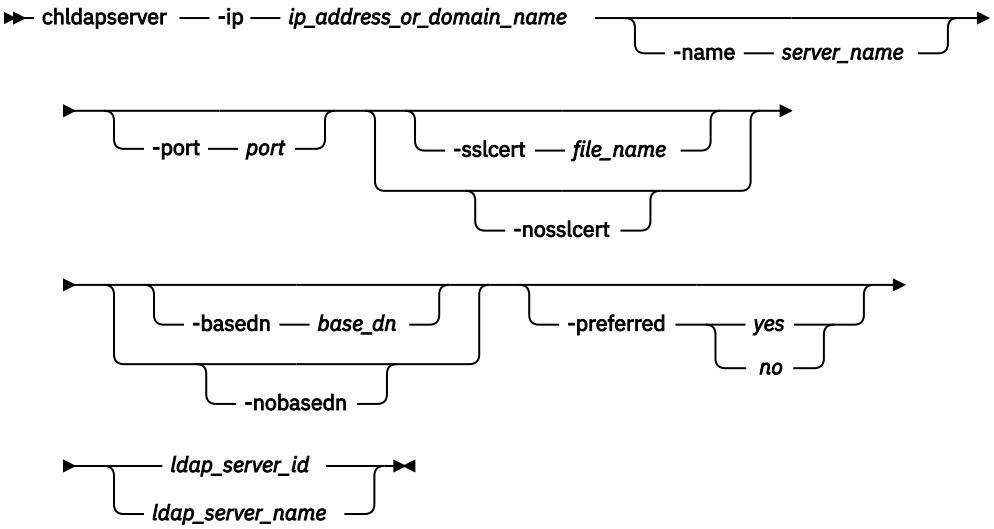
The resulting output:

```
No feedback
```

chldapserver

Use the **chldapserver** command to modify a Lightweight Directory Access Protocol (LDAP) server.

Syntax



Parameters

-ip *ip_address_or_domain_name*
(Required) Specifies either a server IP address or domain name, or a fully qualified domain name (FQDN) (maximum of 256 characters).

-name *server_name*
(Optional) Specifies the LDAP server name.

-port *port*
(Optional) Specifies the LDAP server port. The default value (if you do not specify a value) is 389. If you specify TLS security the value is 389 and if you specify SSL security the value is 636.

-sslcert *file_name* **| -nosslcert**
(Optional) Set (**-sslcert**) or clear (**-nosslcert**) the secure socket layer (SSL) certificate.

-basedn *base_dn* **| -nobasedn**
(Optional) Use the base distinguished name (DN) for search (**-nobasedn** indicates to use the default DN).

-preferred *yes | no*
(Optional) Specifies whether the server is preferred over other configured LDAP servers (or not preferred).

ldap_server_id* | *ldap_server_name
(Required) Specifies the LDAP server ID or name.

Description

Important: During normal operation, LDAP requests are sent to **-preferred** servers depending on availability. If no servers are marked as **-preferred**, LDAP requests are sent to configured servers based on availability.

If **-sslcert** is specified, the server certificate is verified while authenticating. The SSL certificate must exist on the current node. If **-nosslcert** is specified, any certificate file is deleted and the server certificate is not checked.

The **-basedn** parameter indicates the distinguished name (DN) to use as a base from which to search for users in the LDAP directory. If Transport Layer Security (TLS) is enabled and **-sslcert** is specified, the

server certificate is verified during authentication. The secure socket layer (SSL) certificate must exist on the node being used. Otherwise, a server certificate is not checked.

Note: The Microsoft AD server must use **-basedn**.

The system must be configured with an appropriate version IP address when **-ip** is specified. The IP address specified with the **-ip** parameter must be of a version supported by the system. The certificate file must be in valid PEM format and have a maximum length of 12 kilobytes.

Distinguished names must be a sequence of attribute=value pairs separated by a comma (,), semi-colon(;), or plus sign (+) escaping special characters with \ where appropriate, and specified UTF-8 characters using their byte encoding. For example, \, for commas or \C4\87 for the UTF-8 character c acute.

This command runs whether or not LDAP authentication is enabled.

Remember: There can be a maximum of six configured LDAP servers. If you attempt to create a seventh LDAP server an error is returned.

An invocation example with basic server details

```
chldapserver -ip 192.135.60.3 -port 400 ldapserver0
```

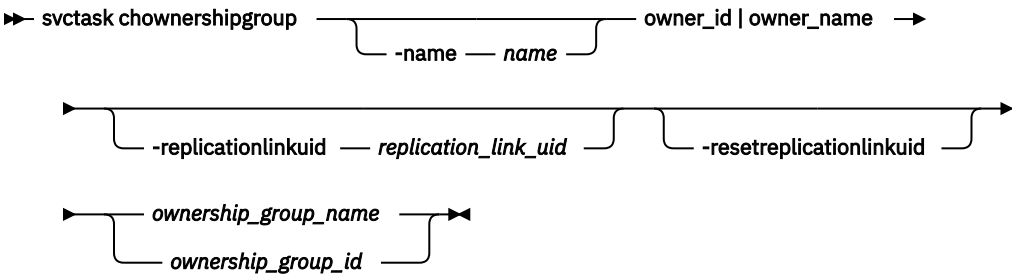
The resulting output:

```
No feedback
```

chownershipgroup

The **chownershipgroup** command modifies the properties of an ownership group.

Syntax



Parameters

-name *name*
(Optional) Specify the new name for the ownership group.

owner_id | owner_name
(Required) Specify the ID or the name of the ownership group to modify.

-replicationlinkuid *replication_link_uid*
(Optional) Specify the new link UID value for this ownership group.

-resetreplicationlinkuid
(Optional) Specifies that the ownership group link UID value for this ownership group should be reset. This action generates a new link UID value for the ownership group. By doing this, all connections between this ownership group and ownership groups on remote systems will be removed.

Note: Items owned by ownership groups (such as user groups, hosts, host clusters and port sets) that are unaffected by replication are not impacted by the ownership group link changes.

Description

Use the **chownershipgroup** command to modify the properties of an ownership group.

An invocation example

```
chownershipgroup -replicationlinkuid 000000000000000000000010024A0C04C ownershipgroup3
```

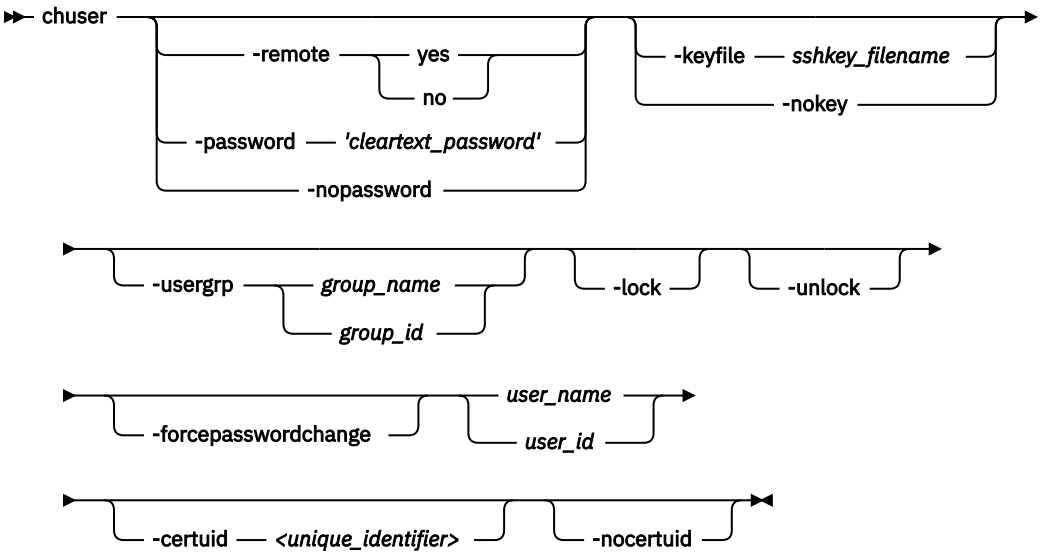
The resulting output:

```
chownershipgroup -resetreplicationlinkuid ownershipgroup3
```

chuser

Use the **chuser** command to change the attributes of an existing user.

Syntax



Parameters

-remote yes | no
(Optional) Specifies whether the user authenticates to the cluster by using a remote authentication service. Either yes or no must be set. If a **password** is already set, you cannot specify **-remote yes**.

-password *'cleartext_password'*
(Optional) Specifies the new password to be associated with the user. This parameter fails if remote is already set. The password cannot start or end with a blank. It must consist of a string of 6 - 64 printable ASCII characters. Enclose the password in single quotation marks. 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. Do not enclose the password in single quotation marks if you use the prompt. 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.

-usergrp ***group_name | group_id***
(Optional) Specifies the new group for the user.

-lock
(Optional) Specifies to lock the account indefinitely and it cannot be logged in to unless unlocked again with the **-unlock** flag.

-unlock
(Optional) Specifies to unlock the account so it can be logged in again.

-forcepasswordchange
(Optional) Specifies the account is unable to run svctask commands until the password is changed.

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

-certuid ***<unique_identifier>***
(Optional) Specifies a Unique ID that can be used by the authentication services that might run on the system. It must consist of a string up to *255* printable ASCII characters.

-nocertuid
(Optional) Deletes the **certuid** associated with the user.

Note: The **-certuid** parameter is mutually exclusive with **-keyfile** and **-nocertuid**.

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. You can issue all commands except for **sainfo** and **satask** commands. These commands can only be issued by user superuser.

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 can have an SSH key but cannot have a password.

If you use the **keyfile** parameter, the SSH key file should be placed in the /tmp directory before you run 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.

When two person integrity (TPI) is enabled, you can still use the **chuser** command as a restricted security administrator with the following exceptions:

- The target user cannot belong to a user group of security administrator role.
- You cannot change the non security-administrator user-group of the target user to a security administrator user group.
- You cannot change a remote user.

To be able to run the two restricted actions mentioned, you need an approved role elevation. But even holding an approved role elevation, changes are not allowed:

- If they do not comply with the minimum security administrator user requirements of TPI, or
- If the target user has an active role elevation.

An invocation example

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

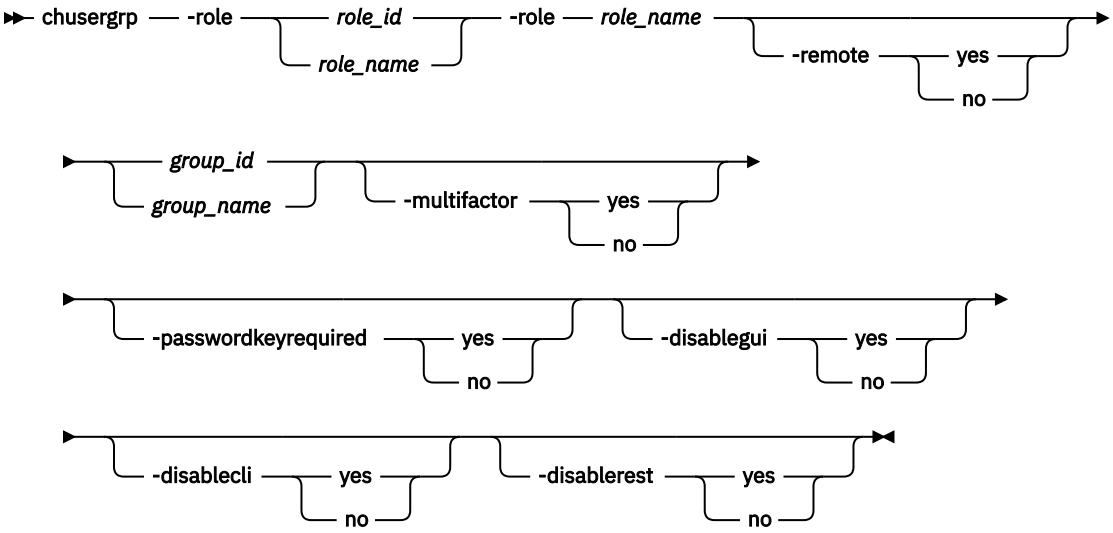
The resulting output:

```
No feedback
```

chusergrp

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

Syntax



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, FlashCopyAdmin, Service, Administrator, or SecurityAdmin.

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

group_id* | *group_name
(Required) The ID or name of the user group whose attributes are to be changed.

-multifactor ***yes* | *no***
(Optional) Specifies whether the multifactor authentication access must be disabled for this user group. Value can be *yes* or *no*.

Note: The **-multifactor** parameter can only set to *yes* when a multifactor authentication service has been enabled first. All users in the user group uses this setting, except for the superuser.

-passwordkeyrequired ***yes* | *no***
(Optional) Specifies if this user group should have password and SSH key required or not. Value can be *yes* or *no*.

-disablegui ***yes* | *no***
(Optional) Specifies whether the GUI access must be disabled for this user group. Value can be *yes* or *no*.

-disablecli ***yes* | *no***
(Optional) Specifies whether the CLI access must be disabled for this user group. Value can be *yes* or *no*.

-disablerest ***yes* | *no***
(Optional) Specifies whether the REST-API access must be disabled for this user group. Value can be *yes* or *no*.

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. You can issue all commands except for **sainfo** and **satask** commands. These commands can only be issued by user superuser.

The roles of the default groups cannot be changed.

When two person integrity (TPI) is enabled, you can still use the **chusergrp** command as a restricted security administrator with the following exceptions:

- the target user group cannot have a current role of security administrator.
- you cannot change the role of a user group from a non security administrator role to a security administrator role.

In order to be able to run the two restricted actions mentioned you will need an approved role elevation. But even holding an approved role elevation, you will not be allowed changes:

- if they do not comply with the minimum security administrator user requirements of TPI, or
- if the target user group has one or more members that have active role elevations.

An invocation example

```
chusergrp -multifactor no
Modifying the authentication setting for this user group will affect logins for all users in
the group.
Are you sure you want to continue? (y/yes to confirm)  yes
```

The resulting output:

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

lscurrentuser

Use the **lscurrentuser** command to display information for the logged-in user.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

Description

When two person integrity (TPI) is enabled, you belong to a security administrator user group, and you run the **lscurrentuser** command, the following information is displayed:

- If you don't have an active role elevation request, the role shows Restricted Security Administrator.
- If you have an active role elevation request, the role shows Security Administrator.

If TPI is not enabled, or it TPI is enabled but you don't belong to the security administrator user group, then the role always matches the user role of the user group.

The Restricted Security Administrator role has access to everything Administrator role has. In addition it allows access to these commands:

- chuser
- mkuser
- rmuser
- chusergrp
- mkusergrp
- rmusergrp

with some restrictions. For example, Restricted Security Administrator role cannot manipulate users or user groups that pertain to Security Administrator role.

The following attribute values can be displayed in the output:

Table 151. ***lscurrentuser*** *output*

Attribute	Description
name	Indicates the name of the current user.
role	Indicates the role of the current user.
owner_id	Indicates the ownership group ID of the current user.
owner_name	Indicates the ownership group name of the current user.
password_expiry_time	Indicates the time when the password expires for the current user.
multi_factor	Indicates whether the user currently logged in is a member of a user group with multi-factor authentication enabled. Value can be *yes* or *no*.
password_sshkey_required	Indicates whether the users currently logged in is a member of a user group with password and SSH key required. Value can be *yes* or *no*.
gui_disabled	Indicates whether the users currently logged in is a member of a user group with GUI access disabled. Value can be *yes* or *no*.
cli_disabled	Indicates whether the users currently logged in is a member of a user group with CLI access disabled. Value can be *yes* or *no*.
rest_disabled	Indicates whether the users currently logged in is a member of a user group with REST-API access disabled. Value can be *yes* or *no*.
cim_disabled	(Deprecated) CIMOM is disabled and cannot be re-enabled.

lsldap

Use the **lsldap** command to display the details for the system-wide Lightweight Directory Access Protocol (LDAP) configuration.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

Description

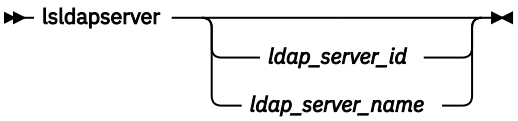
This table provides the attribute values that can be displayed as output view data.

Table 152. lsldap attribute values	
Attribute	**Value**
type	Indicates the LDAP server type. The values are: • `ad` indicates that it is an Active Directory server. • `itds` indicates that it is an IBM Tivoli® Directory Server. • `other` indicates that it is another type of server.
enabled	Indicates whether native LDAP authentication is enabled. The value is `yes` or no/
error_sequence_numbe r	Indicates the sequence number of non-fixed LDAP configuration error log. The value is a number (integer).
username	Indicates the binding user name or distinguished name. The value is an alphanumeric string or blank if there is no name.
security	Indicates the type of security in use. The values are: • `tls` indicates that it is Transport Layer Security. • `ssl` indicates that it is Secure Socket Layer. • `none` indicates that there is no security.
user_attribute	Indicates the LDAP attribute that represents the user login.
group_attribute	Indicates the LDAP attribute that represents the user group membership.
audit_log_attribute	Indicates the LDAP attribute that represents the user name in audit log.
auth_cache_minutes	Indicates the period (in minutes) for which to cache session details.
nested_group_search	Indicates the handling of nested groups. The values are: • `off` indicates that there is no nested group handling search. • `client` indicates that the system must search for nested groups on the client. • `server` indicates that the system must search for nested groups on the server.

lsldapserver

Use the **lsldapserver** command to display the most recent details for all configured Lightweight Directory Access Protocol (LDAP) servers.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

ldap_server_id* | *ldap_server_name
(Optional) Specifies the ID or name for LDAP server that is being used.

Description

Remember:

- The base distinguished name (DN) is at the end of the concise view information; other fields must be added before the base DN.
- The command fails if a server is specified that does not exist.

This table provides the attribute values that can be displayed as output view data.

*Table 153. **lsldapserver** attribute values*

Attribute	Value
id	Specifies the ID of the LDAP server.
name	Specifies the name of the LDAP server.
error_sequence_number	Specifies the sequence number of non-fixed LDAP server error log.
IP_address	Specifies either a server IP address (Internet Protocol Version 4 or 6), or a fully-qualified domain name (maximum of 256 characters).
port	Specifies the LDAP server port. The default value is 389. The value for TLS security is 389 and the value for SSL security is 636.
cert_set	Specifies the certificate setting if a certificate is configured.
preferred	Specifies the server preference (preferred server).
base_dn	Specifies the base distinguished name (DN) that is used in LDAP searches.

Description

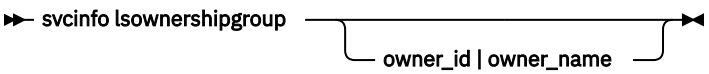
This command displays details for the configured LDAP servers.

Note: There is a maximum of six configured LDAP servers.

lsownershipgroup

The **lsownershipgroup** command displays the ownership groups that are available in the system.

Syntax



Parameters

owner_id | owner_name
(Optional) Specify the ID or the name of the ownership group to display.

Description

Use the **lsownershipgroup** command to display the ownership groups that are available in the system.

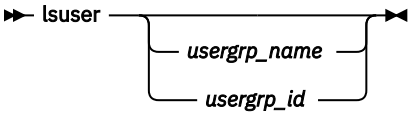
This table provides the attribute values that can be displayed as output view data.

Table 154. lsownershipgroup output	
Attribute	**Description**
replication_link_uid	Indicates the link UID. This value is used when linking ownership groups between systems.

lsuser

Use the **lsuser** command to display a list of the users that are created on the system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

usergrp_name | usergrp_id
(Optional) Specifies the ID or name of the user for which the association is being deleted. If this parameter 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 is created on the system.

This table provides the possible values that are displayed for the **lsuser** command.

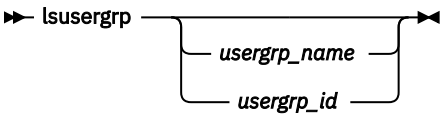
*Table 155. **lsuser** attribute values*	
Attribute	**Value**
id	Specifies the ID of the user.
name	Specifies the name of the user.
password	Specifies if a password is set for the user.
ssh_key	Specifies if an SSH key is set for the user.
remote	Specifies if the user is remote.
usergrp_id	Specifies the ID of the user group that the user is a member of.
usergrp_name	Specifies the name of the user group that the user is a member of.
owner_id	Specifies the owner ID of the user group that the user is a member of.

Table 155. **lsuser** attribute values (continued)	
Attribute	**Value**
owner_name	Specifies the owner name of the user group that the user is a member of.
locked	Specifies whether a user is locked or not. The possible values are: • no - The user is unlocked. • auto (blank) - The user is permanently locked due to exceeding the **max_failed_login_attempts** when the value of **lockout_period** is 0. • auto (timestamp) - The user is temporarily locked until the time specified due to exceeding the **max_failed_login_attempts** when the value of **lockout_period** is in the range 1 - 10080. • manual - The user is permanently locked due to the security administrator using the `chuser -lock id` command. **Note:** When locked is set to manual, the superuser is permanently locked. This permanent locking occurs when the security administrator runs the **chuser -lock superuser** command or enables two person integrity (TPI) with the **chsecurity -twopersonintegrity yes** command.
locked_until	Specifies a timestamp that displays the time in the format YYMMDDHHMMSS when the user account automatically unlocks.
password_expiry_time	Specifies a timestamp that displays the time in the format YYMMDDHHMMSS when the user's password expires.
password_change_required	Specifies whether a user's password is expired. The value is either yes or no. The value is set to yes when the password expires or if the `chuser -forcepasswordchange` command is used.
cert_uid	Specifies a Unique ID that can be used by the authentication services that might run on the system. It must consist of a string up to *255* printable ASCII characters.

lsusergrp

Use the **lsusergrp** command to display a list of the user groups that is created on the system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

usergrp_name | usergrp_id
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 table provides the attribute values that can be displayed as output view data.

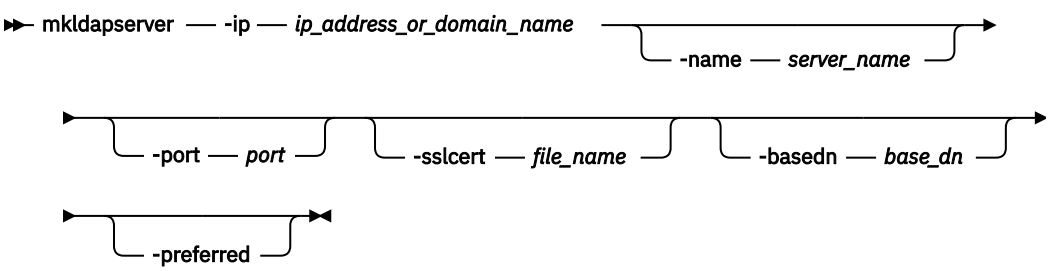
*Table 156. **lsusergrp** output*

Attribute	Description
id	Indicates the ID of the user group.
name	Indicates the name of the user group.
role	Indicates the role of the user group.
remote	Indicates whether users in the user group are authenticated by a remote service.
owner_id	Indicates the ownership group ID of the user group.
owner_name	Indicates the ownership group name of the user group.
multi_factor	Indicates whether the users of this user group should be prompted for second factor authentication. Value can be *yes* or *no*.
password_sshkey_required	Indicates whether the users of this user group should provide both password and SSH public key during authentication. Value can be *yes* or *no*.
gui_disabled	Indicates whether the user currently logged in is a member of a user group with GUI access disabled. Value can be *yes* or *no*.
cli_disabled	Indicates whether the user currently logged in is a member of a user group with CLI access disabled. Value can be *yes* or *no*.
rest_disabled	Indicates whether the user currently logged in is a member of a user group with REST-API access disabled. Value can be *yes* or *no*.
cim_disabled	Indicates whether the user currently logged in is a member of a user group with CIMOM access disabled. Value can be *yes* or *no*.

mkldapserver

Use the **mkldapserver** command to display the data used to create a Lightweight Directory Access Protocol (LDAP) server.

Syntax



Parameters

-ip *ip_address_or_domain_name*
(Required) Specifies either a server IP address or domain name, or a fully qualified domain name (FQDN) (maximum of 256 characters). You can specify an optional fully qualified domain name after you establish a DNS server using the **mkdnsserver** command.

-name *server_name*
(Optional) Specifies the LDAP server name.

-port *port*
(Optional) Specifies the LDAP server port. The default value (if you do not specify a value) is 389. If you specify TLS security the value is 389 and if you specify SSL security the value is 636.

-sslcert *file_name*
(Optional) Set the SSL certificate.

-basedn *base_dn*
(Optional) Use the base distinguished name for search.

-preferred
(Optional) Specifies that this server is preferred over other configured LDAP servers.

Description

Important: During normal operation, LDAP requests are sent to **-preferred** servers depending on availability. If no servers are marked as **-preferred**, LDAP requests are sent to configured servers based on availability.

If **-sslcert** is specified, the server certificate is verified while authenticating.

Note: The SSL certificate must exist on the current node.

The **-basedn** parameter indicates the distinguished name (DN) to use as a base from which to search for users in the LDAP directory. If Transport Layer Security (TLS) is enabled and **-sslcert** is specified, the server certificate is verified during authentication. The secure socket layer (SSL) certificate must exist on the node being used, otherwise a server certificate is not checked.

Note: The Microsoft AD server must use **-basedn**.

The system must be configured with an appropriate version IP address when **-ip** is specified. The IP address specified with the **-ip** parameter must be of a version supported by the system. The certificate file must be in valid PEM format and have a maximum length of 12 kilobytes.

Distinguished names must be a sequence of `attribute=value` pairs separated by a comma (`,`), semi-colon(`;`), or plus sign (+) escaping special characters with a backslash (\) where appropriate, and specified UTF-8 characters using their byte encoding. For example, `\,` for commas or `\C4\87` for the UTF-8 character c acute.

This command runs whether or not LDAP authentication is enabled.

Remember: There is a maximum of six configured LDAP servers. Attempting to create a seventh LDAP server returns an error.

An invocation example

```
mkldapserver -ip 192.135.60.3
```

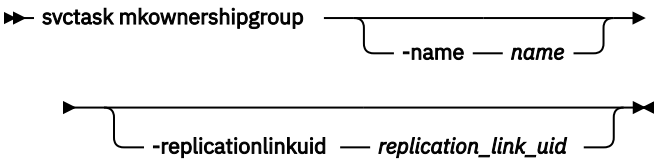
The resulting output:

```
LDAP Server, id [0], successfully created
```

mkownershipgroup

The **mkownershipgroup** command creates an ownership group.

Syntax



Parameters

-name *name*
(Optional) Specify the name of the new ownership group. The default is the ownershipgroup ID.

-replicationlinkuid *replication_link_uid*
(Optional) Specify a link UID value to assign to the new ownership group.

Description

Use the **mkownershipgroup** command to create an ownership group.

An invocation example

```
svctask mkownershipgroup
```

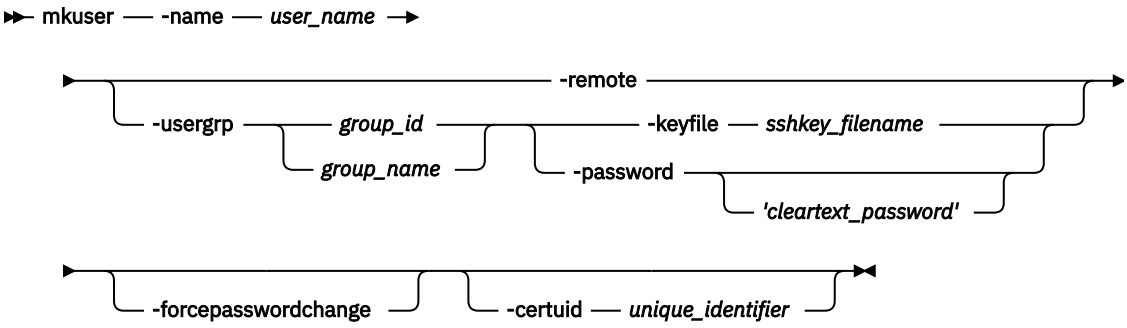
The resulting output:

```
Ownership Group, id[0], successfully created
```

mkuser

Use the **mkuser** command to create either a local or a remote user to access a system.

Syntax



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, except for the following characters: % : " , * '.

-remote | -usergrp
(Required) Specifies whether the user authenticates to the system by using a remote authentication service or system authentication methods. Either the **remote** parameter or the **usergrp** parameter must be set. If **usergrp** is specified, it must be followed by *group_name* or *group_id* (see next parameter).

group_name* | *group_id
(Required if **usergrp** is specified) The ID or name of the user group with which the local user is to be associated.

-password ***'cleartext_password'***
(Optional) Specifies the password to be associated with the user. The password cannot start or end with a blank. It must consist of a string of 6 - 64 printable ASCII characters. You can optionally specify the password with the **password** parameter. Enclose the password in single quotation marks. 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. You cannot specify a password if **remote** is set. Do not enclose the password in single quotation marks if you use the prompt.

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

-forcepasswordchange
(Optional) Specifies the account is required to change the password on the next login.

-certuid ***certuid***
(Optional) Specifies a Ceph Node or VASA Provider or other services certificate unique identifier for authentication. It must consist of a string up to *255* printable ASCII characters.

Note: The **-certuid** parameter is mutually exclusive with **-keyfile**.

Description

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

You must have the security administrator role to create, delete, or change a user. You can issue all commands except for **sainfo** and **satask** commands. These commands can be issued by user superuser only.

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

If you create a remote user, you may specify the **keyfile** parameter. Remote users have their groups that are defined by the remote authentication service.

Up to 150 users can be defined on the system. You can also create new users and assign keys to them.

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

When two person integrity (TPI) is enabled, you can still use the **mkuser** command as a restricted security administrator with the following exceptions:

- You cannot create a new user in a security administrator user group.
- You cannot create a new remote user.

In order to create a new user in a security administrator user-group, or to create a new remote user, you need an approved role elevation. Also when TPI is enabled, the new user name cannot already exist in active role elevations. This can happen if remote LDAP users have active role elevations.

An invocation example

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

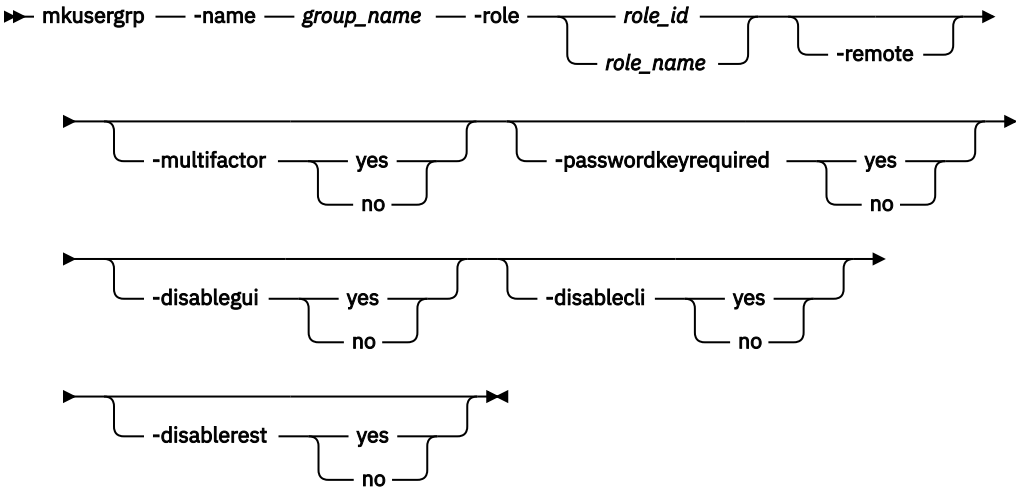
The resulting output:

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

mkusergrp

Use the **mkusergrp** command to create a new user group.

Syntax



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, except for the following characters: `%:",*'`.

-role ***role_id | name***
(Required) Specifies the role (by ID or name) to be associated with all users that belong to this user group. One of the following roles must be selected:

- `Monitor`
- `CopyOperator`
- `Service`
- `Administrator`
- `SecurityAdmin`
- `VasaProvider`
- `RestrictedAdmin`
- 3SiteAdmin

-remote ***yes | no***
(Optional) Specifies if this user group should be used to set the role of remote users. The default value is **no**.

-multifactor ***yes | no***
(Optional) Specifies whether the multifactor authentication access must be disabled for this user group. Value can be *yes* or *no*.

Note: The **-multifactor** parameter can only set to *yes* when a multi-factor authentication service has been enabled first. All users in the user group will use this setting, except for the superuser.

-passwordkeyrequired *yes* **|** *no*
(Optional) Specifies if this user group should have password and SSH key required or not. Value can be *yes* or *no*.

-disablegui *yes* **|** *no*
(Optional) Specifies whether the GUI access must be disabled for this user group. Value can be *yes* or *no*.

-disablecli *yes* **|** *no*
(Optional) Specifies whether the CLI access must be disabled for this user group. Value can be *yes* or *no*.

-disablerest *yes* **|** *no*
(Optional) Specifies whether the REST-API access must be disabled for this user group. Value can be *yes* or *no*.

Description

The **mkusergrp** command creates a new user group to organize users of the system by role. Use the **lsusergrp** command to view a list of user groups that have been created on the clustered system.

You must have the security administrator role (SecurityAdmin role name) to create, delete, or change a user group. For more information, refer to User role.

When two person integrity (TPI) is enabled, a restricted security administrator user will need an approved role elevation to create new user groups of security administrator role.

The command returns the ID of the created user group.

An invocation example

```
mkusergrp -name MFA_Disabled_Group -role 0 -multifactor no -passwordkeyrequired yes
Modifying the authentication setting for this user group will affect logins for all users in
the group.
Are you sure you want to continue? (y/yes to confirm)  yes
```

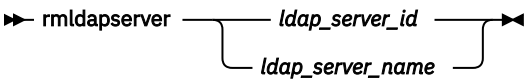
The resulting output:

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

rmldapserver

Use the **rmldapserver** command to delete a Lightweight Directory Access Protocol (LDAP) server.

Syntax



Parameters

ldap_server_id* | *ldap_server_name
(Required) Specifies the LDAP server ID or name to delete.

Description

Remember:

- If remote authentication with LDAP is enabled, the final LDAP server cannot be deleted. To delete the final LDAP server disable LDAP authentication by specifying **chauthservice -enable no**.
- The **rmldapserver** command can be specified whether or not LDAP authentication is enabled.

An invocation example

```
rmldapserver ldapserver0
```

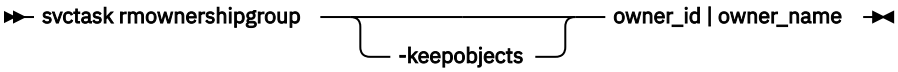
The resulting output:

```
No feedback
```

rmownershipgroup

The **rmownershipgroup** command deletes an ownership group.

Syntax



Parameters

-keepobjects
(Optional) If specified, then the objects that currently belong to the ownership group will be kept but will be moved to noownershipgroup.

owner_id | owner_name
(Required) Specify the ID or the name of the ownership group to delete.

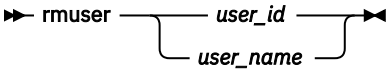
Description

Use the **rmownershipgroup** command to delete an ownership group.

rmuser

Use the **rmuser** command to delete a user.

Syntax



Parameters

user_id* or *user_name
(Required) Specifies the user to be removed.

Description

Use the **rmuser** command to delete a user.

You must have the security administrator role to create, delete, or modify a user. You cannot remove yourself by using this command.

When two person integrity (TPI) is enabled, a restricted security administrator can still use the **rmuser** command except that it cannot remove a target user that belongs to a security administrator user-group and it cannot remove a target user that is remote. To remove a target user that belongs to a security administrator user group, or to remove a target user that is remote, an approved role elevation is required. But even with an approved role elevation, **rmuser** is not allowed:

- If the target user has an active role elevation.

- If removing the target user results in no compliance of the minimum security administrator user requirements of TPI.

An invocation example

```
rmuser jane
```

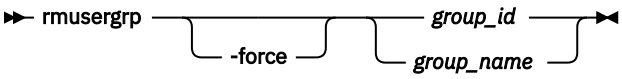
The resulting output:

```
No feedback
```

rmusergrp

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

Syntax



Parameters

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

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

group_id* | *group_name
(Required) The ID or name of the user group to be removed.

Description

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

You must have the security administrator role to create, delete, or change a user group.

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

When two person integrity (TPI) is enabled, a restricted security administrator can still use the **rmusergrp** command except that it cannot remove a target user group that has security administrator role.

In order to remove a target user group that has security administrator role, an approved role elevation is required. But even holding an approved role elevation, one cannot remove a user group that has security administrator role, if the end result will compromise the minimum security administrator user requirements of TPI.

An invocation example

```
rmusergrp support
```

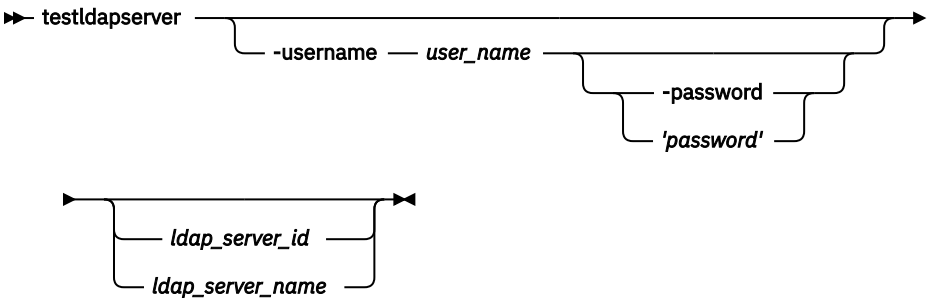
The resulting output:

```
No feedback
```

testldapserver

Use the **testldapserver** command to test a Lightweight Directory Access Protocol (LDAP) server.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-username user_name
(Optional) Specifies the user name to test.

-password *'password'*
(Optional) Specifies the password to test. You can optionally specify the password with this parameter. You must enclose the password in single quotation marks. 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. Do not enclose the password in single quotation marks if you use the prompt.

Note: The **-password** parameter is only valid if **-username** is specified. The actual password does not need to be supplied.

ldap_server_id|*ldap_server_name*
(Optional) Specifies the LDAP server ID or name to test.

Description

The **testldapserver** command allows three levels of testing:

- Server connection test (issue **testldapserver** without supplying username or password). This verifies that a connection can be established with the server while authenticating using the configured administrator credentials according to the LDAP configuration.
- Server connection, LDAP configuration, and user authorization test (issue **testldapserver** with a username). This verifies that:
  - A connection can be established with the server while authenticating using the configured administrator credentials.
  - The LDAP attributes are correctly configured on the system.
  - The user has been assigned a role.
- Server connection, LDAP configuration, and user authentication test (issue **testldapserver** with a username and password). This verifies that:
  - A connection can be established with the server while authenticating using the configured administrator credentials.
  - The user authenticates with the supplied password

No specific server errors indicates success.

Important: This command works whether or not LDAP authentication is selected or enabled with the **chauthservice** command.

This table provides the attribute values that can be displayed as output view data.

*Table 157. **testldapserver** attribute values*

Attribute	Value
id	LDAP server ID
name	LDAP server name
error	Critical server error (or success, depending on situation) encountered

An invocation example with one LDAP server and no specific user information

```
testldapserver -delim ":" ldapserver1
```

The resulting output:

```
id:name:error
1:ldapserver1:CMMVC7075I The LDAP task completed successfully
```

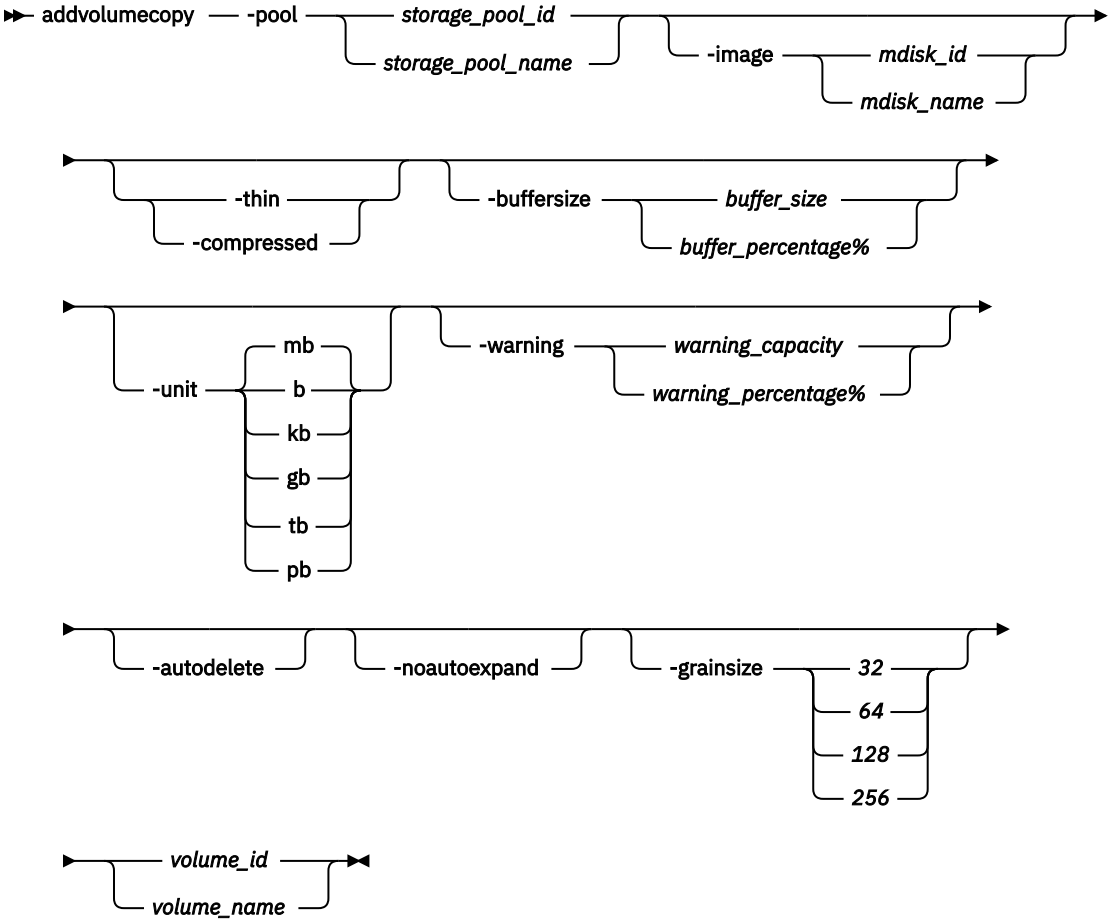
Chapter 45. Volume commands

Use the volume commands to work with volume options for the system.

addvolumecopy

Use the **addvolumecopy** command to add a copy to an existing volume.

Syntax



Parameters

-pool *storage_pool_id | storage_pool_name*
(Required) Specifies the storage pool in which to create the new volume copy.

-image *mdisk_id | mdisk_name*
(Optional) Specifies that the volume copy is to be created (on any topology) in image mode and it specifies which currently unused MDisk to use.

-thin
(Optional) Specifies that the volume copy is to be created with thin provisioning. You cannot specify this parameter with **-compressed**. If you are adding a volume to a pool that uses a provisioning policy, this parameter is not enabled.

© Copyright IBM Corp. 2003, 2025

-compressed
(Optional) Specifies that the volume copy is to be created compressed. You cannot specify this parameter with **-thin**. If you are adding a volume to a pool that uses a provisioning policy, this value is disabled.

-buffersize ***buffer_size | buffer_percentage%***
(Optional) Specifies the pool capacity the volume attempts to reserve as a buffer for thin-provisioned and compressed volumes. You must specify either **-thin** or **-compressed** with this parameter. If you are adding a volume to a pool that uses a provisioning policy, this parameter uses the default value of 2%.

-warning ***warning_capacity | warning_percentage%***
(Optional) Specifies a threshold at which a warning error log is generated for the volume copy. A warning is generated when the used disk capacity on the thin-provisioned or compressed copy exceeds the specified threshold. You can specify the threshold by using *warning_capacity* to specify a size, which defaults to MB unless the **-unit** parameter is specified.If you are adding a volume to a pool that uses a provisioning policy, this parameter uses the default value of 80%.

Important: You can also specify a *warning_percentage%* to use a percentage of the volume size. If you do not specify a warning threshold, a default value of 80% is used. To disable warnings, specify 0.

You must specify either **-thin** or **-compressed** with this parameter.

-unit b | kb | mb | gb | tb | pb
(Optional) Specifies the data units for the buffer_size and -warning parameters.

-noautoexpand
(Optional) Specifies that the volume copy not automatically expand as it is written to; the available buffer capacity decreases as the used capacity increases. The copy goes offline if the buffer capacity is consumed.

The buffer capacity can be increased by specifying expandvdisksize -rsize. You must specify either **-thin** or **-compressed** with this parameter. If you do not specify this keyword, the copy automatically expands as it is written to.

If you are adding a volume to a pool that uses a provisioning policy, the volume automatically expands as it is written to.

-autodelete
(Optional) Specifies that the source volume is deleted after the new volume copy is synchronized.

-grainsize ***32 | 64 | 128 | 256***
(Optional) Sets the grain size (KB) for a thin-provisioned volume. If you are using the thin-provisioned volume in a FlashCopy map, use the same grain size as the map grain size for best performance. If you are using the thin-provisioned volume directly with a host system, use a small grain size. The grain size value must be 32, 64, 128, or 256 KB. The default is 256 KB.

volume_id | volume_name
(Required) Specifies the volume to add the volume copy to.

Description

Use the **addvolumecopy** command to add a copy to an existing volume. The new volume copy is synchronized with the current copy.

You can also use the **addvolumecopy** command to add a volume copy to an existing volume that has snapshots provided all the snapshots are mirrored. The parent pool of at least one volume copy of each snapshot must be in the same parent pool as the copy being added. You must use the **addvdiskcopy** command to add a copy to a VDisk that maintains snapshots.

When you create a volume copy in a storage pool with a provisioning policy, some parameters use default settings or the settings from the provisioning policy.

The **-autodelete** parameter deletes the source volume when the new volume copy is synchronized.

You cannot specify **buffer_size** if the volume copy is created in a data reduction storage pool. Specify **-thin** or **-compressed** to enable thin provisioning or compression.

You cannot specify **-noautoexpand** when you create thin-provisioned or compressed volume copies from a data reduction storage pool.

You cannot create a volume copy that is a thin-provisioned or compressed volume in a data reduction storage pool, and the volume caching mode is `none` or `readonly`. You must specify **chvdisk** to change the volume caching mode to `readwrite`.

You cannot specify **-warning** for a thin-provisioned or compressed volume copy in a data reduction storage pool.

You cannot specify **grain_size** for thin-provisioned and compressed volume copies in data reduction storage pools. This type of volume copy is created with a size of 8 KB.

Thin-provisioned or compressed volume copies in data reduction pools cannot be created if the data reduction storage pool is offline and requires recovery. If the recovery still in progress, you must wait until the recovery is complete,theis and pool is in `online` state.

Add a volume copy to an existing volume

```
addvolumecopy -pool 2 volume5
```

The detailed resulting output:

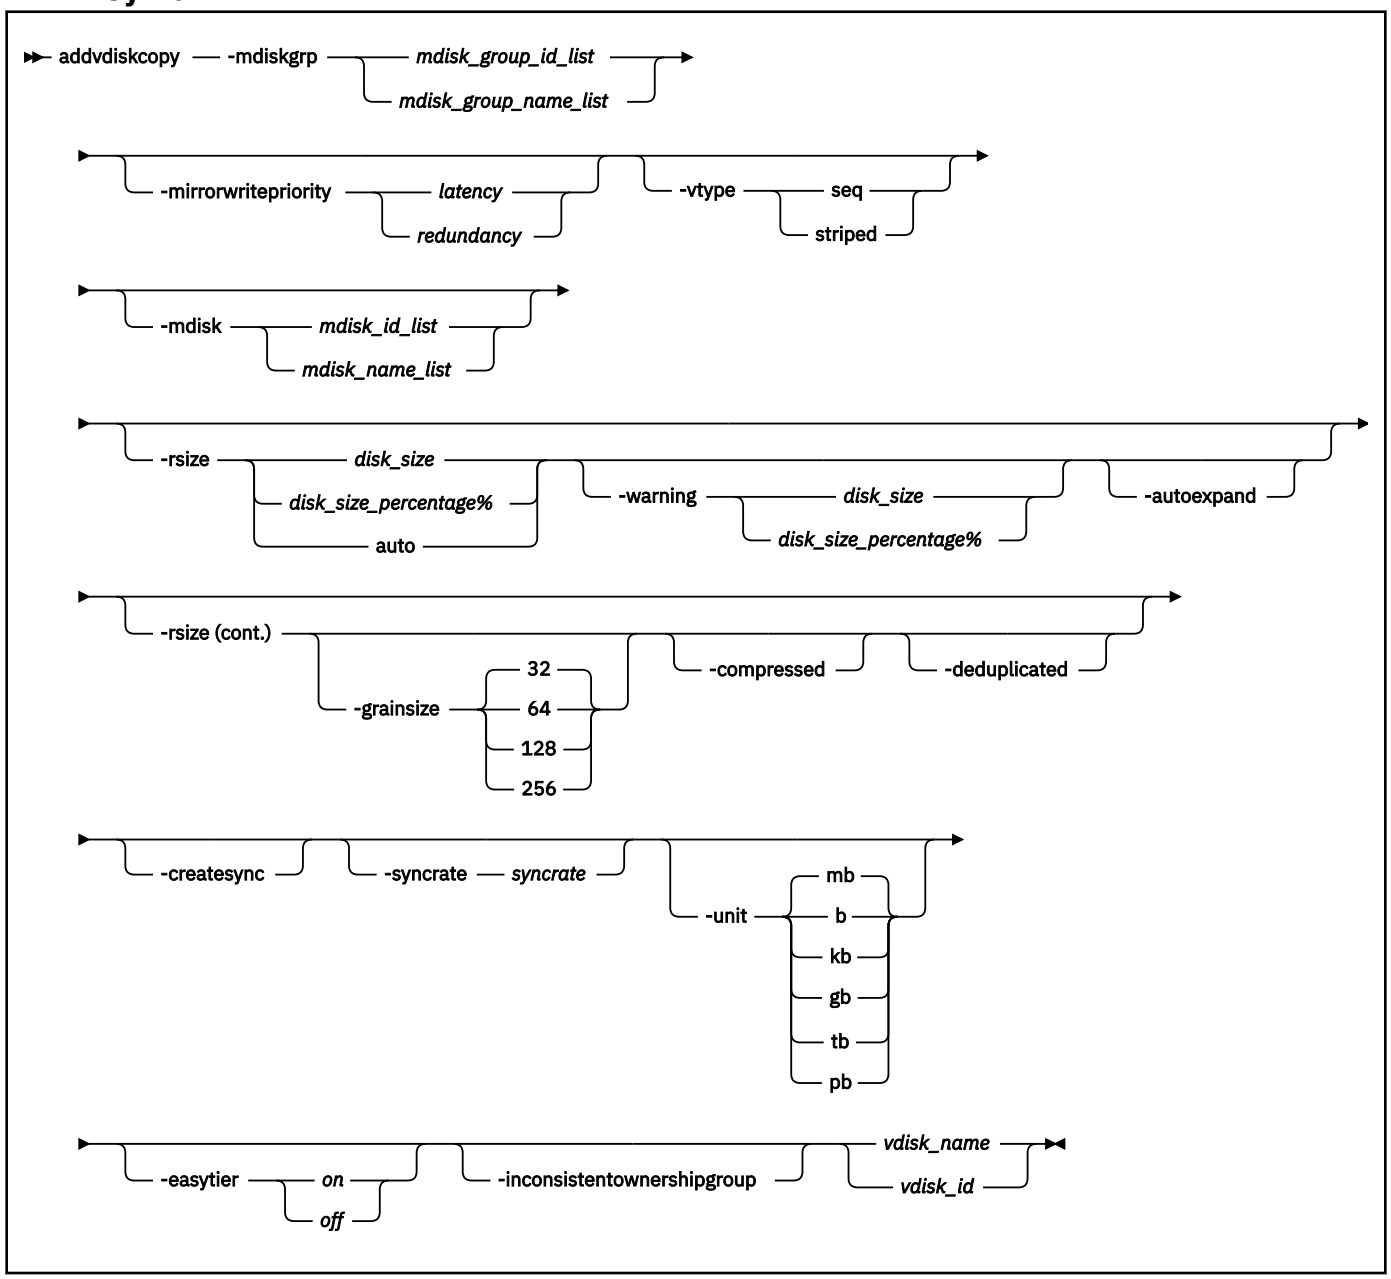
```
No feedback
```

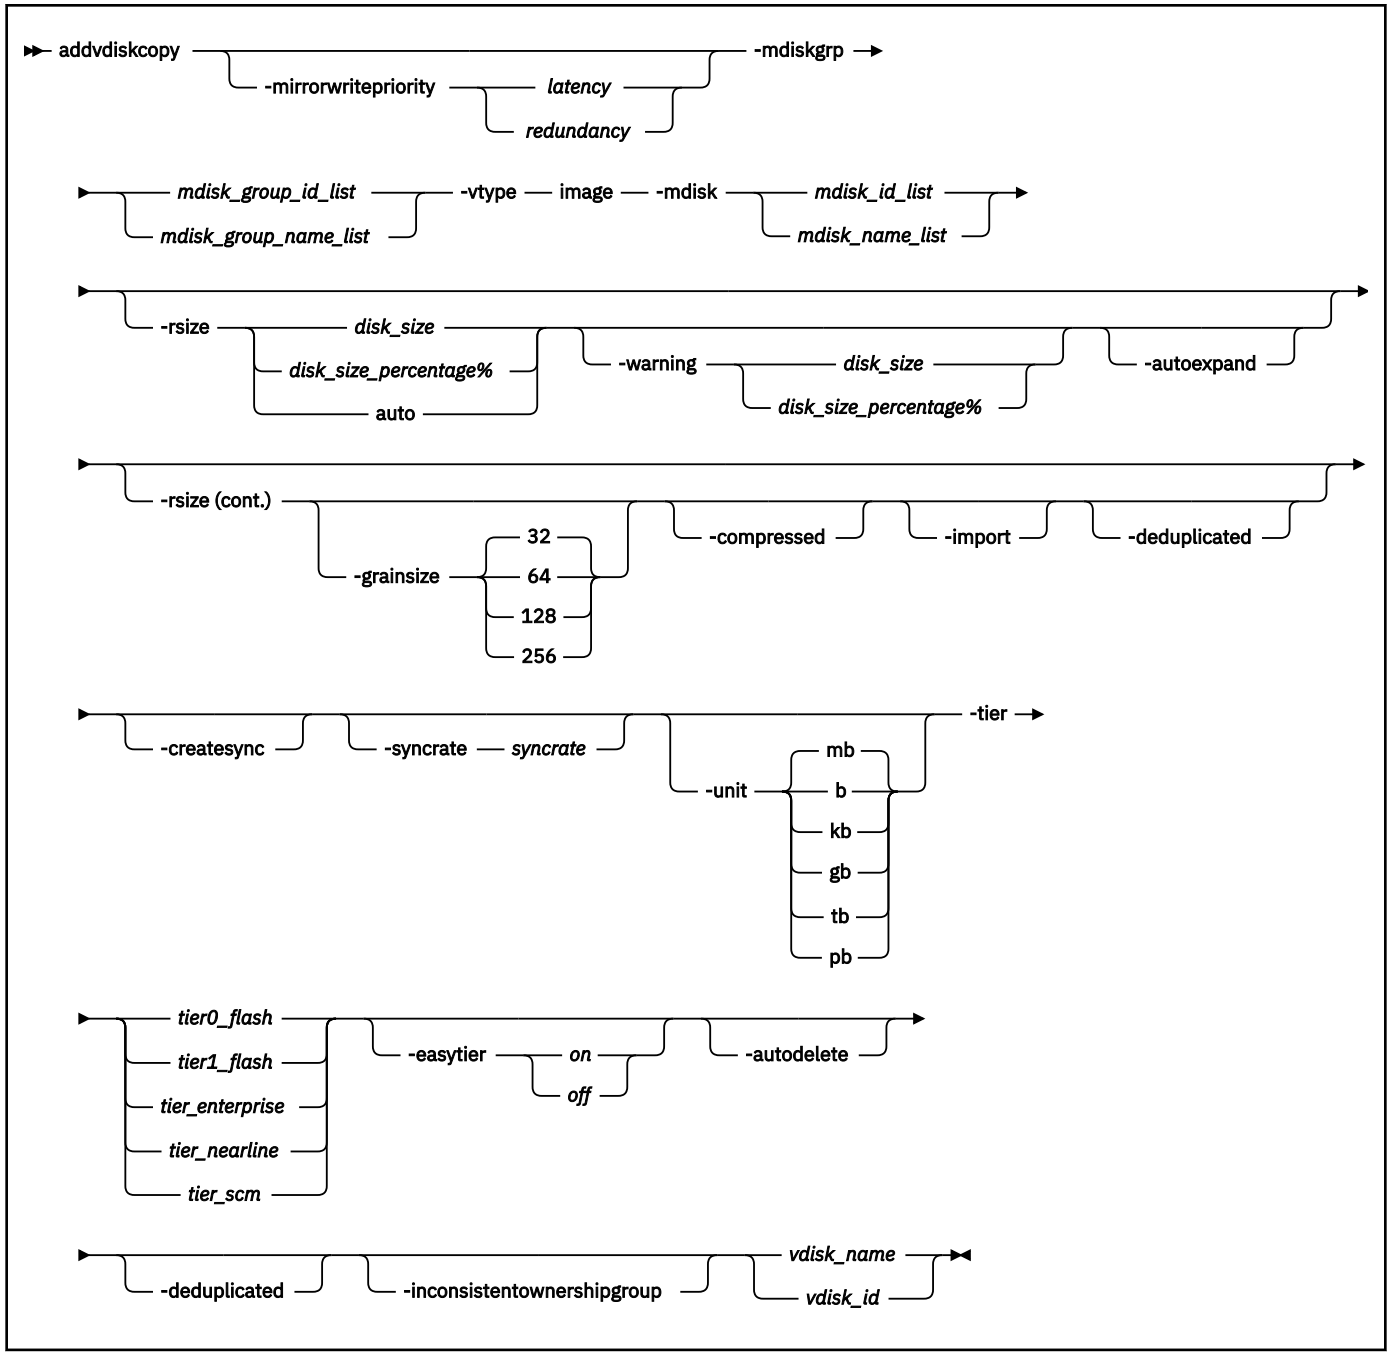
addvdiskcopy

Use the **addvdiskcopy** command to add a copy to an existing volume. In this way, you change a non-mirrored volume into a mirrored volume.

Note: The first syntax diagram depicts the addition of a sequential or striped mode volume. The second syntax diagram depicts the addition of an image mode volume.

Syntax


addvdiskcopy
-mdiskgrp
mdisk_group_id_list
mdisk_group_name_list
-mirrorwritepriority
latency
redundancy
-vtype
seq
striped
-mdisk
mdisk_id_list
mdisk_name_list
-rsize
disk_size
disk_size_percentage%
auto
-warning
disk_size
disk_size_percentage%
-autoexpand
-rsize (cont.)
-grainsize
32
64
128
256
-compressed
-deduplicated
-createsync
-syncrate
syncrate
-unit
mb
b
kb
gb
tb
pb
-easytier
on
off
-inconsistentownershipgroup
vdisk_name
vdisk_id



Parameters

-mdiskgrp *mdisk_group_id_list | mdisk_group_name_list*
(Required) Specifies the storage pools to use to create copies for the volume. You must specify a group for each copy that is being added.

Important: If the MDisk group is from a child pool, the **-vtype** parameter must be set to `striped`.

-mirrorwritepriority *latency | redundancy*
(Optional) Specifies how to configure the mirror write algorithm priority.

- Choosing *latency* means that a copy that is slow to respond to a write input/output (I/O) becomes unsynchronized, and the write I/O completes processing if the other copy successfully writes the data.
- Choosing *redundancy* means that a copy that is slow to respond to a write I/O synchronizes completion of the write I/O with the completion of the slower I/O to maintain synchronization.

-vtype seq | striped | image
(Optional) Specifies the virtualization type for the copy: sequential, striped, or image. The type can be different from the virtualization types for other copies on the volume. The default virtualization type is `striped`. If you specify the `-rsize auto` option or the `-import` option, you must also specify the `-vtype image` option.

Important: You cannot create an image or sequential mode copy from a child pool or data reduction pools.

Restriction: You cannot use `-vtype image` parameter for the volumes with policy-based replication.

-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 managed mode.

-syncrate ***syncrate***
(Optional) Specifies the copy synchronization rate. A value of 0 prevents synchronization. For the supported `-syncrate` values and their corresponding rates, see the table in the command description section. If not specified, the current value is unchanged.

-createsync
(Optional) Suppresses the synchronization of the new volume copy with the primary copy.

Attention: 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.

Remember: You cannot specify **-createsync** for a volume that is fast formatting and for a volume with policy-based replication.

-rsize ***disk_size | disk_size_percentage%*** **| auto**
(Optional) Makes the copy thin-provisioned and specifies the real size of the copy. Specify the *disk_size* | *disk_size_percentage* value by using an integer, or an integer immediately followed by the percent character (**%**). The default units for *disk_size* are megabytes (MB). To specify different units, use the `-unit` parameter. The `auto` option creates a volume copy that uses the entire size of the MDisk; if you specify the `-rsize auto` option, you must also specify the `-vtype image` option.

-deduplicated
(Optional) Adds a deduplicated volume. If you specify **-deduplicated**, you must also specify **-rsize** because it applies only to thin-provisioned or compressed volumes.

Restriction: Data deduplication works only with data reduction storage pools. You can create deduplicated volumes and volume copies in an I/O group only if no compressed volumes or volume copies exist in regular storage pools.

-compressed
(Optional) Add exactly one copy to an existing volume that already has one copy a volume, and enables compression. If you specify the **-compressed** parameter, you must also specify the **-rsize** parameter.

Important:

- You cannot specify this parameter with the `grain_size` parameter.
- When you specify this parameter with the `-import` parameter, you must specify the `-rsize auto` parameter.

-warning ***disk_size | disk_size_percentage%***
(Optional) Generates a warning when the used disk capacity on the thin-provisioned copy first exceeds the specified threshold. You can specify a *disk_size* integer, which defaults to megabytes (MB) unless the `-unit` parameter is specified, or you can specify a *disk_size*%, which is a percentage of the volume size. If `-autoexpand` is enabled, the default value for `-warning` parameter is 80% of

the volume capacity. If `-autoexpand` is not enabled, the default value for warning is 80% of the real capacity. To disable warnings, specify **0**.

Restriction: Requires that the `-rsize` parameters also are specified.

-autoexpand
(Optional) Specifies that thin-provisioned copies automatically expand their real capacities by allocating new extents from their storage pool. If the `-autoexpand` parameter is specified, the `-rsize` parameter specifies a capacity that is reserved by the copy. It protects the copy from going offline when its storage pool runs out of space by allowing it to consume this reserved space first.

Restriction: Requires that the `-rsize` parameters also are specified.

-grainsize 32 | 64 | 128 | 256
(Optional) Sets the grain size (KB) for a thin-provisioned volume copy. The grain size value must be 32, 64, 128, or 256 KB. The default is 256 KB.

Restriction: Requires that the `-rsize` parameters also be specified.

-unit b | kb | mb | gb | tb | pb
(Optional) Specifies the data units for the `-rsize` and `-warning` parameters.

-import
(Optional) Imports an image mode disk that contains a thin-provisioned volume into the system. Requires that the `-rsize` and `-vtype image` parameters also be specified.

-tier *tier0_flash | tier1_flash | tier_enterprise | tier_nearline | tier_scm*
(Optional) Specifies the MDisk tier when an image mode copy is added.

tier0_flash
Specifies a `tier0_flash` hard disk drive or an external MDisk for the newly discovered or external volume.

tier1_flash
Specifies a `tier1_flash` (or flash drive) hard disk drive or an external MDisk for the newly discovered or external volume.

tier_enterprise
Specifies a `tier_enterprise` hard disk drive or an external MDisk for the newly discovered or external volume.

tier_nearline
Specifies a `tier_nearline` hard disk drive or an external MDisk for the newly discovered or external volume.

tier_scm
Specifies a `tier_scm` hard disk drive or an external MDisk for the newly discovered or external volume.

-easytier *on | off*
(Optional) Determines whether the IBM Easy Tier function can move extents for this volume. If a volume copy is striped and not being migrated, see the settings in the following table:

-autodelete
(Optional) Specifies that the primary copy is deleted after the secondary copy is synchronized.

-inconsistentownershipgroup
(Optional) Specifies that inconsistent ownership is permitted for this object.

jvdisk_name | vdisk_id
(Required) Specifies the volume to add the volume copy to, either by ID or by name.

Description

The **addvdiskcopy** command adds a copy to an existing volume. In this way, a non-mirrored volume is changed into a mirrored volume. Use the **mkdiskgrp** parameter to specify the storage pools that provide storage for the copy; the **lsmdiskgrp** command lists the available storage pools and the amount of available storage in each group.

You can also use the **addvdiskcopy** command to add a VDisk copy to an existing VDisk that has snapshots provided all the snapshots are mirrored. The parent pool of at least one VDisk copy of each snapshot must be in the same parent pool as the copy being added.

The **addvdiskcopy** command can be specified with a file system volume, but must be used with the same storage pool for that volume.

Remember: Only compressed copies can be added to file system volumes.

The **addvdiskcopy** command adds a different volume copy, such as a copy created from an uncompressed to compressed conversion or a compressed to uncompressed conversion.

Restriction: A volume cannot have volume copies in different storage pools if cloud snapshot feature is enabled on the volume.

A thin-provisioned or compressed volume copy in a data reduction storage pool must not be a sequential or image mode volume. On some node types, you can create a compressed volume copy in a data reduction storage pool for an I/O group. A compressed volume copy in a data reduction pool can be created only in an I/O group with V5030, V7000, or SVC node types. You can create thin-provisioned volume copies on any node type. Use the **-autoexpand** parameter to create thin-provisioned or compressed volume copies from a data reduction storage pool. Volumes can also have standard-provisioned volume copies in data reduction storage pools.

Compressed volumes/copies cannot be created if the I/O group in which the copy is being created supports only software compression and:

- The I/O group already has Real-time compression copies and the copy is data reduction.
- The I/O group already has data reduction copies and the copy uses Real-time compression.

You cannot create a volume copy that is a thin-provisioned or compressed volume in a data reduction storage pool, and the volume caching mode is `none` or `readonly`. You must specify **chvdisk** to change the volume caching mode to `readwrite`.

You cannot specify **-warning** for a thin-provisioned or compressed volume copy in a data reduction storage pool.

For thin-provisioned and compressed volume copies in data reduction storage pools, the Easy Tier mode for the volume is taken from the data reduction storage pool. The Easy Tier mode cannot be configured on these volume types.

You cannot specify **grain_size** for thin-provisioned and compressed volume copies in data reduction storage pools. This type of volume copy is created with a size of 8 KB.

Thin-provisioned or compressed volume copies in data reduction pools cannot be created if the data reduction storage pool is offline and requires recovery. If the recovery is still in progress, you must wait until the recovery is complete and the pool is in `online` state.

An encryption key cannot be used when you add an image mode MDisk. To use encryption (when the MDisk has an encryption key), the MDisk must be self-encrypting.

Remember: You cannot add a volume copy if the volume to be copied is being formatted.

The virtualization types are defined as follows:

sequential (seq)
This policy requires the **-mdisk** parameter with a single managed disk as its argument. This MDisk must be in managed mode.

It creates the volume by using extents from the given managed disk (assuming that enough free extents exist on the managed disk).

striped
The `striped` policy is the default policy. If the **-vtype** parameter is not specified, this policy is used in its default form. That is, all managed disks in the storage pool are used to create the volume. The striping is at an extent level; one extent from each managed disk in the group is used. For example, a

storage pool with 10 managed disks uses one extent from each managed disk, and then the storage pool uses the 11th extent from the first managed disk, and so on.

If the **-mdisk** parameter is also specified, you can supply a list of managed disks to use as the stripe set. This list can include two or more managed disks from the same storage pool. The same circular algorithm is used across the striped set. However, a single managed disk can be specified more than once in the list. For example, if you enter `-m 0:1:2:1`, the extents are from the following managed disks: 0, 1, 2, 1, 0, 1, 2, and so on. All MDisks that are specified in the **-mdisk** parameter must be in managed mode.

image
This policy allows image mode volumes to be created when a managed disk already has data on it, perhaps from a previrtualized subsystem. When an image mode volume is created, it directly corresponds to the (previously unmanaged) managed disk that it was created from; therefore, volume logical block address (LBA) *x* equals managed disk LBA *i*. You can use this command to bring a nonvirtualized disk under the control of the system. After it is under the control of the system, you can migrate the volume from the single managed disk. When it is migrated, the volume is no longer an image mode volume.

You can add image mode volumes to an already populated storage pool with other types of volumes, such as a striped or sequential.

Important: An image mode copy must be at least as large as the volume that it is being added to, but any capacity beyond the size of the volume is not accessible.

The command returns the ID of the newly created volume copy.

Create the first compressed volume copy for an I/O group to activate compression. You cannot create or move a compressed volume copy to an I/O group that contains one or more nodes that do not support compressed volumes. You must use another I/O group, but note that it does not affect moving to the recovery I/O group.

Important:

- If the volume (or volume copy) is a target of a FlashCopy mapping with a source volume in an `active-active` relationship, the new storage pool must be in the same site as the source volume.
- If this command is used for a volume that is a master volume, an auxiliary volume, or a change volume of an `active-active` relationship, the new copy must be created in a storage pool of the same site as the existing volume copy.
- When you add an image mode copy, the site information for the MDisk that is being added must be well-defined and match the site information for any other MDisks in the storage pool.

The rate at which the volume copies resynchronize after loss of synchronization can be specified by using the **-syncrate** parameter. This following table provides the relationship of the *syncrate* value to the amount of data copied per second.

Remember: These settings also affect the initial rate of formatting.

Table 158. Relationship between the syncrate value and the amount of data copied per second

User-specified *syncrate* attribute value	**Data copied per second**
1 - 10	128 KB
11 - 20	256 KB
21 - 30	512 KB
31 - 40	1 MB
41 - 50	2 MB
51 - 60	4 MB
61 - 70	8 MB

Table 158. Relationship between the syncrate value and the amount of data copied per second (continued)	
User-specified *syncrate* attribute value	**Data copied per second**
71 - 80	16 MB
81 - 90	32 MB
91 - 100	64 MB
101 - 110	128 MB
111 - 120	256 MB
121 - 130	512 MB
131 - 140	1 GB
141 - 150	2 GB

The **addvdiskcopy** command can run on a Safeguarded backup volume with the following limitations:

- The specified pool must be a Safeguarded backup location.
- The volume will be thin-provisioned if the Safeguarded backup location is a standard pool no matter what parameters the user specifies.
- The volume will be compressed and deduplicated if the Safeguarded backup location is a Data Reduction Pool no matter what parameters the user specifies.
- The volume will have **autoexpand** set to yes whether specified or not. The user may specify **-autoexpand** parameter if they so choose but doing so will not make any difference.
- The user can specify **-rsize** and the specified value honored otherwise a default **-rsize** value will be chosen.
  - The default rsize value is 2% of the volume size with a minimum rsize of 1 GiB and maximum rsize 100 GiB.
  - The user cannot specify a **-rsize** of 0.
- The user may optionally specify the following parameters:
  - -grain_size

    Important: If you specify the **grain_size** parameter, you must also specify the **-rsize** parameter.
  - -syncrate
  - -mirrorwritepriority
  - -autodelete
  - -unit
- Other than the previously listed parameters, all other parameters are not permitted. The default value of the -warning parameter is 0, which specifies that warnings are disabled.

Note: You cannot use the **addvdiskcopy** command to create copy of image-mode volumes if the specified volume is assigned with a replication policy. You must consider any associated change volumes when you add or remove volume copies. You must maintain the change volume so it is consistent with the volume it protects.

Failure Scenarios

The following failure scenarios can occur when the command fails:

Scenario 1:

The command fails because of insufficient resources if the I/O group contains any of the following elements:

- At least one 8 GiB node.
- At least one thin-provisioned or compressed volume in a data reduction pool.
- A FlashCopy bitmap size for I/O group exceeds 1.5 GiB.

Scenario 2:

When a thin-provisioned or compressed volume is created within a data reduction pool, the pool must have enough capacity to create more volumes that track SCSI unmap operations from the host. If this capacity is not available, the command fails.

Scenario 3:

Volumes cannot be created in a data reduction pool if offline thin-provisioned or compressed volumes exist in a data reduction pool, either because of thin provisioning (an out-of-space or data corruption issue), or a component that is associated with the thin provisioning keeping a pool volume offline.

An invocation example

```
addvdiskcopy -mdiskgrp 0 -easytier off vdisk8
```

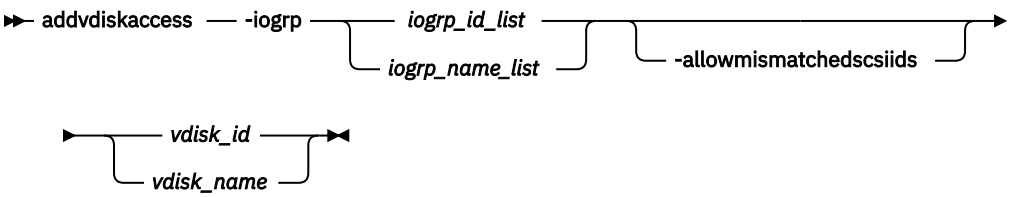
The resulting output:

```
Vdisk [8] copy [1] successfully created
```

addvdiskaccess

Use the **addvdiskaccess** command to add one or more I/O groups to the set of I/O groups in which a volume can be made accessible to hosts.

Syntax



Parameters

-iogrp *iogrp_id_list | iogrp_name_list*
(Required) Specifies a list of I/O groups to add to the I/O group volume access set.

-allowmismatchedscsiids
(Optional) When you add new access I/O groups, this parameter allows the SCSI LUN IDs to be different from the currently assigned SCSI LUN IDs in existing access I/O groups.

vdisk_id | vdisk_name
(Required) Specifies the volume to which to add access through the specified I/O groups.

Description

If an I/O group is already a member of the access set, no error is generated and no action is taken for that I/O group. All host mappings for the volume are added to the I/O groups in the list. The **-force** option is not required to extend additional mappings to other I/O groups.

When an I/O group is added to the access set, access is granted to the volume from the hosts that are mapped to the volume from the nodes in the I/O group. If the volume is mapped twice, it is also mapped twice through all additional I/O groups.

You can add I/O groups to the volume access list if the groups are mapped to iSCSI hosts. This means that iSCSI hosts can access volumes that are accessible through single or multiple I/O groups.

Two mappings are created if a host is mapped to a volume with two I/O groups. Hosts are limited to 512 host-to-volume mappings, which means a host can be mapped to the following volumes:

- 512 volumes in a single I/O group
- 256 volumes across two I/O groups
- 64 volumes across four I/O groups

Restriction:

The **-addvdiskaccess** command fails if one or more of the following issues exist:

- The host (for which the volume has a host mapping) is not associated with an I/O group in the list.
- The host volume mapping limit is exceeded.
- The number of extra mappings that are added exceeds the clustered system limit for host volume mappings.
- The same SCSI LUN ID is not available in the new access I/O group.

The system assigns the same SCSI LUN ID allocated to the volume in other access I/O groups if it is available in the new access I/O groups. Using the **-allowmismatchedscsiids** parameter allows the system to allocate non-identical SCSI LUN IDs in the access I/O group. The lowest value available in each access I/O group is used. It might not be the same in all access I/O groups. Ensure that the host supports this configuration. This parameter is not supported for NVMe hosts. When a volume is mapped to a host that uses the NVMe protocol, the volume namespace ID (NSID) is assigned by the system. When you add access to a volume that already has different SCSI LUN IDs in its existing access I/O group, you must enable the **-allowmismatchedscsiids** parameter, or the command fails.

An invocation example

This example adds I/O group 2 to the volume access set for DB_Volume:

```
addvdiskaccess -iogrp 2 DB_Volume
```

The resulting output:

```
No feedback
```

analyzevdisk

Use the **analyzevdisk** command to queue or cancel volume analysis.

Syntax



Parameters

-cancel
(Optional) Cancels an ongoing compression estimation.

vdisk_id | vdisk_name
(Required) Specifies the volume ID or name to queue for analysis.

Description

This command queues or cancels volume analysis. The order is based on the *vdisk_id* value.

Important: You cannot specify `analyzevdisk -cancel` for a volume that is not currently being analyzed (or is queued for analysis).

You can schedule an offline volume for analysis (no error message is displayed). The volume remains scheduled until it is back online and is analyzed according to its *vdisk id* value.

A concise invocation example to queue vdisk 0 for analysis

```
analyzevdisk 0
```

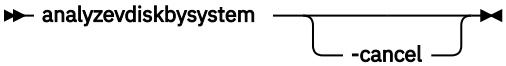
The detailed resulting output:

```
No feedback
```

analyzevdiskbysystem

Use the **analyzevdiskbysystem** command to schedule all existing volumes in system for available capacity analysis.

Syntax



Parameters

-cancel
(Optional) Cancels a scheduled or pending compression estimate.

Description

This command schedules all existing volumes in system for available capacity analysis.

Volumes that are created after you specify the command are not evaluated. Use **analyzevdisk** to evaluate specific volumes.

A concise invocation example to queue vdisk 0 for analysis

```
analyzevdiskbysystem
```

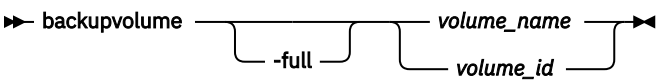
The detailed resulting output:

```
No feedback
```

backupvolume

Use the **backupvolume** command to create a volume snapshot.

Syntax



Parameters

-full
(Optional) Specifies that the snapshot generation for the volume should be a full snapshot.

volume_name | volume_id
(Required) Specifies the volume name or ID for the volume being backed up. The value for the volume name must be an alphanumeric string and the value for the volume ID must be a number.

Description

This command creates a volume snapshot.

The command completes processing when the volume snapshot is taken, and the snapshot is asynchronously transferred to the cloud system.

Important: If a volume belongs to a volume group, you must issue the **backupvolumegroup** command instead of the **backupvolume** command.

An invocation example for creating a full snapshot generation that has existing volume snapshots on the cloud system

```
backupvolume -full vdisk7
```

The resulting output:

```
No feedback
```

backupvolumegroup

Use the **backupvolumegroup** command to create a new snapshot for all of the volumes in a volume group.

Syntax



Parameters

-full
(Optional) Specifies a full backup for volume group members.

volumegroup_name | volumegroup_id
(Required) Specifies a volume group ID or name for the volume to back up. The value must be a number for the volume group ID and an alphanumeric string for the volume group name.

Description

This command creates a new snapshot for all of the volumes in a volume group.

The command completes as soon as a snapshot of the volume group is taken. The backup is asynchronously transferred to the cloud. If any of the volume members have a backup or restore in progress, any new volume group backup cannot be taken. The volume backup on each volume member needs to be enabled by using the **chvdisk** command to enable the volume backup for entire volume group.

An invocation example

To create a backup of a volume group for the first time:

```
backupvolumegroup volgroup1
```

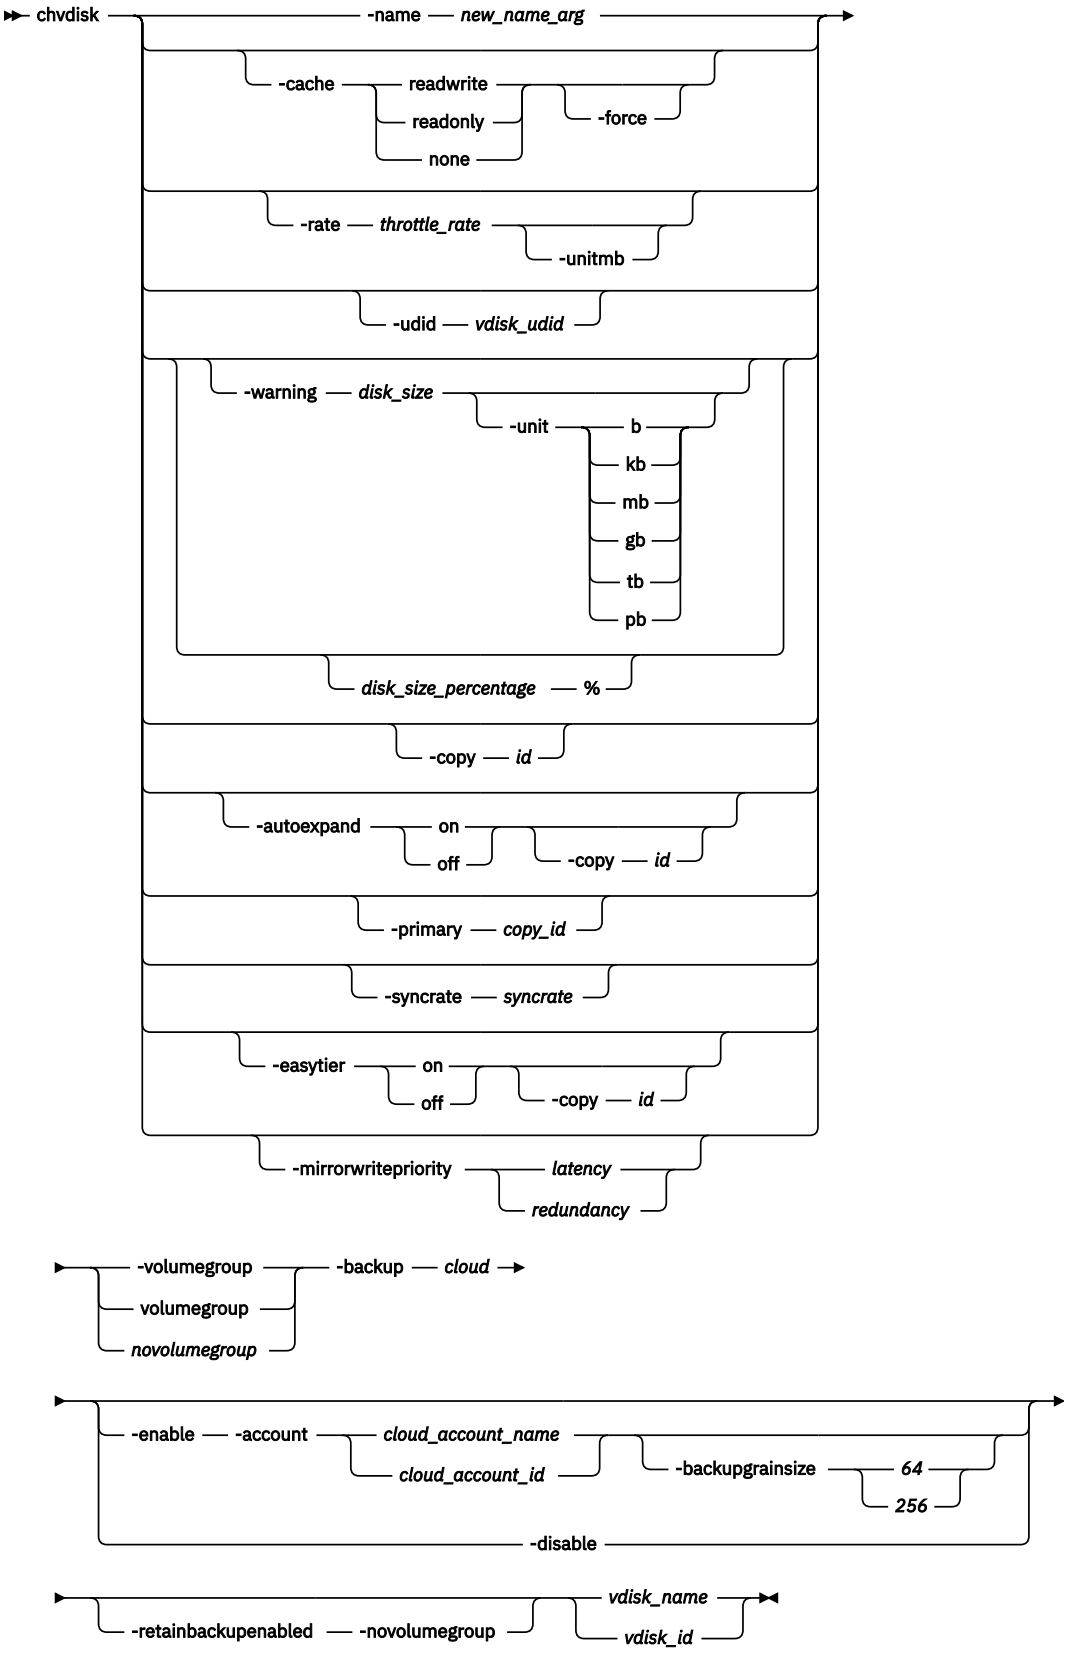
The resulting output:

```
No feedback
```

chvdisk

Use the **chvdisk** command to modify the properties of a volume, such as the disk name, I/O governing rate, or unit number. You can also change IBM Easy Tier settings.

Syntax


chvdisk
-name
new_name_arg
-cache
readwrite
readonly
none
-force
-rate
throttle_rate
-unitmb
-udid
vdisk_udid
-warning
disk_size
-unit
b
kb
mb
gb
tb
pb
disk_size_percentage
%
-copy
id
-autoexpand
on
off
-copy
id
-primary
copy_id
-syncrate
syncrate
-easytier
on
off
-copy
id
-mirrorwritepriority
latency
redundancy
-volumegroup
volumegroup
novolumegroup
-backup
cloud
-enable
-account
cloud_account_name
cloud_account_id
-backupgrainsize
64
256
-disable
-retainbackupenabled
-novolumegroup
vdisk_name
vdisk_id

Parameters

-name *new_name_arg*
(Optional) Specifies a new name to assign to the volume. You cannot use this parameter with the **-rate** or **-udid** parameters. This parameter is required if you do not use the **-rate** or **-udid** parameters.

Note: Do not use this parameter with file system volumes.

-cache readwrite | readonly | none
(Optional) Specifies the caching options for the volume. Valid entries are:

- Use *readwrite* to enable the cache for the volume.
- Use *readonly* to disable write caching and allow read caching for a volume.
- Use *none* to disable the cache mode for the volume.

The default is *readwrite*.

-force
(Optional) The **force** parameter can be used only for changing the caching mode.

Use the **force** parameter with the **cache** parameter to specify that you want the system to change the cache mode of the volume even if the I/O group is offline. This option overrides the cache flush mechanism.

Note:

1. Only a user with the Security Admin role or the superuser is permitted to use **-force** to change the I/O group or preferred node when the specified volume is a Safeguarded source of a Safeguarded snapshot.
2. The **-volumegroup** parameter is not allowed, if the volume group is associated with a snapshot policy and the volume does not meet the criteria to create a snapshot. A snapshot cannot be created if a volume:
   - Has zero size.
   - Target of a FlashCopy map.
   - The ownership group of the pool of the volume does not match the ownership group of the volume group.



Attention: If the **force** parameter is used for changing the caching mode, the contents of the cache are discarded and the volume might be corrupted by the loss of the cached data. This corruption might occur if the system is able to destage all write data from the cache or not. Use the **force** parameter with caution.

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

-rate *throttle_rate* **-unitmb**
(Optional) Specifies the I/O governing rate for the volume, which caps the amount of I/O that is accepted. The default *throttle_rate* units are I/Os. By default the *throttle_rate* is disabled. To change the *throttle_rate* units to megabits per second (MBps), specify the **-unitmb** parameter. The governing rate for a volume can be specified by I/Os or by MBps, but not both. However, you can set the rate to I/Os for some volumes and to MBps for others. When the input/output operations per second (IOPS) limit is configured on a volume, and it is smaller than 100 IOPS, the throttling logic rounds it to 100 IOPS. Even if throttle is set to a value smaller than 100 IOPs, the actual throttling occurs at 100 IOPs.

Note: To disable the throttling on a specific volume, set the `throttle_rate` value to zero.

You cannot use this parameter with the **-name** or **-udid** parameters.

-udid *vdisk_udid*
(Optional) Specifies the unit number (**-udid**) for the disk. The *vdisk_udid* is an identifier that is required to support OpenVMS hosts; no other systems use this parameter. Valid options are a decimal number in the range 0 - 32,767, or a hexadecimal number from 0 to 0x7FFF. A hexadecimal number

must be preceded by 0x (for example, 0x1234). If you do not use the **-udid** parameter, the default **-udid** is 0.

You cannot use this parameter with the **-name** parameters.

-warning *disk_size | disk_size_percentage%*
(Optional) Generates a warning when the used disk capacity on the thin-provisioned copy first exceeds the specified threshold. You can specify a *disk_size* integer, which defaults to MBs unless the **-unit** parameter is specified; or you can specify a *disk_size***%**, which is a percentage of the volume size. To disable warnings, specify 0 or 0%.

-unit b | kb | mb | gb | tb | pb
(Optional) Specifies the data units to use for the **-warning***disk_size* parameter. The default unit value is MB.

-autoexpand on | off
(Optional) Specifies whether thin-provisioned volume copies automatically expand their real capacities by allocating new extents from their storage pool. To use this parameter, the volume must be thin-provisioned.

-copy *id*
(Optional) Specifies the copy to apply the changes to. You must specify this parameter with the **-autoexpand** or **-warning** parameter. The **-copy** parameter is required if the specified volume is mirrored and only one volume copy is thin-provisioned. If both copies are thin-provisioned and the **-copy** parameter is not specified, the specified **-autoexpand** or **-warning** parameter is set on both copies.

-primary *copy_id*
(Optional) Specifies the primary copy. Changing the primary copy takes effect only when the new primary copy is online and synchronized. If the new primary is online and synchronized when the command is issued, the change takes effect immediately. You cannot change the volume's primary copy if that primary copy has its autodelete flag is set to yes (on).

Important: You cannot use this parameter with a volume that is fast formatting.

-syncrate *syncrate*
(Optional) Specifies the copy synchronization rate. A value of zero (0) prevents synchronization. The default value is 50. See Table 159 on page 685 for the supported **-syncrate** values and their corresponding rates. Use this parameter to alter the rate at which the standard-provisioned volume or mirrored volume format before synchronization.

-easytier on | off
(Optional) Enables or disables the IBM Easy Tier function.



Attention: Thin-provisioned and compressed volumes in a data reduction pool always have IBM Easy Tier on, regardless of the pool setting.

-mirrorwritepriority *latency | redundancy*
(Optional) Specifies how to configure the mirror write algorithm priority. A change to the mirror write priority is reflected in the volume's view immediately and in the volume's behavior after all prior input and output (I/O) completes.

1. Choosing *latency* means a copy that is slow to respond to a write I/O becomes unsynchronized, and the write I/O completes if the other copy successfully writes the data.
2. Choosing *redundancy* means a copy that is slow to respond to a write I/O synchronizes completion of the write I/O with the completion of the slower I/O to maintain synchronization.

-volumegroup *volumegroup_name | volumegroup_id*
(Optional) Specifies a new volume group for a volume. This parameter is mutually exclusive with **-novolumegroup**.

Note: If a vdisk is added to a volume group having a snapshot policy with cloud backup schedule, the cloud backup is enabled for the **vdisk**.

-novolumegroup
(Optional) Specifies that a volume does not belong in any volume group. This parameter is mutually exclusive with **-volumegroup**.

Note: If a vdisk is removed from a volume group having a snapshot policy with cloud backup schedule, the cloud backup is disabled for the **vdisk**.

-backup *cloud*
(Optional) Specifies the cloud snapshot type to enable or disable. The value must be `cloud`.

-enable
(Optional) Enables the backup or snapshot type that is specified with the **-backup** parameter.

-disable
(Optional) Disables the backup or snapshot type that is specified with the **-backup** parameter.

-account *cloud_account_id | cloud_account_name*
(Optional) Specifies the cloud account to use for the volume. You must specify **-enable** with this parameter.

-backupgrainsize *64 | 256*
(Optional) Specifies the grain size (in KB) for volume mappings. The values are 64 and 256. You must specify **-enable**to use this parameter.

You can enable a volume for a cloud snapshot with one account. You cannot enable cloud backup on a volume for a second time on the same or different cloud account.

You cannot turn off the cloud snapshot function if a snapshot in progress. Any snapshot that is in progress must complete or be canceled.

vdisk_name | vdisk_id
(Required) Specifies the volume to modify, either by ID or by name.

-retainbackupenabled
(Optional) Retains cloud backup enabled for the **vdisk**.

Case 1
If a vdisk is removed from a volume group without using **-retainbackupenabled**, the cloud backup is disabled for that vdisk.

Case 2
If a vdisk is removed from the volume group using **-retainbackupenabled**, the cloud backup remains enabled for that vdisk.

Description

The **chvdisk** command modifies a single property of a volume. To change the volume name and modify the synchronization rate, for example, you must issue the command twice. If the volume is offline, use **recovervdisk** command to recover the volume and bring it back online.

Important: To change the caching I/O group for a volume or preferred node, use the **movevdisk** command.

A thin-provisioned or compressed copy that is in a data reduction storage pool must enable **-autoexpand**. If a volume contains a copy that is in a data reduction storage pool, the cache mode must be set to `readwrite`.

A thin-provisioned or compressed copy that is in a data reduction storage pool cannot have a warning threshold set. To change the warning threshold, you must specify the **-copy**.

You can specify a new name or label. Then, you can use the new name to refer to the volume.

You can set a limit on the number of I/O transactions that is accepted for this volume. It is set in terms of I/Os per second or MBs per second. By default, no I/O governing rate is set when a volume is created.



Attention: All capacities, including changes, must be in multiples of 512 bytes. An error occurs if you specify a capacity that is not a multiple of 512, which can happen only when byte units are used. The default capacity is in MB.

When the volume is created, no throttling is applied to it. Use the **-rate** parameter to change it. To change the volume back to an unthrottled state, specify 0 (zero) with the **-rate** parameter.

For thin-provisioned and compressed volume copies in data reduction pools, the Easy Tier status is derived from the data reduction pool as the data is managed by a central data disk. Therefore, the Easy Tier mode cannot be turned off on these volume types. The thin-provisioned and compressed volumes in a data reduction pool always have Easy Tier on, regardless of the pool setting. The Easy Tier setting is a combination of pool and volume setting, as detailed in a table in the help for **mkvdisk**. Leaving easy tier always on for the volume allows the pool setting to be the one that switches it on or off. For standard-provisioned volumes in a data reduction pool, Easy Tier can be turned both on and off.

The rate at which the volume copies resynchronize after loss of synchronization can be specified by using the **-syncrate** parameter. This table provides the relationship of the *syncrate* value to the data copied per second.

Note: These settings also affect the initial rate of formatting.

Table 159. Relationship between the syncrate value and the amount of data copied per second

User-specified *syncrate* attribute value	**Data copied per second**
1 - 10	128 KB
11 - 20	256 KB
21 - 30	512 KB
31 - 40	1 MB
41 - 50	2 MB
51 - 60	4 MB
61 - 70	8 MB
71 - 80	16 MB
81 - 90	32 MB
91 - 100	64 MB
101 - 110	128 MB
111 - 120	256 MB
121 - 130	512 MB
131 - 140	1 GB
141 - 150	2 GB

Description

The **chvdisk** command modifies a single property of a volume. To change the volume name and modify the synchronization rate, for example, you must issue the command twice.

Note: If the volume is offline, use one of the **recovervdisk** commands to recover the volume and bring it back online.

Important: To change the caching I/O group for a volume, use the **movevdisk** command.

You can specify a new name or label. Then, you can use the new name to refer to the volume.

An invocation example

```
chvdisk -rate 2040 1
```

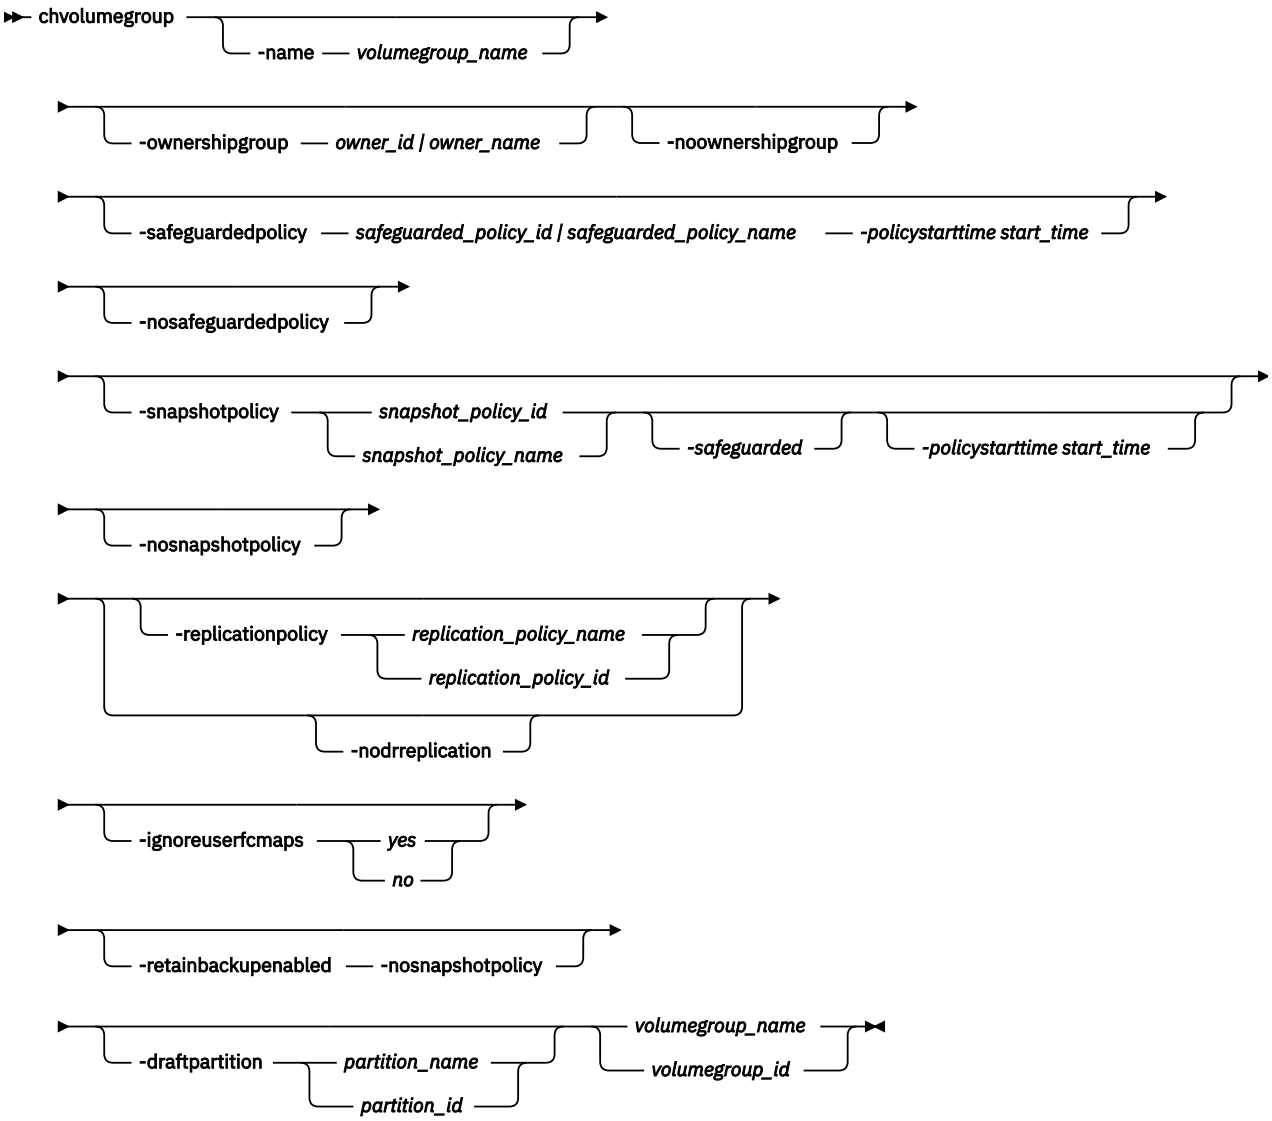
The following output is displayed:

```
No feedback
```

chvolumegroup

Use the **chvolumegroup** command to change volume group properties.

Syntax



Parameters

-name *volumegroup_name*
(Optional) Specifies a new volume group name. The value must be an alphanumeric string.

-ownershipgroup *owner_id | owner_name*
(Optional) Specifies the name or ID of the ownership group to which the object is being added.

-noownershipgroup
(Optional) If specified, the object is removed from the ownership group to which it belongs.

volumegroup_name | volumegroup_id
(Required) Specifies the volume group name or group ID for the volume you want to modify. The value must be a number for the volume group ID, and an alphanumeric string for the volume group name.

-safeguardedpolicy ***safeguarded_policy_id | safeguarded_policy_name***
(Optional) Specifies the name or ID of the assigned Safeguarded backup policy to a volume group.

-policystarttime ***start_time***
(Optional) Specifies when the first copy is taken using the specified policy. Indicates the time in *YYMMDDHHMM* format.

-nosafeguardedpolicy
(Optional) Specifies the removal of the Safeguarded policy that is assigned to a volume group.

-snapshotpolicy ***snapshot_policy_id | snapshot_policy_name***
(Optional) Specifies the snapshot policy to assign to the volume group.

Note: If snapshot policy has cloud backup schedule, cloud backup is enabled for all **vdisks** in the **volumegroup**.

-nosnapshotpolicy
(Optional) Specifies the removal of snapshot policy assigned to a volume group.

Note: If snapshot policy has cloud backup schedule, cloud backup is disabled for all **vdisks** in the **volumegroup**. Additionally, **-retainbackupenabled** allows user to retain cloud backup enabled for the **vdisks**.

-replicationpolicy ***replication_policy_name | replication_policy_id***
(Optional) Specifies the replication policy to assign to the volume group.

When the replication policy has an async-dr topology then the volume group must be in a storage partition and the async-dr replication policy must be associated with that storage partition.

If the volumegroup is in a 2-site-ha storage partition and an associated async-dr replication policy is specified then the volumegroup will be HA with async DR replication configured. If the storage partition is not HA then the volume group will be created with only async DR replication configured.

Note:

- This command can run whether the remote system in the replication policy is connected or disconnected. If the command succeeds, the policy is assigned immediately on the local system. The policy is assigned asynchronously on the remote system when connectivity is restored.
- You cannot assign a replication policy to the volume group if the group contains any volumes with the following characteristics:
  - Volumes are in a different I/O group than the I/O group that is specified in the replication policy.
  - Volumes have disabled cache.
  - Volumes are image mode volumes.
  - Volumes are used in Transparent Cloud Tiering (TCT).

-nodrreplication
(Optional) Removes DR replication for the volume group.

Note:

- This parameter is mutually exclusive with **-replicationpolicy**.
- DR replication can be removed from a volume group configured with a 2-site-async-dr or async-dr replication policy.

-safeguarded
(Optional) Specifies that the snapshot policy creates Safeguarded snapshots of source volumes in the volume group.

Note: The **-safeguarded** option is not allowed while applying the cloud backup policy to a volume group.

-ignoreuserfcmaps
(Optional) Allow a user to create snapshots through the scheduler or manually with the **addsnapshot** command. It can only be done if a volume in the volume group is used as a source of a user FlashCopy mapping. Values can be *yes* or *no*.

-draftpartition *partition_name | partition_id*
Attempt to associate this volume group with the specified draft partition.

-retainbackupenabled
(Optional) Retains cloud backup enabled for the **vdisk**.

Case 1
If a vdisk is removed from a volume group without using **-retainbackupenabled**, the cloud backup is disabled for that vdisk.

Case 2
If a vdisk is removed from the volume group using **-retainbackupenabled**, the cloud backup remains enabled for that vdisk.

Creating snapshot policies

The value in the **-policystarttime** parameter overrides the **-backup_start_time** parameter in the snapshot policy. It is used when the first snapshot is created. The frequency information in the snapshot policy determines the timing of all the subsequent snapshots. If the value is in the past, the frequency information in the snapshot policy is used to determine the appropriate time for the first snapshot to be taken.

Creating Safeguarded backup policies

The value in **-policystarttime** overrides the **backup_start_time** in the Safeguarded policy. It is used for determining when the first Safeguarded backup is created. After this the frequency information in the Safeguarded policy determines the timing of all subsequent Safeguarded backups. If the value is in the past, the frequency information in the Safeguarded policy is used to determine the appropriate time for the first Safeguarded backups to be taken.

- The value of **-policystarttime** should be specified in the local time of the system in *YYMMDDHHMM* format. Seconds are not to be specified.
- If **-safeguardedpolicy** is specified and **-policystartime** is not specified any existing value of **policy_start_time** is cleared and the **backup_start_time** from the Safeguarded policy is used to determine when Safeguarded backup copies are created.
- If **-nosafeguardedpolicy** is specified, any existing value of **-policy_start_time** is cleared.
- The **-safeguardedpolicy** parameter is mutually exclusive with **-nosafeguardedpolicy**, **-ownershipgroup** and **-noownershipgroup**.
- The **-nosafeguardedpolicy** parameter is mutually exclusive with **-ownershipgroup**.
- If the volume group is associated with an ownership group, the **-safeguardedpolicy** parameter fails.
- If the volume group is associated with a Safeguarded policy, the **-ownershipgroup** parameter fails.

A volume group can be associated with a Safeguarded backup policy only if all volumes it contains meet all the requirements to be a Safeguarded source volume. If any volume does not meet those requirements, the **-safeguardedpolicy** parameter fails. For more information, refer to Snapshots.

Note: You cannot modify the name of the volume group if the volume group is in a storage partition with an HA replication policy assigned. You cannot assign or remove a replication policy from a volume group in a storage partition. Volume groups in a storage partition can not be in an ownership group. These restrictions apply to a single-system partition and a partition with an HA replication policy assigned.

Note: Policing on other commands ensures that a replication policy with async-dr topology can only be created and associated with a partition if that partition has DR links configured. There is no requirement to police this as part of the **chvolumegroup** command when using **-replicationpolicy** parameter.

An invocation example

```
chvolumegroup -name newname1 1
```

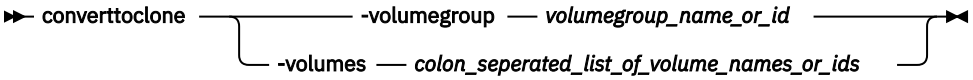
The resulting output:

```
No feedback
```

converttoclone

Use the **converttoclone** command to convert a thinclone volume to a clone volume.

Syntax



Parameters

-volumegroup ***volumegroup_name_or_id***
(Required) The volumegroup within which any thinclone volumes must be converted to clone. If the specified volume group has no thinclone volumes the command will return an error.

-volumes ***colon_seperated_list_of_volume_names_or_ids***
(Optional) A list of thinclone volumes that must be converted to clone. If the user supplied list of volumes contains any volumes that are not thinclones then the command will fail with an error.

An invocation example

This invocation will convert any thinclone volumes within volume group 0 to clones

```
converttoclone -volumegroup 0
```

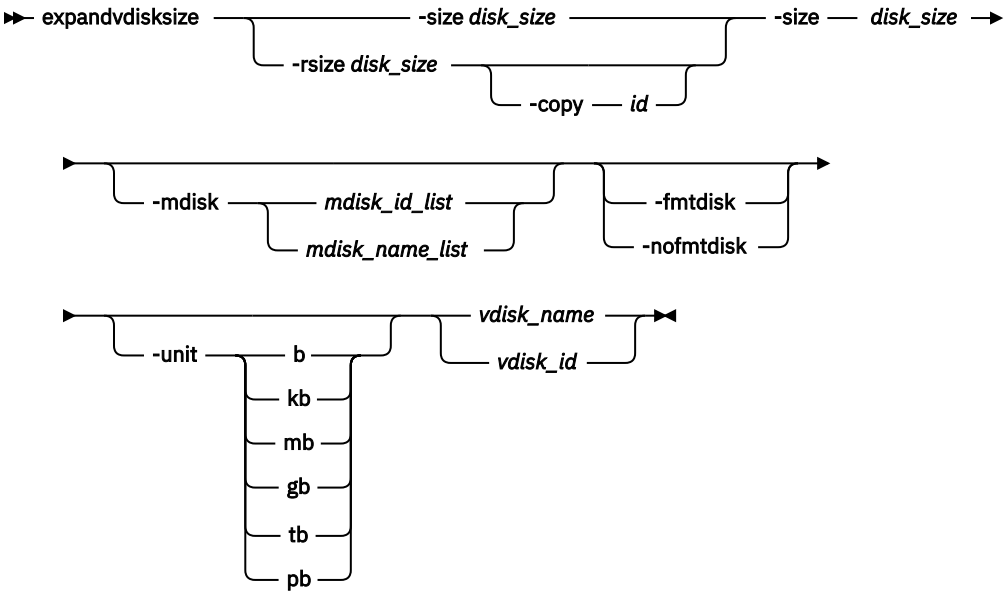
The resulting output:

```
No feedback.
```

expandvdisksize

Use the **expandvdisksize** command to expand the provisioned capacity of a volume by a specified amount.

Syntax



Parameters

-size *disk_size*
(Optional) Specifies the capacity by which the volume is expanded. If you specify a capacity that is not a multiple of 1 MB, it is aligned automatically. You must specify either **-size** or **-rsize**. The default disk_size unit is megabytes (MB). Use the **-unit** parameter to specify a different unit.

-rsize *disk_size*
(Optional) Specifies the capacity by which to increase the real size of a thin-provisioned volume. Specify the *disk_size* value by using an integer. Specify the unit for a *disk_size* integer by using the **-unit** parameter; the default unit is megabytes (MB). The **-rsize** value can be greater than, equal to, or less than the size of the volume. You cannot specify the **-rsize** parameter with the **-size** parameter. You must specify either **-size** or **-rsize**. This option applies only to volume copies in a standard pool.

-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 volume copy. The **-copy** parameter is required if the specified volume is mirrored and only one copy is thin-provisioned. If the volume is mirrored, both copies are thin-provisioned and **-copy** is not specified, both copies are modified by the same amount.

-mdisk *mdisk_id_list | mdisk_name_list*
(Optional) Specifies the list of one or more MDisks to be used as the stripe set. The extents that expand the volume come from the specified list of MDisks. All MDisks in the list must be part of the same storage pool. The **-mdisk** parameter cannot be used if the specified volume is mirrored or if the volume is thin-provisioned.

-fmtdisk
(Optional) Specifies that the volume be formatted before use. This flag formats the new extents that have been added to the volume as a result of the **expandvdisksize** command. The **expandvdisksize** command completes asynchronously if you use this parameter. Fast formatting is on by default for standard-provisioned volumes.

-nofmtdisk
(Optional) Disables fast format for the newly expanded region of the standard-provisioned volume. This option is ignored when expanding a mirrored volume that has a fully allocated copy. The newly expanded region is always fast formatted.

-unit b | kb | mb | gb | tb | pb
(Optional) Specifies the *disk_size* unit for the **-size** or **-rsize** parameter. The default value is megabytes (MB).

vdisk_name | vdisk_id
(Required) Specifies the volume to modify, either by ID or by name.

Description

Use the **expandvdisksize** command to expand the writable capacity that is allocated to a particular volume by the specified amount. You can use expandvdisksize while the volume is fast formatting and while the volume is synchronizing.

The command can also be used to expand the provisioned capacity of a thin-provisioned volume without altering the writable capacity that is assigned to the volume. To change the capacity of a non-thin-provisioned volume, or the provisioned capacity of a thin-provisioned volume, use the **-size** parameter. To change the real capacity of a thin-provisioned volume, use the **-rsize** parameter.

Note:

1. You cannot resize (expand) an image mode volume.
2. You cannot resize (expand) volume if that volume is being fast formatted.
3. You cannot resize (expand) a volume if cloud snapshot is enabled on that volume.
4. You cannot specify `expandvdisksize -rsize` to expand (resize) a thin or compressed volume copy that is in a data reduction pool.
5. You cannot specify `expandvdisksize -mdisk` to resize (expand) a volume when a volume is being migrated.
6. You cannot resize (expand) a volume that is replicated, or in a highly available storage partition.

An invocation example to increase the capacity of vdisk1 by 2048 bytes using extents from two MDisks (and to format the new part of the volume)

```
expandvdisksize -size 2048 -unit b -mdisk mdisk0:mdisk1 -fmtdisk vdisk1
```

The resulting output:

```
No feedback
```

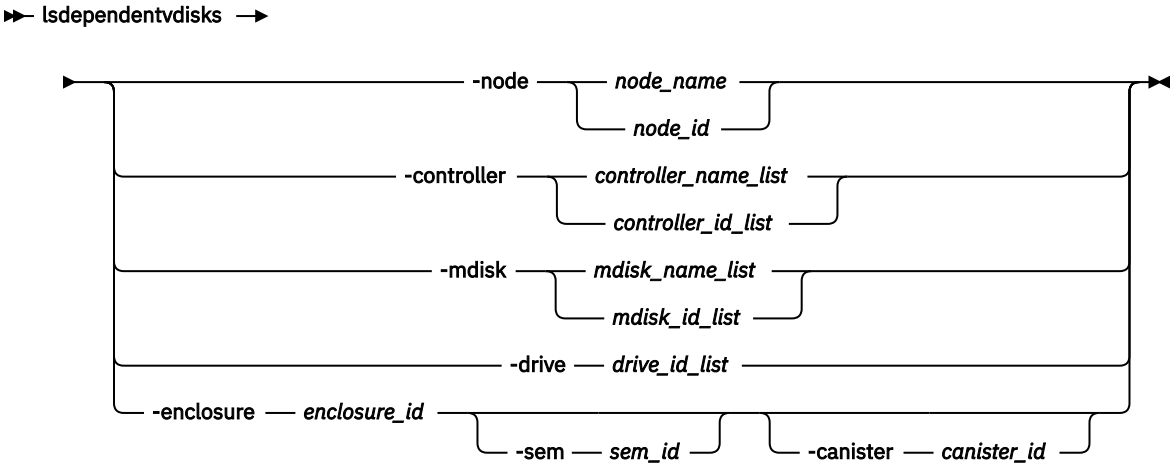
expandvolume - Deprecated

This command is deprecated as Hyperswap configurations are no longer supported. Use **expandvdisksize** to expand volumes.

lsdependentvdisks

Use the **lsdependentvdisks** command to view which volumes go offline if you remove a specific piece of hardware from the system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-node *node_name | node_id*
(Optional) Specifies the node for which volume dependency is required.

-controller *controller_name_list | controller_id_list*
(Optional) Specifies the controllers for which volume dependency is required.

-mdisk *mkdisk_name_list | mkdisk_id_list*
(Optional) Specifies the MDisks for which volume dependency is required.

-drive
(Optional) Specifies the drives for which volume dependency is required. There is a maximum of 128 entries.

-enclosure *enclosure_id*
(Optional) Specifies the enclosure for which volume dependency is required.

-sem *sem_id*
(Optional) Specifies the volumes that depend on secondary expander module availability. The possible values are 1 and 2.

Note: If you do not specify **-sem** this command shows the volumes that depend on the enclosure.

-canister *canister_id*
(Optional) Specifies an enclosure canister. The possible values are 1 and 2.

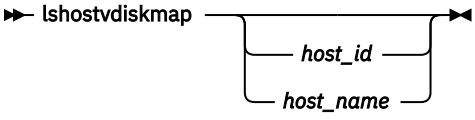
Description

Use this command to view which volumes go offline if you remove a specific piece of hardware from the system. Use this command to determine which volumes are affected before undergoing maintenance.

lshostvdiskmap

Use the **lshostvdiskmap** command to display a list of volumes that are mapped to a host. These volumes are the volumes that are recognized by the specified host.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

host_id | host_name
(Optional) Specifies the host in terms of its ID or name. The command displays a list of all the volumes that are mapped to the specified host and the Small Computer System Interface (SCSI) ID by which they are mapped. If you do not specify a host ID or name , the command displays a list of all recognized volume mappings.

Description

This command displays a list of volume IDs and names. These volumes are the volumes that are mapped to the specified host; that is, they are visible to the specified host. The SCSI_id is also displayed. This SCSI_id is the Logical Unit Number (LUN) by which the volume is recognized by the host in the `Host, Bus, Target, LUN` scheme. The SCSI_id or LUN is used for determining the path of the volume, but the volume carries an EUI-64 identifier that is tagged as the vdisk_UID.

Shared mappings, or volumes that are mapped to a host cluster visible by the **lshostclustervolumemap** command, are still listed in the output of **lshostvdiskmap** with *shared* as **mapping_type**. Shared mappings have a consistent SCSI_id across all host mappings of the same Host Cluster. To limit the number of paths, a physical host might have multiple host-objects that are defined as a subset of host ports. If you do not follow the best practices or inadvertently map the same volume to multiple sub-host objects, operating software of the host might see the same vdisk_UID with multiple paths that cover different SCSI_ids. This is functional still for some operating software, but must be avoided.

For a specific volume based on which operating system and multipath software are used, you can use different commands to determine the volume serial number, or Volume Unique Identifier (vdisk_UID). For example, in AIX with AIXPCM, **lsmpio -ql hdisk** finds the Volume Serial Number for volumes that are mapped to that operating software.

Find the host that is defined to the system that corresponds with the host that you are working with.

1. The worldwide port names (WWPNs) are an attribute of the host bus adapter (HBA). You can find the WWPNs by looking at the device definitions that are stored by your operating system. For example, on AIX® they are in the Object Data Manager (ODM), in Windows® they are in the Device Manager details for the specified HBA.
2. Verify which host is defined to the system that these ports belong to. The ports are stored as part of the detailed view, so you must list each host in turn by issuing the following command:

   `lshost` *`host_name | host_id`*

   where *host_name* | *host_id* is the name or ID of the host. Check for matching WWPNs.

Note: Consider this when you name your hosts. For example, if the actual host is called `orange`, also name the host that is defined to the system `orange`.

The command returns the following values:

id
Indicates the host ID in the output for **lshostvdiskmap**.

name
Indicates the host name in the output for **lshostvdiskmap**.

SCSI_id
Specifies the SCSI ID for SCSI mappings.

host_cluster_id
Indicates the unique ID for a host system.

host_cluster_name
Indicates the unique name for a host system.

vdisk_id
Specifies the UID of the volume.

vdisk_name
Indicates the name of the volume.

vdisk_UID
Indicates the UID of the volume.

IO_group_id
Indicates the ID of the input/output (I/O) group in which the host volume mapping exists.

IO_group_name
Specifies the name of I/O group in which the host volume mapping exists.

NVMe_NSID
Indicates NVMe namespace ID.

lsmetadatavdisk

Use the **lsmetadatavdisk** command to display the information for metadata volumes on the system.

Syntax



Description

This command displays the information for metadata volumes on the system.

This table provides the attribute values that can be displayed as output view data.

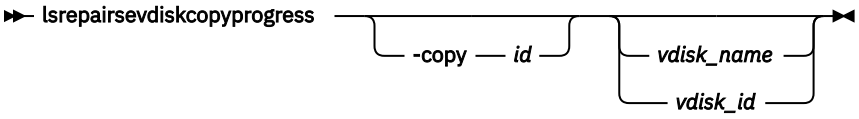
*Table 160. **lsmetadatavdisk** output*

Attribute	Description
vdisk_id	Indicates the ID of the metadata volume.
vdisk_name	Indicates the name of the metadata volume.
status	Indicates the running status of the metadata volume.
version	Indicates the version of the metadata volume. The values can be 1 or 2.

lsrepairsevdiskcopyprogress

The **lsrepairsevdiskcopyprogress** command lists the repair progress for thin-provisioned volume copies or compressed volume copies.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-copy *id*
(Optional) Lists the repair progress for the specified copy.

vdisk_name* | *vdisk_id
(Optional) Specifies the volume name or ID to list repair progress for. You must specify this parameter last on the command line. If you do not enter this parameter, the command lists progress for all thin-provisioned copies in the clustered system.

Description

The **lsrepairsevdiskcopyprogress** command lists the repair progress for thin-provisioned or compressed copies for the specified volume. If you do not specify a volume, the command lists the repair progress for all thin-provisioned or compressed copies in the clustered system.

Remember: Run this command after you specify the **repairsevdiskcopy** command, which you must run as required by the fix procedures or by your product support information.

The command returns values for the following volume copy attributes:

task
Specifies the active task.

- `repairing` indicates repair of a thin-provisioned volume copy
- `compressed_repairing` indicates repair of a compressed volume copy.

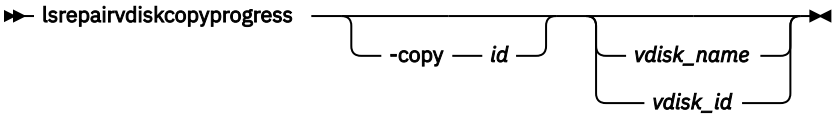
progress
Specifies the task completion percentage.

estimated_completion_time
Specifies the expected duration of the task in the format YYMMDDHHMMSS (or blank if the estimated completion is unknown).

lsrepairvdiskcopyprogress

Use the **lsrepairvdiskcopyprogress** command to display the progress of volume repairs and validations.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-copy *id*
(Optional) Specifies the volume copy ID to list repair progress for. If you do not specify this parameter, progress is displayed for all copies.

vdisk_name* | *vdisk_id
(Optional) Specifies the volume name or ID to list repair progress for. You must specify this parameter last on the command line.

Description

The **lsrepairvdiskcopyprogress** command displays the progress of repairs and validations that are made to mirrored volumes. Use this command to track progress after you run the **repairvdiskcopy** command. You can specify a volume copy by using the **-copy** parameter. To display the volumes that have two or more copies with an active task, specify the command with no parameters; it is not possible to have only one volume copy with an active task.

The command displays progress for the following types of volume copies:

- All volume copies display the same task; validate, medium, or resync, depending on the specified parameter.
- All volume copies display the same percentage and estimated completion time.
- If specified, non-mirrored volumes are displayed as a single copy with a blank task; they are not displayed in the full concise view.
- Once a task completes, the task is blank for all copies.
- If the task is blank, the percentage and the completion time are also blank.

The command returns values for the following volume repair attributes:

vdisk_id
Indicates the volume ID.

vdisk_name
Indicates the volume name.

copy_id
Indicates the system-assigned identifier for the volume copy.

task
Indicates the active task. The values can be `repairing` or `compressed_repairing`.

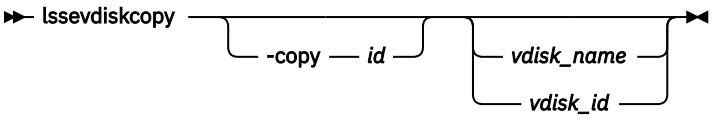
progress
Indicates the task completion percentage. This value is 0 when task is in `compressed_repairing` state.

estimated_completion_time
Indicates the expected time (duration) the task completion time. The value is in the YYMMDDHHMMSS format, and is blank if the duration is not known.

lssevdiskcopy

Use the **lssevdiskcopy** command to list the thin-provisioned copies of the specified volumes.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-bytes
(Optional) Displays all capacities as bytes. Capacity values that are displayed in units other than bytes might be rounded.

-copy *id*
(Optional) Specifies the volume copy to list thin-provisioned copies for. You must specify a *vdisk_name* | *vdisk_id* value with this parameter.

vdisk_name* | *vdisk_id
(Optional) Specifies the volume name or ID to list thin-provisioned copies for. You must specify this parameter last on the command line. If you do not enter this parameter, the command lists all thin-provisioned copies in the clustered system.

Description

The `lssevdiskcopy` command lists all thin-provisioned copies of the specified volume. If you do not specify a volume, the command lists all thin-provisioned volume copies in the clustered system.

The command provides a concise view of the thin-provisioned properties of the selected volume copies. Run the **lsvdiskcopy** command to see a concise view of the properties that are common to thin-provisioned and non-thin-provisioned volume copies. See the description of the **lsvdisk** command for a description of the fields that are shown in the view.

The command returns values for the following volume copy attributes:

copy_id
Indicates a system-assigned identifier for the volume copy. The value can be `0` or `1`.

status
Indicates the system status. The value can be `online` or `offline`. A copy is offline if all nodes cannot access the storage pool that contains the copy.

sync
Indicates whether the volume copy is synchronized.

auto_delete
Indicates that the primary copy is deleted after the secondary copy is synchronized. The values are `yes` or `no`.

primary
Indicates whether the volume copy is the primary copy. A volume has exactly one primary copy. The value can be `yes` or `no`.

mdiskgrp_id
Indicates the ID of the storage pool that the volume copy belongs to.

mdiskgrp_name
Indicates the name of the storage pool that the volume copy belongs to.

type
Indicates the virtualization type of the volume. The value can be `striped`, `sequential`, or `image`.

mdisk_id
Indicates the MDisk ID that is used for sequential and image mode volumes.

mdisk_name
Indicates the MDisk name that is used for sequential and image mode volumes.

fast_write_state
Indicates the cache state of the volume copy. The value can be `empty`, `not_empty`, `corrupt`, or `repairing`. The value is always empty for non-thin-provisioned copies. A cache state of `corrupt` indicates that the volume is thin-provisioned and requires repair that is initiated by a **recovervdisk** command or the **repairsevdiskcopy** command.

used_capacity
Indicates the portion of `real_capacity` that is being used to store data. For non-thin-provisioned copies, this value is the same as the volume capacity. If the volume copy is thin-provisioned, the value increases from zero to the `real_capacity` value as more of the volume is written to.

Remember: This field is blank for thin-provisioned and compressed volume copies in a data reduction pool.

real_capacity
Indicates the amount of physical storage that is allocated from a storage pool to this volume copy. If the volume copy is not thin-provisioned, the value is the same as the volume capacity. If the volume copy is thin-provisioned, the value can be different.

Remember: This field is blank for thin-provisioned and compressed volume copies in a data reduction pool.

free_capacity
Indicates the difference between the `real_capacity` and `used_capacity` values.

Remember: This field is blank for thin-provisioned and compressed volume copies in a data reduction pool.

overallocation
Expressed as a percentage, indicates the ratio of volume capacity to `real_capacity` values. This value is always `100` for non-thin-provisioned volumes.

Remember: This field is zero for thin-provisioned and compressed volume copies in a data reduction pool.

autoexpand
Indicates whether `autoexpand` is enabled on a thin-provisioned volume. The value can be on or `off`.

warning
Expressed as a percentage, for thin-provisioned volume copies only, indicates that a warning is generated when the ratio of `used_capacity` to volume capacity reaches the specified level.

Remember: This field is zero for thin-provisioned and compressed volume copies in a data reduction pool.

grainsize
For thin-provisioned volume copies, indicates the grain size that is chosen for the volume copy when it was created.

Remember: This field is blank for compressed volume copies in regular storage pools.

se_copy
Indicates whether the copy is thin-provisioned. The value can be `yes` or no.

easy_tier
Indicates whether Easy Tier is permitted to manage the pool.

Note:

1. If easy_tier is on, then easy_tier_status can take on any value.
2. If easy_tier is `off`, then easy_tier_status is `measured` or `inactive`.

easy_tier_status
Indicates which Easy Tier functions are active for the volume copy:

- `active` indicates that a pool is being managed by Easy Tier to provide tier management performance-based pool balancing. For example, extents of this volume copy can be moved for performance (automatic data placement).
- `inactive` indicates that no Easy Tier function is active.
- `measured` indicates that statistics are being gathered for this volume copy, but no extents are moved.
- `balanced` indicates that a pool is being managed by Easy Tier to provide performance-based pool balancing (for example, extents can be moved).

tier
Indicates which tier information is being reported:

- `tier0_flash`
- `tier1_flash`
- `tier_enterprise`
- `tier_nearline`
- `tier_scm`

tier_capacity
Indicates the total MDisk capacity that is assigned to the volume in the tier.

Note: For thin-provisioned copies, the capacity by tier is the real capacity.

compressed_copy
Indicates whether the volume copy is compressed.

uncompressed_used_capacity
For compressed volume copies, indicates the amount of data that is written to the volume copy before compression. This field is blank for volume copies in a data reduction storage pool.

used_capacity_before_reduction
Indicates the total amount of data that is written to a thin-provisioned or compressed volume copy in a data reduction storage pool before data reduction occurs. This field is blank for fully allocated volume copies and volume copies not in a data reduction pool.

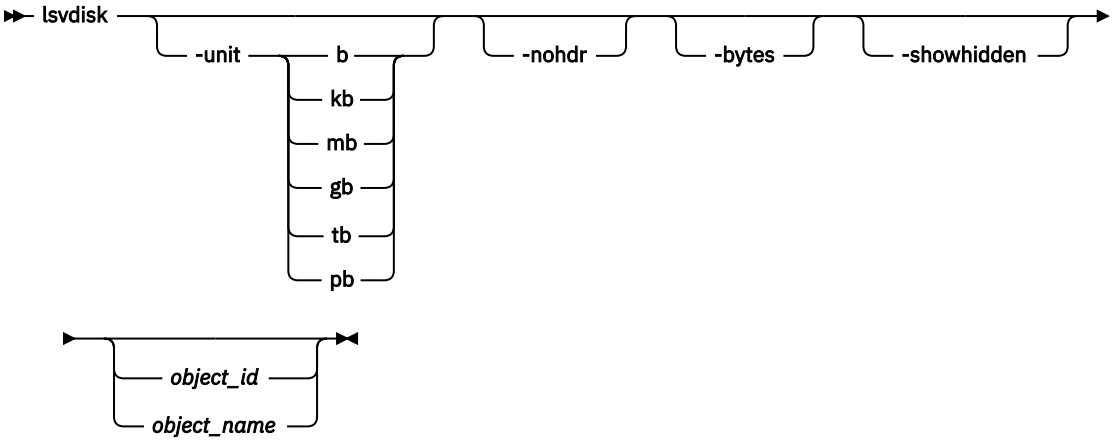
parent_mdisk_grp_id
Indicates the physical storage pool ID that the volume extents are allocated from. This value is a numeric string (in the range 0 - 127 characters) or blank.

parent_mdisk_grp_name
Indicates the physical storage pool name that the volume extents are allocated from. This value is an alphanumeric string (in the range 1 - 63 characters) or blank.

lsvdisk

Use the **lsvdisk** command to display a concise list or a detailed view of volumes that are recognized by the system.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-unit b | kb | mb | gb | tb | pb
(Optional) Specifies the data units for the **-filtervalue** parameter.

-bytes
(Optional) Displays all capacities as bytes. Capacity values that are displayed in units other than bytes might be rounded. When you filter a capacity value, use a unit of bytes, **-unit b**, for exact filtering. For thin-provisioned copies, the capacity by tier contains the real capacities.

-showhidden
(Optional) Specifies that hidden volumes should be displayed. This includes snapshots, system-owned volumes, and volumes that are in the deleting state.

object_id | object_name
(Optional) Specifies the name or ID of an object. When you use this parameter, the detailed view of the specific object is returned and any value that is specified by the **-filtervalue** parameter is ignored. If you do not specify the *object_id | object_name* parameter, the concise view of all objects that match the filtering requirements that are specified by the **-filtervalue** parameter are displayed.

Description

This command displays a concise list or a detailed view of attributes for all volumes and volume copies in the system.

The command returns values for the following volume attributes:

name
Indicates the name of the volume.

id
Indicates the ID of the volume.

capacity
Indicates the provisioned capacity of the volume that is the size of the volume as seen by the host.

Internal_capacity
Indicates the capacity of the volume for which resources are allocated. This value can be greater than the provisioned capacity, if an attempt to expand the size of the volume fails whilst it was replicated.

vdisk_UID
Indicates the UID of the volume. If a volume is mapped to a host using the NVMe protocol, then the `vdisk_UID` value shows the Namespace Globally Unique Identifier (NGUID) for the volume.

NVMe_GUID
Indicates the NVMe UID of the volume.

NVMe_nsid
Indicates NVMe namespace ID.

status
The value can be `online`, `offline`, `degraded`, or `deleting`.

udid
Indicates the unit number for the volume. Only OpenVMS hosts require a unit number.

IO_group_id
Indicates the I/O group (ID) that the volume belongs to.

IO_group_name
Indicates the I/O group (name) that the volume belongs to.

mdisk_grp_id
Indicates the ID of the storage pool that the volume belongs to. If the volume has more than one copy, these fields display `many`.

mdisk_grp_name
Indicates the name of the storage pool that the volume belongs to. If the volume has more than one copy, these fields display `many`.

type
Indicates the virtualization type of the volume. The value can be `striped`, `seq`, `image` or `many`. The value `many` indicates that the volume has more than one copy, which can have different virtualization types.

formatted
Indicates whether the volume was formatted when it was created. The value can be `yes` or `no`.

formatting
Indicates whether the volume is formatting. The value can be `yes` or `no`.

mdisk_id
Indicates the MDisk ID that is used for sequential and image mode volumes. If the volume has more than one copy, these fields display `many`.

mdisk_name
Indicates the MDisk name that is used for sequential and image mode volumes. If the volume has more than one copy, these fields display `many`.

FC_id
Indicates the ID of the FlashCopy mapping that the volume belongs to. The value `many` indicates that the volume belongs to more than one FlashCopy mapping.

FC_name
Indicates the name of the FlashCopy mapping that the volume belongs to. The value `many` indicates that the volume belongs to more than one FlashCopy mapping.

RC_id (Deprecated)

RC_name (Deprecated)

preferred_node_id
Indicates the ID of the preferred node for processing I/O to the volume.

fast_write_state
Indicates the cache state for the volume or volume copy. The value can be `empty`, `not_empty`, `corrupt`, or `repairing`. A cache state of `corrupt` indicates that the volume or volume copy requires repairing or recovery by using either the **recovervdisk** or **repairvdiskcopy** command.

cache
Indicates the cache mode of the volume. The value can be `readonly`, `readwrite`, or `none`.

fc_map_count
Indicates the number of FlashCopy mappings that the volume belongs to.

sync_rate
Indicates the rate for synchronization for mirrored copies.

copy_count
Indicates the number of volume copies that exist for this volume.

se_copy_count
Indicates the number of thin-provisioned copies.

Remember: This value represents only thin-provisioned copies and is not used for compressed volume copies.

filesystem
This field is always blank.

mirror_write_priority
Indicates the mirror write priority algorithm that is used if the volume is mirrored.

RC_change (Deprecated)

compressed_copy_count
Indicates the number of compressed volume copies.

access_IO_group_count
Indicates the number of I/O groups in the volume access set.

last_access_time
Indicates the time (YYMMDDHHMMSS) the volume last received I/O from any mapped host.

parent_mdisk_grp_id
Indicates the physical storage pool ID that the volume extents are allocated from.

parent_mdisk_grp_name
Indicates the physical storage pool name that the volume extents are allocated from.

owner_type
Indicates the type of owning component or object (such as vVols). The values are:

- `host_integration_metadata`
- `vvol`
- `accesspoint`
- `cloud_backup`
- `safeguarded_copy`
- `none`

The value is none if no owner is assigned.

owner_id
Indicates an identification number for the owning object. This value must be a numeric character unless there is no owning object (in which case it is blank).

owner_name
Indicates the name for owning object that owns this volume. The value must be a set of up to 63 alphanumeric characters, but is blank if there is no owning object.

encrypt
Indicates whether all copies of a volume reside in MDisk groups (storage pools) that are reported as encrypting. This means that either one of the following things apply:

- The copies reside in a storage pool that has an encryption key.
- All MDisks in the group are self-encrypting or encrypted for RAID.

The values are `yes` and `no`.

volume_id
Indicates the volume ID (for a high availability volume).

volume_name
Indicates the volume name (for a high availability volume).

function (Deprecated)

throttle_id
Indicates the ID for the throttle object. The field is blank if no throttle is configured.

throttle_name
Indicates the name of the throttle object. The value is blank if no throttle is configured.

IOPs_limit
Indicates the IOPs limit that is configured for the volume. The value is blank if no limit is configured.

bandwidth_limit_MB
Indicates the bandwidth limit configured (in MB) for the volume. The value is blank if no limit is configured.

volume_group_id
Indicates the ID of the volume group this volume belongs to. This value is blank if the volume is not in a volume group.

volume_group_name
Indicates the name of the volume group this volume belongs to. This value is blank if the volume is not in a volume group.

cloud_backup_enabled
Indicates whether the cloud snapshot feature is enabled for the specified volume. The values are `yes` or `no`.

cloud_account_id
Indicates the cloud account ID.

cloud_account_name
Indicates the cloud account name.

backup_status
Indicates whether a new cloud snapshot can be started. If a backup is in progress, the status of the backup operation is given. The values are:

- `off`
- `ready`
- `copying`
- `copying_error`
- `not_ready`

last_backup_time
Indicates the time (`YYMMDDHHMMSS`) of the most recent cloud snapshot for the volume.

restore_status
Indicates whether a restore can be performed for the volume. If a restore is in progress, the status of the restore operation is given. The values are:

- none
- available
- restoring
- restoring_error
- committed
- committing
- committing_error

backup_grain_size
Indicates the grain size for the FlashCopy mappings that are used for the cloud snapshot function. The value is blank if cloud snapshot is not enabled.

Note: This sizing does not reflect the size of the grains that are stored in the cloud (which are fixed at 256 KB).

changevolume1_id
Indicates the ID of the associated change volume for the policy-based replication.

preferred_node_name
Indicates the name of the preferred node for processing I/O to the volume.

is_snapshot
Indicates the volume is a snapshot. Snapshots are hidden by default.

snapshot_count
Indicates the number of cloud snapshots associated with this volume.

volume_type
Indicates the type of volume.

source_volume_id
Indicates the ID of the source volume.

source_volume_name
Indicates the name of the source volume.

source_snapshot
Indicates the name of the source snapshot.

source_snapshot_timestamp
Indicates the timestamp of source snapshot (YYMMDDHHMMSS).

protection_provisioned_capacity
Indicates the amount of capacity provisioned to store point-in-time copies of this volume. This includes snapshots and consistency protection for replication.

protection_written_capacity
Indicates the amount of capacity written to store point-in-time copies of this volume before any data reduction. This includes snapshots and consistency protection for replication.

display_name
Indicates the field to show a recognizable name for a Virtual Volume (vVol). This value is only used for vVols and is automatically generated by the system. It must be an alphanumeric string (in the range 1 - 63 characters) or blank.

location1_volume_id
Indicates the ID of the volume in the first position of the topology diagram.

location2_volume_id
Indicates the ID of the volume in the second position of the topology diagram.

location3_volume_id
Indicates the ID of the volume in the third position of the topology diagram.

replication_mode
Indicates the replication mode of the volume on the local system. The field is blank if policy-based replication is not configured.

is_safeguarded_snapshot
Indicates the volume is a Safeguarded snapshot. Safeguarded snapshots are hidden by default.

safeguarded_snapshot_count
Indicates the number of Safeguarded snapshots associated with the volume.

The command returns values for the following volume copy attributes:

copy_id
Indicates a system-assigned identifier for the volume copy. The value can be `0` or `1`.

status
Indicates the status. The value can be `online`, `offline`, `degraded`, or `deleting`.

sync
Indicates whether the volume copy is synchronized.

auto_delete
Indicates that the primary copy is deleted after the secondary copy is synchronized. The values are `yes` or `no`.

primary
Indicates whether the volume copy is the primary copy. A volume has exactly one primary copy. The value can be `yes` or `no`.

mdisk_grp_id
Indicates the ID of the storage pool that the volume copy belongs to.

mdisk_grp_name
Indicates the name of the storage pool that the volume copy belongs to.

type
Indicates the virtualization type of the volume. The value can be `striped`, `seq`, or `image`.

mdisk_id
Indicates the MDisk ID that is used for sequential and image mode volumes.

mdisk_name
Indicates the MDisk name that is used for sequential and image mode volumes.

fast_write_state
Indicates the cache state for the volume or volume copy. The value can be `empty`, `not_empty`, `corrupt`, or `repairing`. A cache state of `corrupt` indicates that the volume or volume copy requires repairing or recovery by using either the **`recovervdisk`** or **`repairvdiskcopy`** command.

used_capacity
Indicates the portion of `real_capacity` that is being used to store data. For non-thin-provisioned copies, this value is the same as the volume capacity. If the volume copy is thin-provisioned, the value increases from zero to the `real_capacity` value as more of the volume is written to. This field is blank for volume copies that are thin-provisioned or compressed volume copies in a data reduction pool.

real_capacity
Indicates the amount of physical storage that is allocated from a storage pool to this volume copy. If the volume copy is not thin-provisioned, the value is the same as the volume capacity. If the volume copy is thin-provisioned, the value can be different. This field is blank for volume copies that are thin-provisioned or compressed volume copies in a data reduction pool.

free_capacity
Indicates the difference between the `real_capacity` and `used_capacity` values. This field is blank for storage pools that are not thin-provisioned or compressed volume copies in a data reduction pool.

overallocation
Expressed as a percentage of the volume capacity, indicates the ratio of capacity to `real_capacity` values. This value is always `100` for non-thin-provisioned or compressed volumes. This field is blank for volume copies that are thin-provisioned or compressed volume copies in a data reduction pool.

autoexpand
Indicates whether **autoexpand** is enabled on a thin-provisioned volume. The value can be on or off.

warning
Expressed as a percentage of the volume capacity, this value indicates a warning for thin-provisioned or compressed volume copies. A warning is generated when the ratio of used_capacity to volume capacity reaches the specified level. This field is blank for volume copies that are thin-provisioned or compressed volume copies in a data reduction pool.

grainsize
For thin-provisioned volume copies, indicates the grain size that is chosen for the volume copy when it was created.

se_copy
Indicates whether the copy is thin-provisioned.

easy_tier
This value is set by the user and indicates whether Easy Tier is permitted to manage the pool.

easy_tier_status
Indicates which Easy Tier functions are active.

tier
The tier information being reported.

tier_capacity
Indicates the total MDisk capacity that is assigned to the volume in the tier.

Note: For thin-provisioned copies, the capacity by tier is the real capacity.

compressed_copy
Indicates whether the volume copy is compressed.

uncompressed_used_capacity
For compressed volume copies, indicates the amount of data written to the volume copy before compression. This field is blank for volume copies in a data reduction storage pool.

parent_mdisk_grp_id
Indicates the physical storage pool ID that the volume extents are allocated from.

parent_mdisk_grp_name
Indicates the physical storage pool name that the volume extents are allocated from.

encrypt
Indicates whether all copies of a volume reside in MDisk groups (storage pools) that are reported as encrypting. This means that either one of the following things apply:

- The copies reside in a storage pool that has an encryption key.
- All MDisks in the group are self-encrypting or encrypted for RAID.

The values are yes and no.

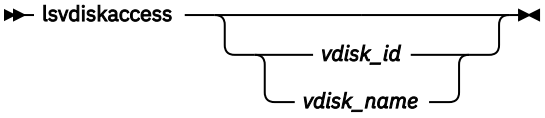
used_capacity_before_reduction
Indicates the total amount of data written to a thin-provisioned or compressed volume copy in a data reduction storage pool before data reduction occurs. This field is blank for fully allocated volume copies and volume copies not in a data reduction pool.

restore_in_progress
Indicates whether the volume is currently restored. The values are yes or no.

lsvdiskaccess

Use the **lsvdiskaccess** command to display a list of all input/output (I/O) groups in the volume access set.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

vdisk_id | vdisk_name
(Optional) Specifies the volume for which to list access I/O groups.

Description

The **lsvdiskaccess** command lists the I/O groups in a volume access set. An accessible volume in an I/O group does not indicate that the volume is mapped to any hosts. There is a detailed view and concise view, but the detailed view does not contain more information than the concise view.

This command returns values for the following volume attributes:

VDisk_id
Identifies the volume ID.

VDisk_name
Identifies the volume name.

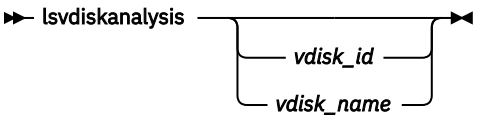
IO_group_id
Identifies an I/O group ID in the volume access set.

IO_group_name
Identifies an I/O group name in the volume access set.

lsvdiskanalysis

Use the **lsvdiskanalysis** command to display information for thin provisioning and compression estimation analysis-report for a single volume or multiple volumes.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

vdisk_id | vdisk_name
Specifies the volume (by ID or name) to display compression estimation status for.

Description

This command displays information for thin provisioning and compression estimation analysis-report for a single volume or multiple volumes.

This table provides the attribute values that can be displayed as output view data.

Table 161. Output from the ***lsvdiskanalysis*** *command*

Attribute	Description
id	Indicates the ID (by number) of the object.
state	Indicates one of the following values: • idle Indicates that the volume was never analyzed. • scheduled Indicates that the volume is scheduled for analysis (the analysis starts based on ascending volume IDs). • active Indicates that the volume is being analyzed. • estimated Indicates that the volume was analyzed and the analysis results reflect the estimated savings from thin provisioning and compression. • sparse Indicates that the volume was analyzed but not enough samples of nonzero data were found. • cancelling Indicates that the analysis is ongoing and a request exists to cancel the analysis, but the analysis is not yet ended. • automatic Indicates that the volume is being analyzed due to automatic scheduling.
started_time	Indicates the date and time that the analysis started, which helps determine how long an estimate takes. The value must be an alphanumeric data string or be empty if analysis never started.
analysis_time	Indicates the date and time that the analysis ends to help determine how current the results are. If you cancel the analysis, the value for the time is invalid (0, and it is not displayed). While active, time is invalid as well and does not reflect expected completion time. The value must be an alphanumeric data string or be empty if analysis never occurred.
capacity	Indicates the provisioned capacity (host size) of the volume.
thin_size	Indicates the estimated size of the data without zero portions (thin-provisioned size).
thin_savings	Indicates how much data saving is expected if it is a thin-provisioned volume.

*Table 161. Output from the **lsvdiskanalysis** command (continued)*	
Attribute	**Description**
thin_savings_ratio	Indicates the percent of data that is saved by thin-provisioned. The number must be a percentage.
compressed_size	Indicates the estimated size of any nonzero data after compression completes.
compression_savings	Indicates how much data to expect to save if the volume is a compressed volume.
compression_savings_ratio	Indicates the amount of data that is saved by compression. The number must be a percentage.
total_savings	Indicates how much data to expect to save by converting a volume to a compressed volume.
total_savings_ratio	Indicates the amount of data that is saved by compression is based on the overall volume capacity, which includes the thin nature of compressed volumes. The number must be a percentage.
margin_of_error	Indicates the accuracy estimation. The number is a percentage.
file_system	Indicates the file system(s) of the volume

lsvdiskanalysisprogress

Use the **lsvdiskanalysisprogress** command to display information about the space analysis progress for an entire system.

Syntax

Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

Description

This command displays information about the space analysis progress for an entire system.

This table provides the attribute values that can be displayed as output view data.

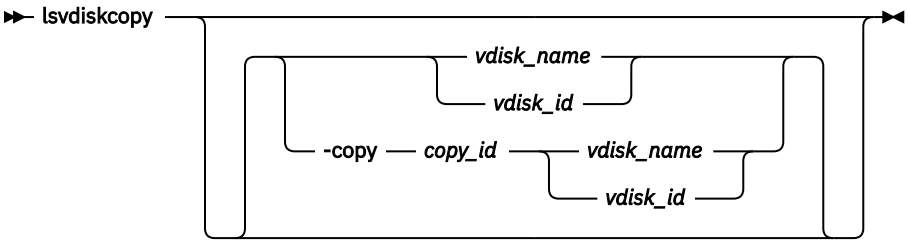
*Table 162. **lsvdiskanalysisprogress** output*	
Attribute	**Description**
vdisk_count	Indicates the number of volumes on this system.
pending_analysis	Indicates the number of volumes that belong to this system and: • Scheduled a available capacity analysis • Have an active available capacity analysis • Are canceling available capacity analysis

Table 162. **lsvdiskanalysisprogress** output (continued)	
Attribute	**Description**
estimated_completion_time	Indicates the estimated time at which analysis is expected to end. This value is calculated based on extrapolation from the actual analysis duration. Estimated completion time does not consider volumes that are offline and displays estimated completion time as if the volumes are online.
last_cycle_start_time	Indicates the start time of the last automatic space analysis time interval. The value is a number (in seconds) or blank if no cycle has previously run.
automatic_pending_analysis	Indicates the number of volumes that belong to this system and are scheduled for automatic space analysis for the current 12-hour time interval. The value is `0` if automatic scheduling is disabled.
automatic_estimated_completion_time	Indicates the automatic analysis estimated time (in seconds) that the analysis is expected to end. This value is calculated based on extrapolation from the actual analysis duration. The value is blank if automatic scheduling is disabled.

lsvdiskcopy

Use the **lsvdiskcopy** command to list volume copy information.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-bytes
(Optional) Displays all capacities as bytes. Capacity values that are displayed in units other than bytes might be rounded.

-copy ***copy_id***
(Optional) Specifies the volume copy to list information for. You must specify a *vdisk_name* | *vdisk_id* value with this parameter.

vdisk_name* | *vdisk_id
(Optional) Specifies the volume to list copy information for. You must specify this parameter last on the command line. If you specify a *vdisk_name* | *vdisk_id* value only, all copies for the volume are listed.

Description

The **lsvdiskcopy** command lists information for volume copies. If you specify the command with no parameters, all volumes and copies in the clustered system are listed.

The command returns values for the following volume copy attributes:

copy_id
Specifies a system-assigned identifier for the volume copy. The value can be 0 or 1.

status
Indicates the status. The value can be online, offline, or deleting. A copy is offline if all nodes cannot access the storage pool that contains the copy.

sync
Indicates whether the volume copy is synchronized.

auto_delete
Indicates that the primary copy is deleted after the secondary copy is synchronized. The values are yes or no.

primary
Indicates whether the volume copy is the primary copy. A volume has exactly one primary copy. The value can be yes or no.

mdiskgrp_id
Indicates the ID of the storage pool that the volume copy belongs to.

mdiskgrp_name
Indicates the name of the storage pool that the volume copy belongs to.

type
Indicates the virtualization type of the volume. The value can be striped, sequential, or image.

mdisk_id
Indicates the MDisk ID that is used for sequential and image mode volumes.

mdisk_name
Indicates the MDisk name that is used for sequential and image mode volumes.

fast_write_state
Indicates the cache state of the volume copy. The value can be empty, not_empty, corrupt, or repairing. The value is always empty for non-thin-provisioned copies. A cache state of corrupt indicates that the volume is thin-provisioned and requires repair that is initiated by a **recovervdisk** command or the **repairsevdiskcopy** command.

used_capacity
Indicates the portion of real_capacity that is being used to store data. For non-thin-provisioned copies, this value is the same as the volume capacity. If the volume copy is thin-provisioned, the value increases from zero to the real_capacity value as more of the volume is written to.

Remember:

- This value is the same as the volume capacity value for standard-provisioned copies.
- This field is blank for volume copies that are thin-provisioned or compressed volume copies in a data reduction pool.

real_capacity
Indicates the amount of physical storage that is allocated from a storage pool to this volume copy. If the volume copy is not thin-provisioned, the value is the same as the volume capacity. If the volume copy is thin-provisioned, the value can be different.

Remember:

- This value is the same as the volume capacity value for standard-provisioned copies.
- This field is blank for volume copies that are thin-provisioned or compressed volume copies in a data reduction pool.

free_capacity
Indicates the difference between the `real_capacity` and `used_capacity` values.

Remember:

- This value is `0` for standard-provisioned copies.
- This field is blank for volume copies that are thin-provisioned or compressed volume copies in a data reduction pool.

Remember: This value is zero for standard-provisioned copies.

overallocation
Expressed as a percentage, indicates the ratio of volume capacity to `real_capacity` values. This value is always `100` for non-thin-provisioned volumes.

Remember: This field is blank for storage pools that are not thin-provisioned or compressed volume copies in a data reduction pool.

autoexpand
Indicates whether `autoexpand` is enabled on a thin-provisioned volume. The value can be on or `off`.

warning
Expressed as a percentage of the volume capacity for thin-provisioned or compressed volume copies, indicates that a warning is generated when the ratio of `used_capacity` to volume capacity reaches the specified level.

Remember: This field is blank for storage pools that are not thin-provisioned or compressed volume copies in a data reduction pool.

grainsize
For thin-provisioned volume copies, indicates the grain size that is chosen for the volume copy when it was created.

Remember: This value is always blank for compressed volume copies in regular storage pools.

se_copy
Specifies whether the copy is thin-provisioned.

easy_tier
Indicates whether Easy Tier is permitted to manage the pool.

Remember:

1. If easy_tier is on, then easy_tier_status can take on any value.
2. If easy_tier is `off`, then easy_tier_status is `measured` or `inactive`.

easy_tier_status
Indicates which Easy Tier functions are active for the volume copy:

- `active` indicates that a pool is being managed by Easy Tier to provide tier management performance-based pool balancing. For example, extents of this volume copy can be moved for performance (automatic data placement).
- `inactive` indicates that no Easy Tier function is active.
- `balanced` indicates that a pool is being managed by Easy Tier to provide performance-based pool balancing (for example, extents can be moved).
- `measured` indicates that statistics are being gathered for this volume copy, but no extents are moved.

tier
Indicates which tier information is being reported:

- `tier0_flash`
- `tier1_flash`
- `tier_enterprise`

- `tier_nearline`
- tier_scm

tier_capacity
Indicates the total MDisk capacity that is assigned to the volume in the tier.

Important: For thin-provisioned copies, the capacity by tier is the real capacity.

Important: By design, **`tier_capacity`** reports blank for thin-provisioned and compressed copies in data reduction pools.

compressed_copy
Indicates whether the volume copy is compressed.

uncompressed_used_capacity
For compressed volume copies, indicates the amount of data written to the volume copy before compression. This field is blank for volume copies in a data reduction storage pool.

used_capacity_before_reduction
Indicates the total amount of data written to a thin-provisioned or compressed volume copy in a data reduction storage pool before data reduction occurs. This field is blank for standard-provisioned volume copies and volume copies not in a data reduction pool.

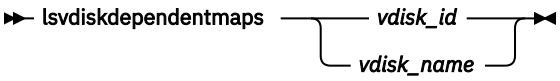
parent_mdisk_grp_id
Indicates the physical storage pool ID that the volume extents are allocated from. This value is a numeric string (in the range 0 - 127 characters) or blank.

parent_mdisk_grp_name
Indicates the physical storage pool name that the volume extents are allocated from. This value is an alphanumeric string (in the range 1 - 63 characters) or blank.

lsvdiskdependentmaps

Use the **`lsvdiskdependentmaps`** command to display all FlashCopy mappings with target volumes that are dependent upon data that is held on the specified volume.

Syntax



Parameters

`-delim` ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

`-nohdr`
(Optional) Suppress the headings in the output.

vdisk_id | vdisk_name
Specifies the name or ID of a volume.

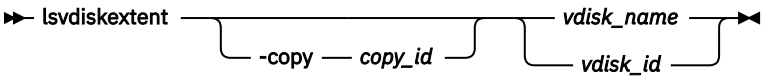
Description

The **`lsvdiskdependentmaps`** command displays FlashCopy mappings that have target volumes that are dependent upon data that is held on the specified *vdisk_id | vdisk_name*. This data can be used to determine whether a FlashCopy mapping can be prepared. Issue the command for the target volume *vdisk_id | vdisk_name* of the FlashCopy mapping to be prepared. If no FlashCopy mappings are returned, the FlashCopy mapping can be prepared. Any FlashCopy mappings that are returned in the list must be stopped or be in the `idle_or_copied` state before the new FlashCopy mapping can be prepared.

lsvdiskextent

Use the **lsvdiskextent** command to list the MDisk extents that are provided for the specified volumes.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-copy *copy_id*
(Optional) Displays a list of MDisks that are members of the specified volume copy.

vdisk_name | vdisk_id
(Required) Specifies one or more volume IDs or names.

Description

Note: You cannot specify this command for a thin-provisioned or compressed volume or a volume copy that is in a data reduction storage pool.

The **lsvdiskextent** command displays a list of MDisk IDs and the number of extents that each MDisk provides to the specified volumes.

Each volume is constructed from one or more MDisks. To determine the relationship between a volume and its MDisks, issue the following command:

```
lsvdiskmember vdisk_name | vdisk_id
```

where *vdisk_name* | *vdisk_id* is the name or ID of the volume. This command displays a list of MDisk IDs that make up the volume.

To determine the number of extents that are provided by each MDisk, issue the following command:

```
lsvdiskextent vdisk_name | vdisk_id
```

where *vdisk_name* | *vdisk_id* is the name or ID of the volume. This command displays a table of MDisk IDs and the corresponding number of extents that each MDisk provides as storage for the specified volume.

To determine the relationship between MDisks and volumes, issue the following command for each MDisk:

```
lsmdiskmember mdisk_name | mdisk_id
```

where *mdisk_name* | *mdisk_id* is the name or ID of the MDisk. This command displays a list of IDs that corresponds to the volumes that are using this MDisk.

To determine the relationship between MDisks and volumes, and the number of extents that are used by each volume, you must use the command-line interface. For each MDisk, issue the following command:

```
lsmdiskextent mdisk_name | mdisk_id
```

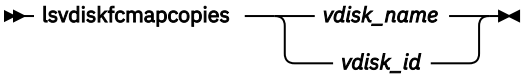
where *mdisk_name* | *mdisk_id* is the name or ID of the MDisk. This command displays a table of volume IDs and the corresponding number of extents that are used by each volume.

Note: If the MDisk specified is in a data reduction pool, the output includes all thin-provisioned and compressed volumes in the pool without displaying the number of extents in each.

lsvdiskfcmapcopies

Use the **lsvdiskfcmapcopies** command to display a list of all FlashCopy mappings with a target volume that contains a valid copy of the specified volume.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

vdisk_name* | *vdisk_id
(Required) Specifies the name or ID of the volume for which the FlashCopy mappings are displayed.

Description

This command returns a list of the FlashCopy mappings that have a target volume with a valid copy of the specified volume. The target volumes of these mappings can be considered as candidate source volumes for mappings to restore from.

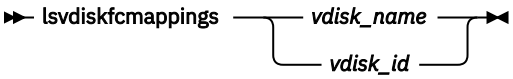
The mappings that are returned are in the copying, idle_copied, or stopping state with 100% progress.

Note: Maps that are rc_controlled are not shown in the view when this command is specified.

lsvdiskfcmappings

Use the **lsvdiskfcmappings** command to display a list of FlashCopy mappings to which the volume belongs. A volume can be part of up to 256 FlashCopy mappings.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

vdisk_name* | *vdisk_id
Specifies the name or ID of the volume with FlashCopy mappings that are listed.

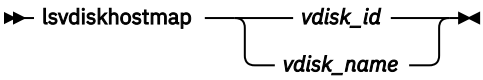
Description

The **lsvdiskfcmappings** command returns a list of all FlashCopy mappings that the volume is a member of. The list is returned in no particular order.

lsvdiskhostmap

Use the **lsvdiskhostmap** command to list the hosts to which a volume is mapped.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

vdisk_id | vdisk_name
Specifies the ID or name of the volume. The clustered system displays a list of all the hosts to which this volume is mapped and the Small Computer System Interface (SCSI) ID by which the volume is mapped.

Description

This command displays a list of host IDs and names. These hosts specify a volume that is mapped to them; that is, the volume is visible to these hosts. The SCSI LUN ID is also displayed. The SCSI LUN ID is the ID by which the volume is recognized by the host.

You can use only one I/O group (the group ID value is always 0).

Determining the hosts to which a volume is mapped:

To list the hosts that this volume is mapped to, issue the following command:

```
lsvdiskhostmap vdisk_id | vdisk_name
```

Where *vdisk_id* | *vdisk_name* is the name or ID of the volume. A list is displayed. Look for the host name or ID to determine which host this volume is mapped to. If no data is displayed, the volume is not mapped to any hosts.

The command returns the following values:

id
Specifies the ID of the volume in the output for **lsvdiskhostmap**.

name
Specifies the name of the volume in the output for **lsvdiskhostmap**.

SCSI_id
Specifies the SCSI ID.

host_id
Specifies the ID of the host.

host_name
Specifies the name of the host.

vdisk_UID
Specifies the UID of the volume.

IO_group_id
Specifies the ID of the input/output (I/O) group in which the host volume mapping exists.

IO_group_name
Specifies the name of I/O group in which the host volume mapping exists.

mapping_type
Indicates the mapping type for a host system. Values are `private` or `shared`.

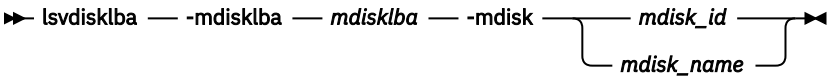
host_cluster_id
Indicates the unique ID for a host system.

host_cluster_name
Indicates the unique name for a host system.

lsvdisklba

Use the **lsvdisklba** command to list the volume and logical block address (LBA) for the specified storage pool LBA.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-mdisklba *mdisklba*
(Required) Specifies the 64-bit hexadecimal LBA on the MDisk. The LBA must be specified in hex, with a 0x prefix.

-mdisk *mdisk_id | mdisk_name*
(Required) Specifies the MDisk name or ID.

Description

The **lsvdisklba** command returns the LBA of the volume that is associated with the MDisk LBA.

If applicable, the command also lists the range of LBAs on both the volume and MDisk that are mapped in the same extent, or for thin-provisioned disks, in the same grain.

Note: If **lsvdisklba** is run during a software upgrade, the command fails and an error message is displayed.

The **vdisk_lba** field provides the corresponding LBA on the virtual capacity for the input LBA. For compressed volume copies, it is blank and the system provides the ranges of virtual LBAs that are compressed into the input LBA.

This table provides command output that depends on several variables.

*Table 163. **lsvdisklba** command output scenarios*

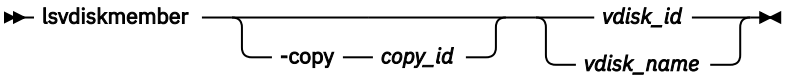
Field	Typical scenario	Quorum disk	Thin-provisioned metadata	Extent not allocated	Formatting extent	Extent allocated to thin-provisioned disk, LBA not used on thin-provisioned disk
copy_id	yes	no	yes	no	yes	yes
vdisk_id	yes	no	yes	no	yes	yes
vdisk_name	yes	no	yes	no	yes	yes
type	allocated	metadata	metadata	unallocated	formatting	unallocated

Table 163. **lsvdisklba** command output scenarios (continued)						
Field	**Typical scenario**	**Quorum disk**	**Thin-provisioned metadata**	**Extent not allocated**	**Formatting extent**	**Extent allocated to thin-provisioned disk, LBA not used on thin-provisioned disk**
vdisk_lba	yes	no	no	no	no	no
vdisk_start	yes	no	no	no	no	no
vdisk_end	yes	no	no	no	no	no
mdisk_start	yes	yes	yes	yes	yes	yes
mdisk_end	yes	yes	yes	yes	yes	yes

lsvdiskmember

Use the **lsvdiskmember** command to display a list of MDisks that are members of the specified volume.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-copy *copy_id*
(Optional) Displays a list of MDisks that are members of the specified volume copy.

vdisk_id | vdisk_name
(Required) Identifies the specific volume to query.

Description

This command displays a list of managed disks, which provide extents that make up the volume that is specified by the ID.

Every volume is constructed from one or more MDisks. At times, you might have to determine the relationship between the two objects. The following procedure determines the relationships.

If you use the **lsmdiskmember** command, the concise view displays a list of volumes. These volumes are the ones 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. If they are in an offline state, they are still displayed.

To determine the relationship between volumes and MDisks, issue the following command:

```
lsvdiskmember vdisk_id | vdisk_name
```

Where *vdisk_id* | *vdisk_name* is the name or ID of the volume. It displays a list of IDs that correspond to the MDisks that make up the volume.

To determine the relationship between volumes and MDisks, and the number of extents that are provided by each MDisk, you must use the command-line interface. Issue the following command:

```
lsvdiskextent vdisk_id | vdisk_name
```

Where *vdisk_id* | *vdisk_name* is the name or ID of the volume. It displays a table of MDisk IDs and the corresponding number of extents that each MDisk provides as storage for the specified volume.

To determine the relationship between MDisks and volumes, issue the following command:

`lsmdiskmember` *mdisk_id* | *mdisk_name*

Where *mdisk_id* | *mdisk_name* is the name or ID of the MDisk. It displays a list of IDs that correspond to the volumes that are using this MDisk.

To determine the relationship between MDisks and volumes, and the number of extents that are used by each volume, you must use the command-line interface. For a specified MDisk, issue the following command:

`lsmdiskextent` *mdisk_id* | *mdisk_name*

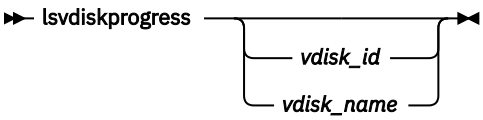
Where *mdisk_id* | *mdisk_name* is the name or ID of the MDisk. It displays a table of volume IDs and the corresponding number of extents that are used by each volume.

Note: You cannot specify this command for a thin-provisioned or compressed volume or a volume copy that is in a data reduction storage pool.

lsvdiskprogress

Use the **lsvdiskprogress** command to track the progress during new volume formatting.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

vdisk_id | vdisk_name
Specifies the volume ID or name. If you do not specify this parameter, the progress of all volumes currently being formatted is displayed.

Description

This command displays the progress of the format of a new volume as a completed percentage. If the volume has multiple copies, the command reports the average progress of the format.

The command returns values for the following volume attributes:

id
Indicates the ID of the volume that is being formatted.

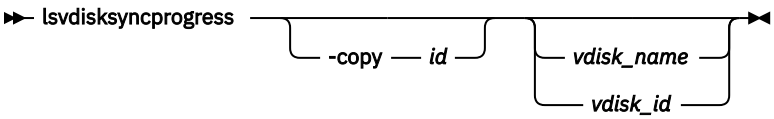
progress
Indicates the formatting progress.

estimated_completion_time
Indicates the estimated time in which the formatting operation completes. The value is in the YYMMDDHHMMSS format, and is blank if the duration is not known.

lsvdisksyncprogress

Use the `lsvdisksyncprogress` command to display the progress of volume copy synchronization.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-copy *id*
(Optional) Specifies the volume copy ID to list synchronization progress. You must also specify a *vdisk_name* | *vdisk_id* value. If you do not specify this parameter, progress is displayed for all copies.

vdisk_name* | *vdisk_id
(Optional) Specifies the volume name or ID to list synchronization progress.

Description

To display the volume copies that require synchronization, specify the command with no parameters. To display the synchronization progress for all copies of a volume, specify the command with the *vdisk_name* | *vdisk_id* parameter. Estimated completion time is displayed in the YYMMDDHHMMSS format. The command displays progress for the following special cases as:

- A synchronized copy displays a progress of 100 and a blank estimated completion time.
- An offline copy or a copy with a zero synchronization rate displays a blank estimated completion time. An offline copy displays (gradually) decreasing progress if the volume is being written to.
- Nonmirrored volumes are displayed as a single copy with a progress of 100, and a blank estimated completion time.

The **lsvdisksyncprogress** command also displays the progress of a mirrored volume synchronization. After you create a mirrored volume by using the **mkvdisk** or **addvdiskcopy** command, you can use the command to monitor the progress of the synchronization.

lsvolumebackup

Use the **lsvolumebackup** command to list the volumes that have cloud snapshot that enabled and volumes that have cloud snapshots in the cloud account.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

Description

This command lists the volumes that use cloud snapshot and also lists volumes that have cloud snapshots in the cloud account.

This view spans all cloud accounts. To refresh the view by reloading what is on the cloud, specify `chcloudaccountawss3 -refresh 0` or `chcloudaccountswift -refresh 0`.

A snapshot that is being copied to the cloud (which means the snapshot status value is `copying` or `copying_error`) does not count towards the generation count total for the volume. Not counting towards the generation means that it is not available for a restore, and if there is a local system failure, the generation no longer exists in the cloud.

The specified volume appears in the displayed output when the volume has cloud snapshot that is enabled. The generation count is initially 0 and remains 0 while the volume copy is in progress. The generation count changes to 1 after the copy completes.

The last snapshot time is blank while the first snapshot is in progress. A snapshot that is being deleted from the cloud counts toward the generation count for the volume even if it is not available for a restore.

This table provides the attribute values that can be displayed as output view data.

Table 164. lsvolumebackup output

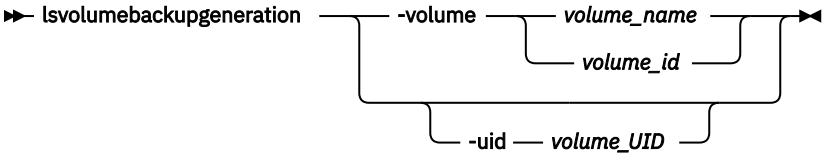
Attribute	Description
volume_UID	Indicates the volume UID.
volume_id	Indicates the volume ID if a volume with the specified UID exists on the local system. The value must be a number (or blank).
volume_name	Indicates the volume name. The value must be an alphanumeric string.
volume_group_id	Indicates the volume group ID that the volume is a member of (if applicable), if a volume with the specified UID exists on the local system. The value must be a number (or blank).
volume_group_name	Indicates the volume group name that the volume is a member of (if applicable), if a volume with the specified UID exists on the local system. The value must be an alphanumeric string (or blank).
cloud_account_id	Indicates the ID for the cloud account that contains the volume snapshots.
cloud_account_name	Indicates the cloud account name for the cloud account that contains the volume snapshot. The value must be an alphanumeric string.
last_backup_time	Indicates the timestamp of the most recent snapshot for this volume. The value must be in YYMMDDHHMMSS format or blank.
generation_count	Indicates the number of snapshot generations that exist for the specified volume. The value must be a number. **Note:** • Any generations that are being copied to the cloud do not count towards this number. • Any generations that are being deleted count towards this number until the delete process completes.

Table 164. lsvolumebackup output (continued)	
Attribute	**Description**
backup_size	Indicates the approximate amount of storage (the capacity) that is in use by snapshot generations for the specified volume

lsvolumebackupgeneration

Use the **lsvolumebackupgeneration** command to list any volume snapshots available on the specified volume.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

-volume *volume_name* | *volume_id*
(Required) Specifies the volume to list cloud snapshots for by volume ID or name.

Note: The volume must exist on the local system.

The value for volume ID must be a number and the value for volume name must be an alphanumeric string. This parameter is mutually exclusive with **-uid**.

-uid *volume_UID*
(Optional) Specifies the volume to list cloud snapshots for by volume UID. This parameter is mutually exclusive with **-volume**.

Description

This table provides the attribute values that can be displayed as output view data.

Table 165. lsvolumebackupgeneration output	
Attribute	**Description**
generation_id	Indicates the snapshot generation volume ID. The value must be a number.
backup_time	Indicates the timestamp of the most recent snapshot. The value must be in YYMMDDHHMMSS format (or blank).
auto_backup	Indicates whether the cloud backup is triggered manually or automatically through a policy schedule.

Table 165. lsvolumebackupgeneration output (continued)	
Attribute	**Description**
backup_expiry_time	Indicates expiry time of the backup generation. The value must be in YYMMDDHHMMSS format.
volume_group_name	Indicates the volume group name. The value must be an alphanumeric string (or blank).
volume_size	Indicates the provisioned capacity of the volume during snapshot generation. This value can differ from the current volume size (the capacity in MB or GB) if a volume with the specified UID exists on the local system.
type	Indicates the type of volume snapshot generation. The values are `full` and `incremental`.
state	Indicates the state of the volume backup generation (in the cloud system). The values are: • `copying` • `complete` • `deleting`
cloud_upload_size	Indicates the amount of data (the capacity in MB or GB) that is uploaded from the snapshot generation volume to the cloud system.

Generations that are being copied to the cloud account are included in the `copying` state value. Generations that are being deleted from the cloud account are also included in the `deleting` state value.

lsvolumebackupprogress

Use the **lsvolumebackupprogress** command to display information about the progress of snapshot operations.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

Description

This command display information about the progress of snapshot operations.

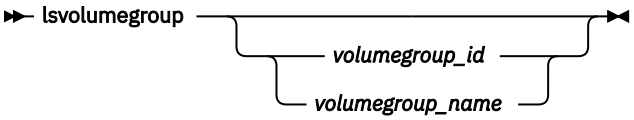
This table provides the attribute values that can be displayed as output view data.

Table 166. lsvolumebackupprogress output	
Attribute	**Description**
volume_UID	Indicates the volume UID. The value must be a number 0 - 32 characters long.
volume_id	Indicates the volume ID (if a volume with the specified UID exists on the local system). The value must be a number (or blank).
volume_name	Indicates the volume name (if a volume with the specified UID exists on the local system). The value must be an alphanumeric string (or blank).
task	Indicates the type of task that is in progress. The values are `backup` and `delete`.
status	Indicates the task status. The values are: • `copying` • `copying_error` • `deleting` • `deleting_error`
generation_id	Indicates the generation ID for the volume that is created or deleted. The value must be a number.
backup_time	Indicates the snapshot time for the volume that is copied to the cloud system. The value must be in YYMMDDHHMMSS format for snapshot tasks or blank for deletion tasks.
progress	Indicates the task progress as a percentage. The value must be a number 0 - 99.
error_sequence_number	Indicates a particular error number. The value must be a number (or blank).

lsvolumegroup

Use the **lsvolumegroup** command to display information about configured volume groups.

Syntax



Parameters

-delim *delimiter*
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

volumegroup_id | volumegroup_name
(Optional) Specifies a volume group ID or volume group name.

Description

This command displays information about configured volume groups.

This table provides the attribute values that can be displayed as output view data.

Table 167. lsvolumegroup output

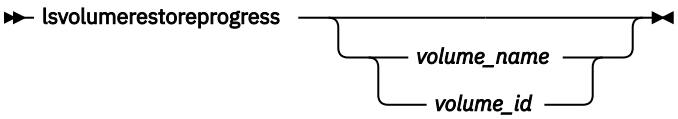
Attribute	Description
volume_count	Indicates the number of volume members in a group.
backup_status	Indicates whether a new (volume) group snapshot can be started. The values are: • `empty` • `off` • `not_ready` • `ready` • `copying` • `copying_error` If a group snapshot is in progress, the value represents the snapshot operation status.
last_backup_time	Indicates the time of the most recent snapshot of the volume group (`YYMMDDHHMMSS`).
owner_type	Indicates the type of ownership for the volume group. This field displays **vvol** if the volume group is used for VMware Virtual Volumes, or **none** otherwise.
safeguarded_policy_start_time	The value specified with **-policytime**.
snapshot_policy_id	Indicates the snapshot policy ID associated with volume group.
snapshot_policy_name	Indicates the snapshot policy name that is associated with volume group.
volume_group_type	Indicates the type of volume group to create. This view includes all volume groups of type clone or thinclone.
source_volume_group_id	Indicates the ID of the source volume group. Displays *blank* if group doesn't exist.
source_volume_group_name	Indicates the name of the source volume group. Displays *blank* if group doesn't exist.
parent_uid	Indicates the UID of the parent volume group this volume group was created from.
source_snapshot	Indicates the name of the source snapshot.
snapshot_count	Indicates the number of snapshots that are associated with this volume group. This count includes snapshot created for the volumes.
protection_provisioned_capacity	The amount of capacity provisioned to store point-in-time copies of volumes in the volume group. This includes snapshots and consistency protection for replication.
protection_written_capacity	The amount of capacity written to store point-in-time copies of volumes in the volume group before any data reduction. This includes snapshots and consistency protection for replication.
uid	Indicates the UID to uniquely identify the volume group.

Table 167. lsvolumegroup output (continued)	
Attribute	**Description**
replication_policy_id	Indicates the ID of the replication policy that is assigned to the volume group. By default, this field displays the DR replication policy assigned to the volume group.
replication_policy_name	Indicates the name of the replication policy that is assigned to the volume group. By default, this field displays the DR replication policy assigned to the volume group. **Note:** If the volume group is in a storage partition with an HA replication policy, and the HA partner system is running a version of software earlier than 8.7.0.0, this field displays the HA replication policy to maintain backwards compatibility.
ignore_user_flash_copy_maps	Indicates whether or not the volume group contains source volumes that are used in FlashCopy mappings. Value can be *yes* or *no*.
safeguarded_snapshot_count	Indicates the number of Safeguarded snapshots associated with the volume group.
partition_id	Indicates the ID of the partition associated with the volume group.
partition_name	Indicates the name of the partition associated with the volume group.
restore_in_progress	Indicate whether one or more volumes within the volume group are currently being restored.
draft_partition_id	Indicates the ID for draft partition associated with the volume group.
draft_partition_name	Indicates the name of draft partition associated with the volume group.
last_restore_time	Indicates the time that the volume group was last restored. The timestamp is in the format YYMMDDhhmmss. It will be blank if the **volumegroup** has never been restored.
ha_replication_policy_id	Indicates the ID of the HA replication policy assigned to the storage partition containing this volume group.
ha_replication_policy_name	Indicates the name of the HA replication policy assigned to the storage partition containing this volume group.
dr_replicated	Indicates if the volume group has replication configured for disaster recovery.
partition_default	Indicates whether the volume group is tagged as a default volume group. If it is in a partition, it will act as default volume group for the partition

lsvolumerestoreprogress

Use the **lsvolumerestoreprogress** command to display information about restore operation progress.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

volume_name | volume_id
(Optional) Indicates the volume name or ID for the volume that is restored. The value for volume name must be an alphanumeric string and the value for volume ID must be a number.

Description

This command displays information about restore operation progress.

This table provides the attribute values that can be displayed as output view data.

Table 168. lsvolumerestoreprogress output

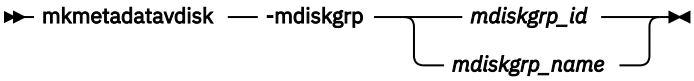
Attribute	Description
volume_id	Indicates the volume ID for the volume that is restored. The value must be a number (or blank).
volume_name	Indicates the volume name for the volume that is restored. The value must be an alphanumeric string (or blank).
task	Indicates the type of task that is in progress. The value is `restore`.
status	Indicates the task status. The values are: • `restoring` • `restoring_error`
generation_id	Indicates the generation ID for the volume snapshot generation that is restored. The value must be a number.
backup_time	Indicates the time for the volume snapshot generation that is being restored to the cloud system (or available on the restore volume). The value must be in YYMMDDHHMMSS format for snapshot tasks or blank for deletion tasks.
progress	Indicates the task progress as a percentage. The value must be a number 0 - 100. The status is blank when the status is `uncommitted`.
error_sequence_number	Indicates a particular error number. The value must be a number (or blank).
volume_backup_id	Indicates the UID of the volume snapshot that is restored. The value must be a number 1 - 32.
restore_volume_id	Indicates the ID of the volume that is the target of the restore operation. This volume is either the production volume (which means that the restore volume ID is the same as the volume ID) or a temporarily restored volume that is automatically provisioned by the restore process. The value must be a number.

Table 168. lsvolumerestoreprogress output (continued)	
Attribute	**Description**
restore_volume_name	Indicates the name of the volume that is the target of the restore operation. This volume is either the production volume (which means that the restore volume name is the same as the volume name) or a temporarily restored volume that is automatically provisioned by the restore process. The value must be an alphanumeric string.

mkmetadatavdisk

Use the `mkmetadatavdisk` command to create a metadata volume for use with VMware Virtual Volumes (vVols).

Syntax



Parameters

-mdiskgrp *mdiskgrp_id* | *mdiskgrp_name*
(Required) Specifies the storage pool in which to create the metadata volume.

Note: Ensure that the storage pool is a parent pool.

Restriction: You cannot specify a data reduction pool with this command.

Description

This command creates a metadata volume for use with VMware Virtual Volumes (vVols).

An invocation example

```
mkmetadatavdisk -mdiskgrp pool_a
```

The resulting output:

```
Virtual Disk, id [2], successfully created
```

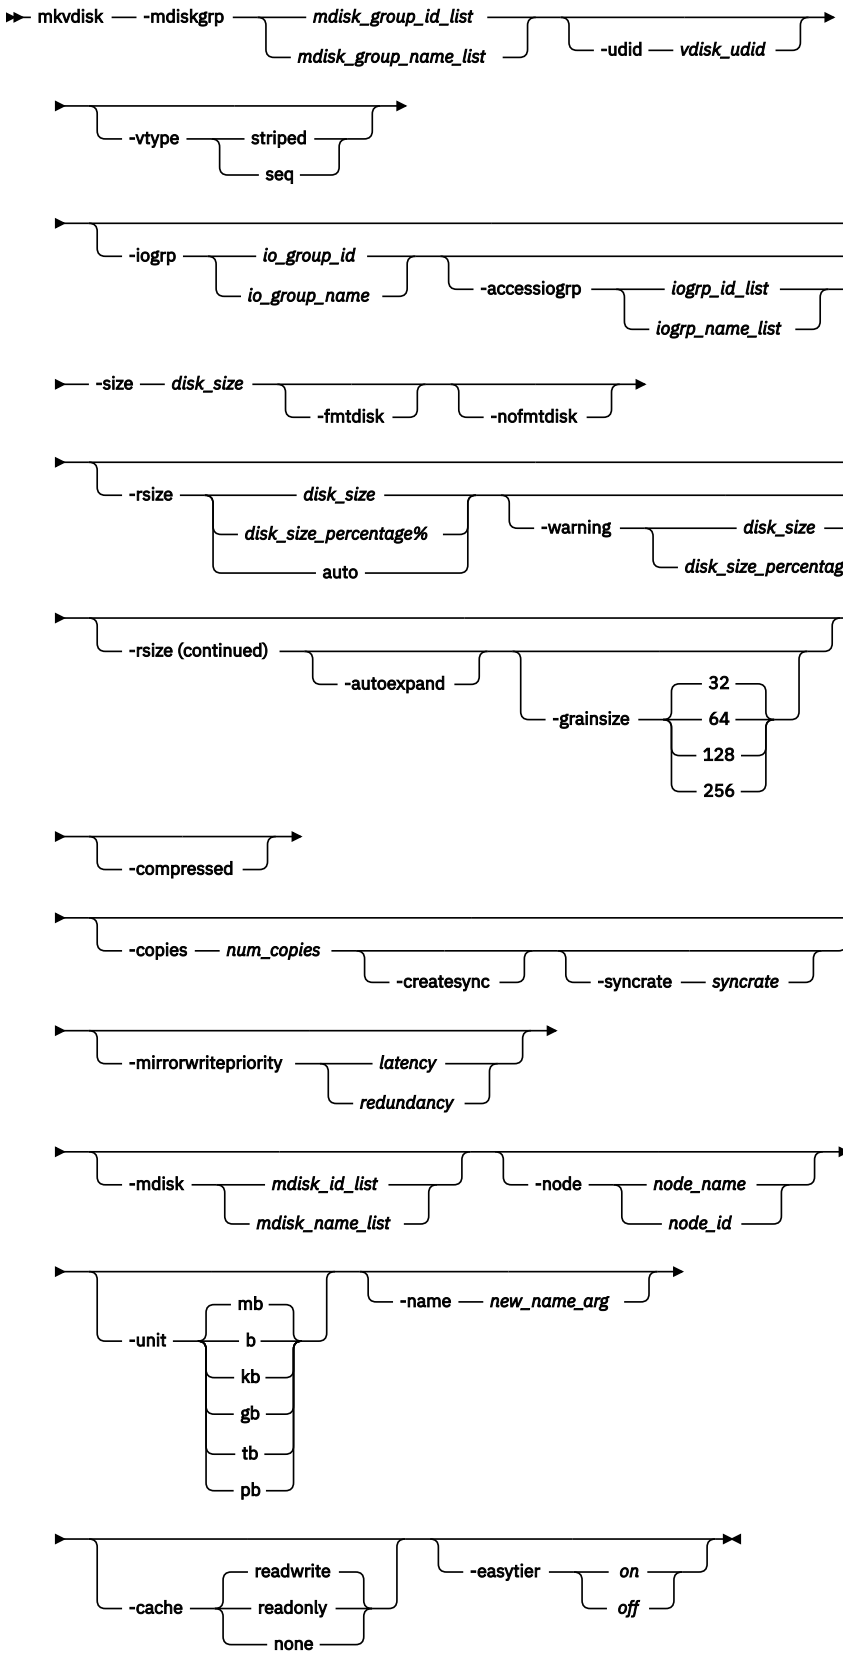
mkvdisk

Use the **mkvdisk** command to create sequential, striped, or image mode volume objects. When they are mapped to a host object, these objects are seen as disk drives with which the host can run I/O operations. The first syntax diagram is for striped or sequential volumes and the second syntax diagram for image mode volumes. Use the **mkvolume** command for a simplified way of creating high availability volumes. Use the **mkimagevolume** command for a simplified way of creating an image mode volume, importing existing data from a managed disk.

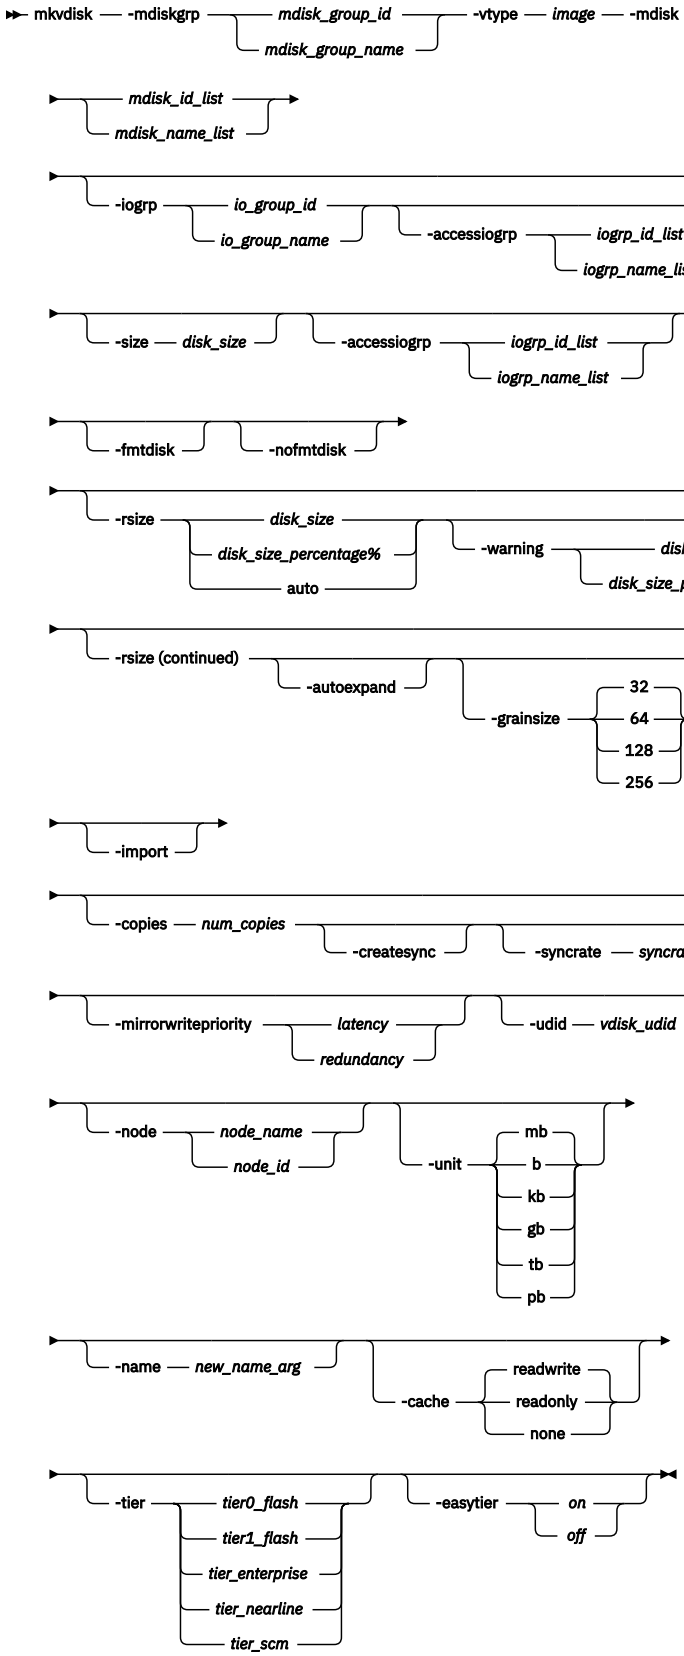
Syntax

Note: The first syntax diagram depicts the creation of a **sequential** or **striped** mode volume. The second syntax diagram depicts the creation of an **image** mode volume.

Create a sequential or striped mode volume.



Create an `image` mode volume.


mkvdisk
-mdiskgrp
mdisk_group_id
mdisk_group_name
-vtype
image
-mdisk
mdisk_id_list
mdisk_name_list
-iogrp
io_group_id
io_group_name
-accessiogrp
iogrp_id_list
iogrp_name_list
-size
disk_size
-accessiogrp
iogrp_id_list
iogrp_name_list
-fmtdisk
-nofmtdisk
-rsize
disk_size
disk_size_percentage%
auto
-warning
disk_size
disk_size_percentage%
-rsize (continued)
-autoexpand
-grainsize
32
64
128
256
-import
-copies
num_copies
-createsync
-syncrate
syncrate
-mirrorwritepriority
latency
redundancy
-udid
vdisk_udid
-node
node_name
node_id
-unit
mb
b
kb
gb
tb
pb
-name
new_name_arg
-cache
readwrite
readonly
none
-tier
tier0_flash
tier1_flash
tier_enterprise
tier_nearline
tier_scm
-easytier
on
off

Parameters

-mdiskgrp ***mdisk_group_id_list | mdisk_group_name_list***
(Required) Specifies one or more storage pools to use when you are creating this volume. If you are creating multiple copies, you must specify one storage pool per copy. The primary copy is allocated from the first storage pool in the list.

Note: If you are creating volumes in a pool that is used as a Safeguarded backup location, only one pool can be specified. Mirrored volumes are not supported.

-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`).

-vtype seq | striped | image
(Optional) Specifies the virtualization type. When you create sequential or image mode volumes, you must also specify the **-mdisk** parameter. You cannot use `-vtype` *`seq`* or `-vtype` *`image`* in a data reduction pool. The default virtualization type is striped.

-iogrp ***io_group_id | io_group_name***

(Optional) Specifies the I/O group (node pair) with which to associate this volume. If you specify **-node**, you must also specify **-iogrp**.

Remember:

- Create the first compressed volume copy for an I/O group to activate compression.
- You cannot create or move a volume copy that is compressed to an I/O group that contains at least one node that does not support compressed volumes. You must select another I/O group to move the volume copy to (but it does not affect moving to the recovery I/O group).

The system auto-selects an I/O group based on the number of volumes in other I/O groups, with exception to systems with V5030 nodes where regular pool compression cannot be used with data reduction pool compression within the same I/O group. On a system that has at least one V5030 in it, auto I/O group selection can be used. However, the compression type must match the compression type already in use. If you create an opposite compression type, then **-iogrp** must be specified.

Note: If the user does not specify an I/O group then the system automatically selects the I/O group that supports the capabilities required (e.g. compression, deduplication) with the fewest volumes in.

-accessiogrp ***iogroup_id_list | iogroup_name_list***

(Optional) Specifies the members of the volume I/O group access set. If this option is not specified, only the caching I/O group is added to the volume I/O group access set. If this option is specified, then the **-iogrp** option must also be specified and the I/O group access set must include the I/O group that is specified by the **-iogrp** parameter.

-size ***disk_size***
(Required for sequential or striped volume creation) (Optional for image volume creation) Specifies the capacity of the volume, which is used with the value of the unit. All capacities, including changes, must be in multiples of 512 bytes. An error occurs if you specify a capacity that is not a multiple of 512. It can happen only 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.

-fmtdisk
(Optional) Specifies that the volume is formatted. This parameter is no longer required for any volumes.

This parameter is not required when you create standard-provisioned volumes. The format operation is automatically applied to standard-provisioned volumes unless you specify **-nofmtdisk** parameter.

The format operation sets the extents that make up this volume to all zeros after it is created. This process takes place in the background concurrently with host I/O operations on the new volume.

Remember: Formatting is on by default for single copy, standard-provisioned, and non-image mode volumes. You cannot format an image mode volume.

The format operation completes asynchronously. You can query the status by using the **lsvdiskprogress** comm. You cannot specify this parameter with the **-vtype image** parameter.

This parameter is not required when you create thin-provisioned volumes. Thin-provisioned volumes return zeros for extents that are not written to. No format operation is required. This parameter also synchronizes mirrored copies by default.

-nofmtdisk
(Optional) Specifies that formatting be turned off for the new volume.

Remember: Formatting is on by default for single copy, standard-provisioned, and non-image mode volumes, and you can specify this parameter to turn it off.

-rsize *disk_size | disk_size_percentage%* **| auto**
(Optional) Defines how much physical space is initially allocated to the thin-provisioned or compressed volume. This parameter makes the volume thin-provisioned; otherwise, the volume is standard-provisioned. Specify the *disk_size* | *disk_size_percentage* value by using an integer, or an integer immediately followed by the percent character (**%**). Specify the units for a *disk_size* integer by using the **-unit** parameter; the default is MB. The **-rsize** value can be greater than, equal to, or less than the size of the volume. The **auto** option creates a volume copy that uses the entire size of the MDisk.

If you specify the **-rsize auto** option, you must also specify the **-vtype image** option. If you specify , **-import** you must specify **-rsize**.

If the volume is in a data reduction storage pool, the value of the **-rsize** parameter will be ignored in **mkvdisk**. Only its presence or absence is used to determine whether the disk is a data reduction volume copy or a thick volume copy.

-warning *disk_size | disk_size_percentage%*
(Optional) Requires that the **-rsize** parameter also be specified. Specifies a threshold at which a warning error log is generated for volume copies. A warning is generated when the used disk capacity on the thin-provisioned copy first exceeds the specified threshold.

Note: You cannot specify this parameter for thin-provisioned or compressed volumes that are in data reduction pools.

You can specify a *disk_size* integer, which defaults to MBs unless the **-unit** parameter is specified. Or you can specify a *disk_size***%**, which is a percentage of the volume size.

Important: If **-autoexpand** is:

1. Enabled, the default value for **-warning** is 80% of the volume capacity.
2. Not enabled, the default value for **-warning** is 80% of the real capacity.

To disable warnings, specify `0`.

-autoexpand
(Optional) Specifies that thin-provisioned copies automatically expand their real capacities by allocating new extents from their storage pool. Requires that the **-rsize** parameter also be specified. If the **-autoexpand** parameter is specified, the **-rsize** parameter specifies a capacity that is reserved by the copy. It protects the copy from going offline when its storage pool runs out of space by having the storage pool to use this reserved space first.

The parameter has no immediate effect on image mode copies. However, if the image mode copy is later migrated to managed mode, the copy is then automatically expanded.

-grainsize 32 | 64 | 128 | 256
(Optional) Sets the grain size (KB) for a thin-provisioned volume. This parameter also requires that the **-rsize** parameter be specified. If you are using the thin-provisioned volume in a FlashCopy map, use

the same grain size as the map grain size for best performance. If you are using the thin-provisioned volume directly with a host system, use a small grain size. The grain size value must be 32, 64, 128, or 256 KB. The default is 256 KB.

If the volume to be created is a thin-provisioned and compressed volume in a data reduction storage pool, the **-grainsize** parameter cannot be used. This type of volume will be created with a grain size of 8 KB.

-compressed
(Optional) Enables compression for the volume. This parameter must be specified with **-rsize** and cannot be specified with **-grainsize**.

-import
(Optional) Imports a thin-provisioned volume from the MDisk. If you specify, **-import** you must also specify **-rsize**.

-copies *num_copies*
(Optional) Specifies the number of copies to create. The *num_copies* value can be 1 or 2. Setting the value to 2 creates a mirrored volume. The default value is 1.

-createsync
(Optional) Creates copies in sync. Use this parameter if you already formatted the MDisks, or when read stability to unwritten areas of the volume is not required.

-syncrate *syncrate*
(Optional) Specifies the copy synchronization rate. A value of (0) prevents synchronization. The default value is 50. See Table 169 on page 736 for the supported **-syncrate** values and their corresponding rates. Use this parameter to alter the rate at which the standard-provisioned volume or mirrored volume format before synchronization.

-mirrorwritepriority *latency | redundancy*
(Optional) Specifies how to configure the mirror write algorithm priority. If not specified, the default value is *latency*.

1. Choosing *latency* means a copy that is slow to respond to a write input/output (I/O) becomes unsynchronized, and the write I/O completes if the other copy successfully writes the data.
2. Choosing *redundancy* means a copy that is slow to respond to a write I/O synchronizes completion of the write I/O with the completion of the slower I/O to maintain synchronization.

-mdisk *mdisk_id_list | mdisk_name_list*
(Optional) Specifies one or more managed disks. For sequential and image mode volumes, the number of MDisks must match the number of copies. For sequential mode volumes, each MDisk must belong to the specified storage pool. For striped volumes, you cannot specify the **-mdisk** parameter if the **-copies** value is greater than 1.

When you create a single copy striped volume, you can specify a list of MDisks to stripe across.

You must use this parameter to specify an MDisk that has a mode of unmanaged.

-node *node_id | node_name*
(Optional) Specifies the preferred node ID or the name for I/O operations to this volume. You can use the **-node** parameter to specify the preferred access node. If you specify **-node**, you must also specify **-iogrp**.

Note: This parameter is evaluated by multipath device drivers. The system chooses a default if you do not supply this parameter.

-unit b | kb | mb | gb | tb | pb
(Optional) Specifies the data units to use along with the capacity that is specified by the **-size** and **-rsize** parameters. The default unit type is MB.

-name *new_name_arg*
(Optional) Specifies a name to assign to the new volume.

-cache readwrite | readonly | none
(Optional) Specifies the caching options for the volume. Valid entries are:

- *readwrite* enables the cache for the volume.
- *readonly* disables write caching while allowing read caching for a volume.
- *none* disables the cache mode for the volume.

The default is *readwrite*.

-tier ***tier0_flash | tier1_flash | tier_enterprise | tier_nearline | tier_scm***
(Optional) Specifies the MDisk tier when an image mode copy is added. This parameter is valid only when **-vtype** is image.

tier0_flash
Specifies a `tier0_flash` hard disk drive or an external MDisk for the newly discovered or external volume.

tier1_flash
Specifies an `tier1_flash` (or flash drive) hard disk drive or an external MDisk for the newly discovered or external volume.

tier_enterprise
Specifies a `tier_enterprise` hard disk drive or an external MDisk for the newly discovered or external volume.

tier_nearline
Specifies a `tier_nearline` hard disk drive or an external MDisk for the newly discovered or external volume.

tier_scm
Specifies a `tier_scm` hard disk drive or an external MDisk for the newly discovered or external volume.

ssd
Specifies an SSD (or flash drive) hard disk drive or an external MDisk for the newly discovered or external volume.

nearline
Specifies a nearline hard disk drive or an external MDisk for the newly discovered or external volume.

enterprise
Specifies an enterprise hard disk drive or an external MDisk for the newly discovered or external volume.

Note: This action applies to both copies if you are creating a mirrored volume with two image mode copies by using this command.

-easytier ***on | off***
Determines whether the IBM Easy Tier function is allowed to move extents for this volume.

Note: The **-easytier** parameter must be followed by either *on* or *off*:

- If set to *on*, then Easy Tier functions are active.
- If set to *off*, then Easy Tier functions are inactive.

-udid ***vdisk_udid***
(Optional) Specifies the unit number *udid* for the disk. The `udid` is an identifier that is required to support OpenVMS hosts; no other systems use this parameter. Valid options are a decimal number 0 - 32 767, or a hexadecimal number 0 - 0x7FFF. A hexadecimal number must be preceded by `0x` (for example, `0x1234`).

-size ***disk_size***
(Required for sequential or striped volume creation) (Optional for image volume creation) Specifies the capacity of the volume, which is used with the value of the unit. All capacities, including changes, must be in multiples of 512 bytes. An error occurs if you specify a capacity that is not a multiple of 512. It happens 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.

-unit b | kb | mb | gb | tb | pb
(Optional) Specifies the data units to use with the capacity that is specified by the **-size** and **-rsize** parameters. The default unit type is MB.

-name *new_name_arg*
(Optional) Specifies a name to assign to the new volume.

Description

This command creates a new volume object. You can use the command to create various types of volume objects, making it one of the most complex commands.

Remember: You can create a striped volume only in a child pool - you cannot create sequential or image volumes in a child pool.

You must decide which storage pool or storage pools provide the storage for the volume. Use the **lsmdiskgrp** command to list the available storage pools and the amount of free storage in each pool. If you are creating a volume with more than one copy, each storage pool that you specify must have enough space for the size of the volume.

If you create a thin-provisioned or compressed volume from a data reduction storage pool, that volume uses the same properties as the data reduction storage pool. You can create standard-provisioned volumes from data reduction pools, but these volumes use different data reduction properties.

A thin-provisioned or compressed volume that is in a data reduction storage pool must:

- Not be in sequential or image mode.
- Not have a warning threshold set (using **-warning**).
- Use `-cache readwrite` when caching.
- Have **-autoexpand** enabled.

For thin-provisioned and compressed volume copies that are in data reduction storage pools, the Easy Tier status is taken from the data reduction pool because that data is managed by a central data disk. It means that you cannot specify **-easytier** with thin-provisioned or compressed volumes.

Important: The extent size for the storage pool can limit volume size. Consider the maximum volume size that you want to use when you create storage pools. Refer to the information on creating storage pools for a comparison of the maximum volume capacity for each extent size. The maximum is different for thin-provisioned volumes.

No restriction exists for the number of compressed volumes within a data reduction storage pool.

Choose an I/O group for the volume. This action determines which nodes in the system process the I/O requests from the host systems. If you have more than one I/O group, ensure that you distribute the volumes between the I/O groups so that the I/O workload is shared evenly between all nodes. Use the **lsiogrp** command to show the I/O groups and the number of volumes that are assigned to each I/O group.

Note: It is normal for systems with more than one I/O group to have storage pools that have volumes in different I/O groups. FlashCopy processing can make copies of volumes whether the source and target volumes are in the same I/O group. However, if you plan to use intra-system Metro or Global Mirror operations, ensure that both the master and auxiliary volume are in the same I/O group.

The command returns the IDs of the newly created volume.

Specify the virtualization type by using the **-vtype** parameter; the supported types are sequential (*seq*), *striped*, and *image*.

sequential (seq)
This virtualization type creates the volume that uses sequential extents from the specified MDisk (or MDisks, if creating multiple copies). The command fails if there are not enough sequential extents on the specified MDisk.

striped
The default virtualization type. If the **-vtype** parameter is not specified, *striped* is the default; all managed disks in the storage pool are used to create the volume. The striping is at an extent level; one extent from each managed disk in the group is used. For example, a storage pool with 10 managed disks uses one extent from each managed disk. It then uses the 11th extent from the first managed disk, and so on.

If the **-mdisk** parameter is also specified, you can supply a list of managed disks to use as the stripe set. It can be two or more managed disks from the same storage pool. The same circular algorithm is used across the striped set. However, a single managed disk can be specified more than once in the list. For example, if you enter `-mdisk 0:1:2:1`, the extents are from the following managed disks: 0, 1, 2, 1, 0, 1, 2, and so forth. All 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 volumes to be created when a managed disk already has data on it, perhaps from a previrtualized subsystem. When an image mode volume is created, it directly corresponds to the (previously unmanaged) managed disk that it was created from. Therefore, except for thin-provisioned image mode volumes, volume logical block address (LBA) *x* equals managed disk LBA *x*. You can use this command to bring a nonvirtualized disk under the control of the system. After it is under the control of the system, you can migrate the volume from the single managed disk. When it is migrated, the volume is no longer an image mode volume.

You can add image mode volumes to an already populated storage pool with other types of volumes, such as a striped or sequential.

Important: An image mode volume must be 512 bytes or greater. At least one extent is allocated to an image mode volume.

Remember: If you create a mirrored volume from two image mode MDisks without specifying a **-size** value, the capacity of the resulting volume is the smaller of the two MDisks, and the remaining space on the larger MDisk is not accessible.



Attention:

1. Do not create a volume in an offline I/O group. You must ensure that the I/O group is online before you create a volume to avoid any data loss. This action applies in particular to re-creating volumes that are assigned the same object ID.
2. To create an image mode disk, you must already have a quorum disk present in the system because an image mode disk cannot be used to hold quorum data. Refer to information on quorum disk creation for more details.
3. The command fails if either limit of 2048 volumes per I/O Group or 8192 volume copies per system is reached.

The rate at which the volume copies resynchronize after loss of synchronization can be specified by using the **-syncrate** parameter. This table provides the relationship of the *syncrate* value to the data copied per second.

Note: These settings also affect the initial rate of formatting.

Table 169. Relationship between the syncrate value and the amount of data copied per second

User-specified *syncrate* attribute value	**Data copied per second**
1 - 10	128 KB
11 - 20	256 KB

Table 169. Relationship between the syncrate value and the amount of data copied per second (continued)	
User-specified *syncrate* attribute value	**Data copied per second**
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
101 - 110	128 MB
111 - 120	256 MB
121 - 130	512 MB
131 - 140	1 GB
141 - 150	2 GB

An invocation example

```
mkvdisk -mdiskgrp Group0 -size 0
-iogrp 0 -vtype striped -mdisk mdisk1 -node 1
```

The resulting output:

```
Virtual Disk, id [1], successfully created
```

An invocation example to turn off formatting while creating volume `Burnley1`

```
mkvdisk -mdiskgrp Burnley1 -iogrp 0 -mdiskgrp 0:1 -size 500 -nofmtdisk -copies 2
```

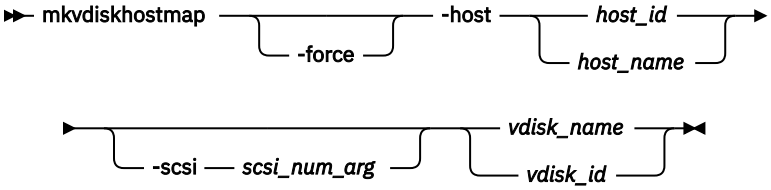
The resulting output:

```
Virtual Disk, id [0], successfully created
```

mkvdiskhostmap

Use the **mkvdiskhostmap** command to create a new mapping between a volume and a host, which makes the volume accessible for input/output (I/O) operations to the specified host.

Syntax



Parameters

-force
(Optional) Allows multiple volume-to-host assignments, which are not normally allowed.

-host *host_id | host_name*
(Required) Specifies the host to map the volume to, either by ID or by name.

-scsi *scsi_num_arg*
(Optional) Specifies the Small Computer System Interface (SCSI) logical unit number (LUN) ID to assign to this volume on the given host. The **scsi_num_arg** parameter contains the SCSI LUN ID that is assigned to the volume on the given host for all I/O groups that provide access to the volume. If you do not specify the **-scsi** parameter, the next available SCSI LUN ID is used. The **-scsi** parameter is not supported for NVMe hosts. When a volume is mapped to a host that uses the NVMe protocol, the volume namespace ID (NSID) is assigned by the system.

vdisk_name | vdisk_id
(Required) Specifies the name of the volume that you want to map to the host, either by ID or by name.

Description

This command creates a new mapping between the volume and the specified host. The volume is presented to the host as if the disk is directly attached to the host. It is only after this command is processed, that the host can perform I/O transactions to the volume.

Optionally, you can assign a SCSI LUN ID to the mapping for SCSI hosts. When the HBA in the host scans for devices that are attached to it, it discovers all volumes that are mapped to its Fibre Channel ports. When the devices are found, each one is allocated an identifier (SCSI LUN ID). For example, the first disk found is usually `SCSI LUN 0`, and so on. You can control the order in which the HBA discovers volumes by assigning the SCSI LUN ID, as required. If you do not specify a SCSI LUN ID, the cluster automatically assigns the next available SCSI LUN ID, if any mappings already exist with that host. When you issue the **mkvdiskhostmap** command, the assigned SCSI LUN ID number is returned.

For NVMe, the SCSI LUN ID is irrelevant and instead the system assigns the NSID that cannot be changed. When you create a mapping to an NVMe volume, the NSID is returned (The NSID is the vdisk_ID + 1).

The **mkvdiskhostmap** command fails if the:

- Host to which this mapping is being made is not associated with any one of the I/O groups in the volume access set
- Volume has more than one I/O group in its access set and the host being mapped to the volume does not support volumes being mapped from multiple I/O groups
- Host protocol is incompatible with the volume protocol

The returned ID is same for the I/O groupall I/O groups to which the volume is mapped.

If you generate different SCSI LUN IDs, only one is returned. The returned ID is for the highest-numbered I/O group to which the volume was mapped. To view other values, issue **lshostvdiskmap** or **lsvdiskhostmap**.

The SCSI LUN ID is used for the highest numbered I/O group to which the volume is mapped.

Some HBA device drivers stop when they find a gap in the SCSI LUN IDs. For example:

- Volume 1 is mapped to Host 1 with SCSI LUN ID 1
- Volume 2 is mapped to Host 1 with SCSI LUN ID 2
- Volume 3 is mapped to Host 1 with SCSI LUN ID 4

When the device driver scans the HBA, it may stop after identifying volumes 1 and 2, because no SCSI LUN is mapped with ID 3. For optimal configuration, ensure that the SCSI LUN ID allocation is contiguous.

You can create multiple volume assignments (assigning the same volume to multiple hosts, for example, which might be necessary for clustered system hosts assigning a volume to multiple hosts). Normally,

multiple volume-to-host assignments are not used because corruption is likely to occur if more than one host can access a disk. To map to more than one host, you must use the **mkvdiskhostmap** command with the **-force** parameter. For example: `mkvdiskhostmap -host host1 -force 4` `mkvdiskhostmap -host host2 -force 4`.

Note: It is recommended to use host clusters if a volume is to be mapped to multiple hosts. This ensures that the same SCSI LUN ID is used for all hosts.

These commands create two host-to-volume mappings for volume 4 that map to `host1` and `host2`. Omitting the **-force** parameter causes the mapping to fail if that volume is already mapped to a host.

The command also fails if the host object (to which this mapping is being made) is not associated with the I/O group containing the volume.

An invocation example

```
mkvdiskhostmap -host host1 -scsi 1 5
```

The resulting output:

```
Virtual Disk to Host map, id [1], successfully created
```

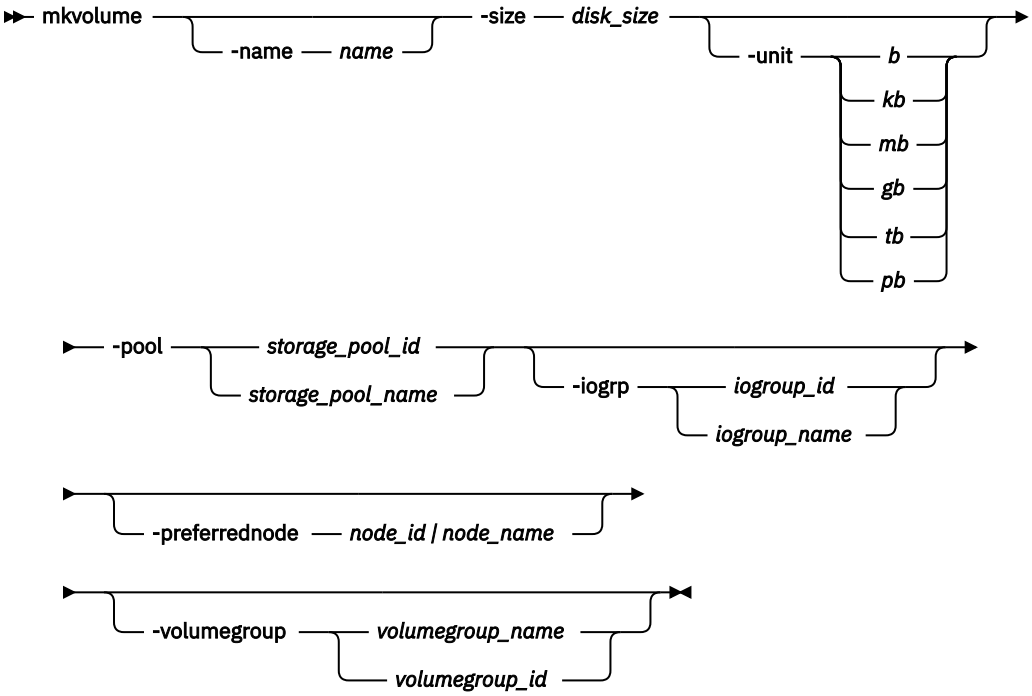
mkvolume

Use the **mkvolume** command to create an empty volume from existing storage pools.

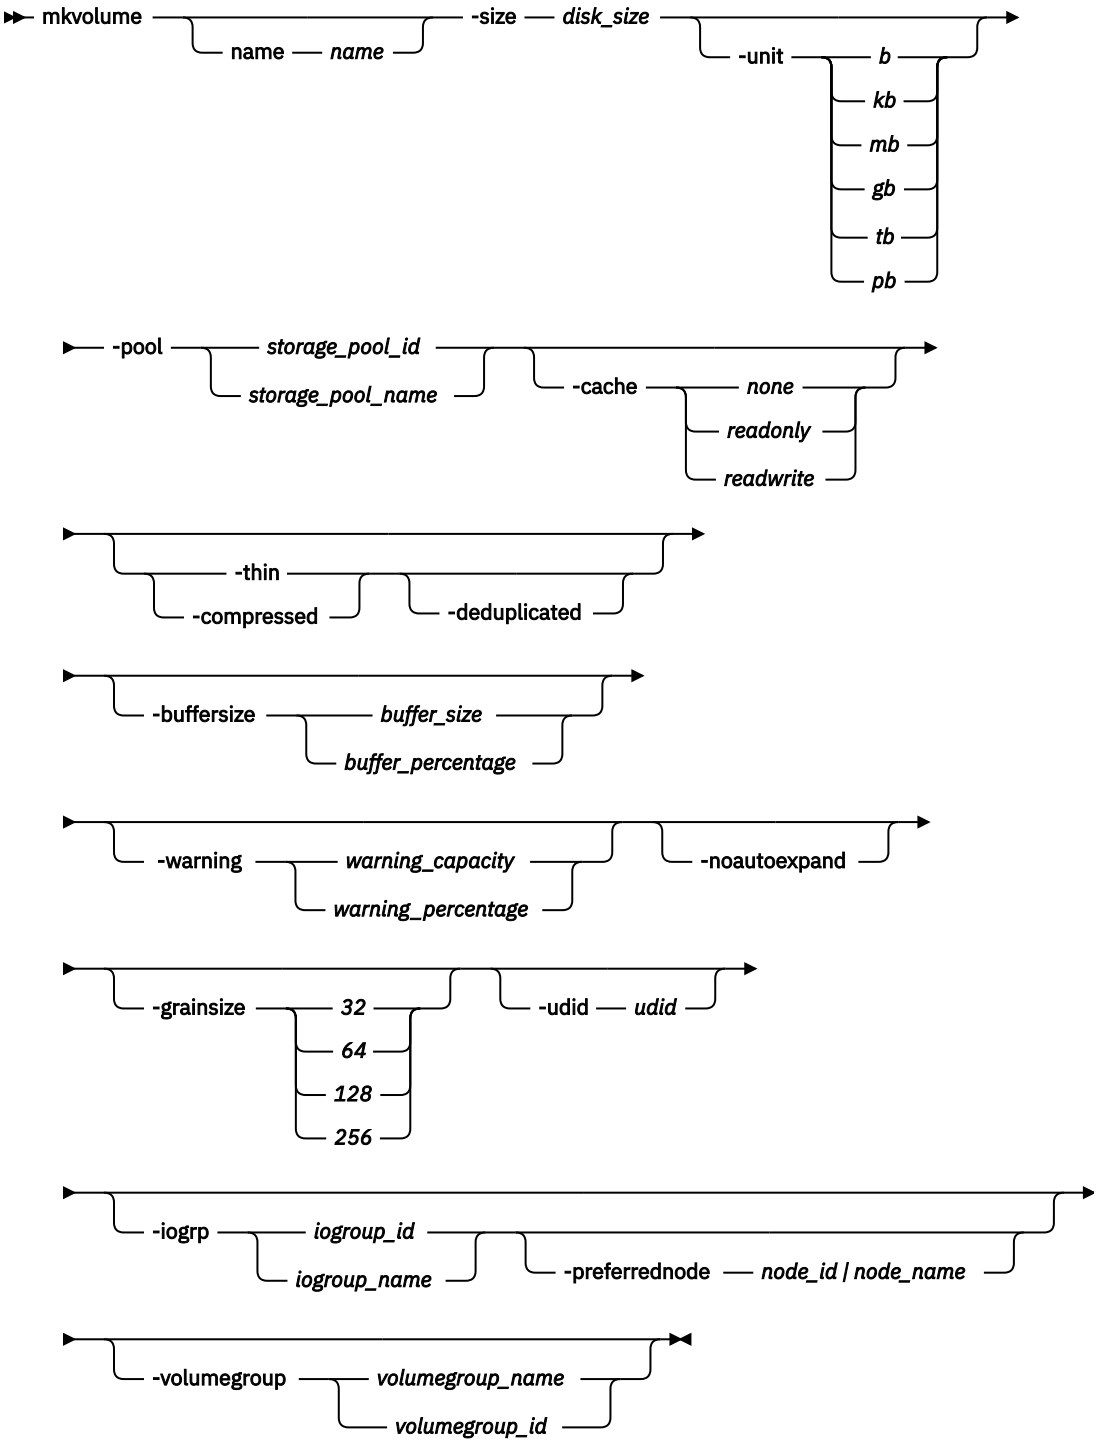
If you are creating volumes in a pool that uses a provisioning policy, the volume is created with the default values and capacity saving characteristics that are defined in the provisioning policy. Any volume that is created in a pool with a provisioning policy adopts those values and alternative values cannot be specified with the **mkvolume** command. For pools with a provisioning policy, the **mkvolume** command parameters are simplified.

Syntax

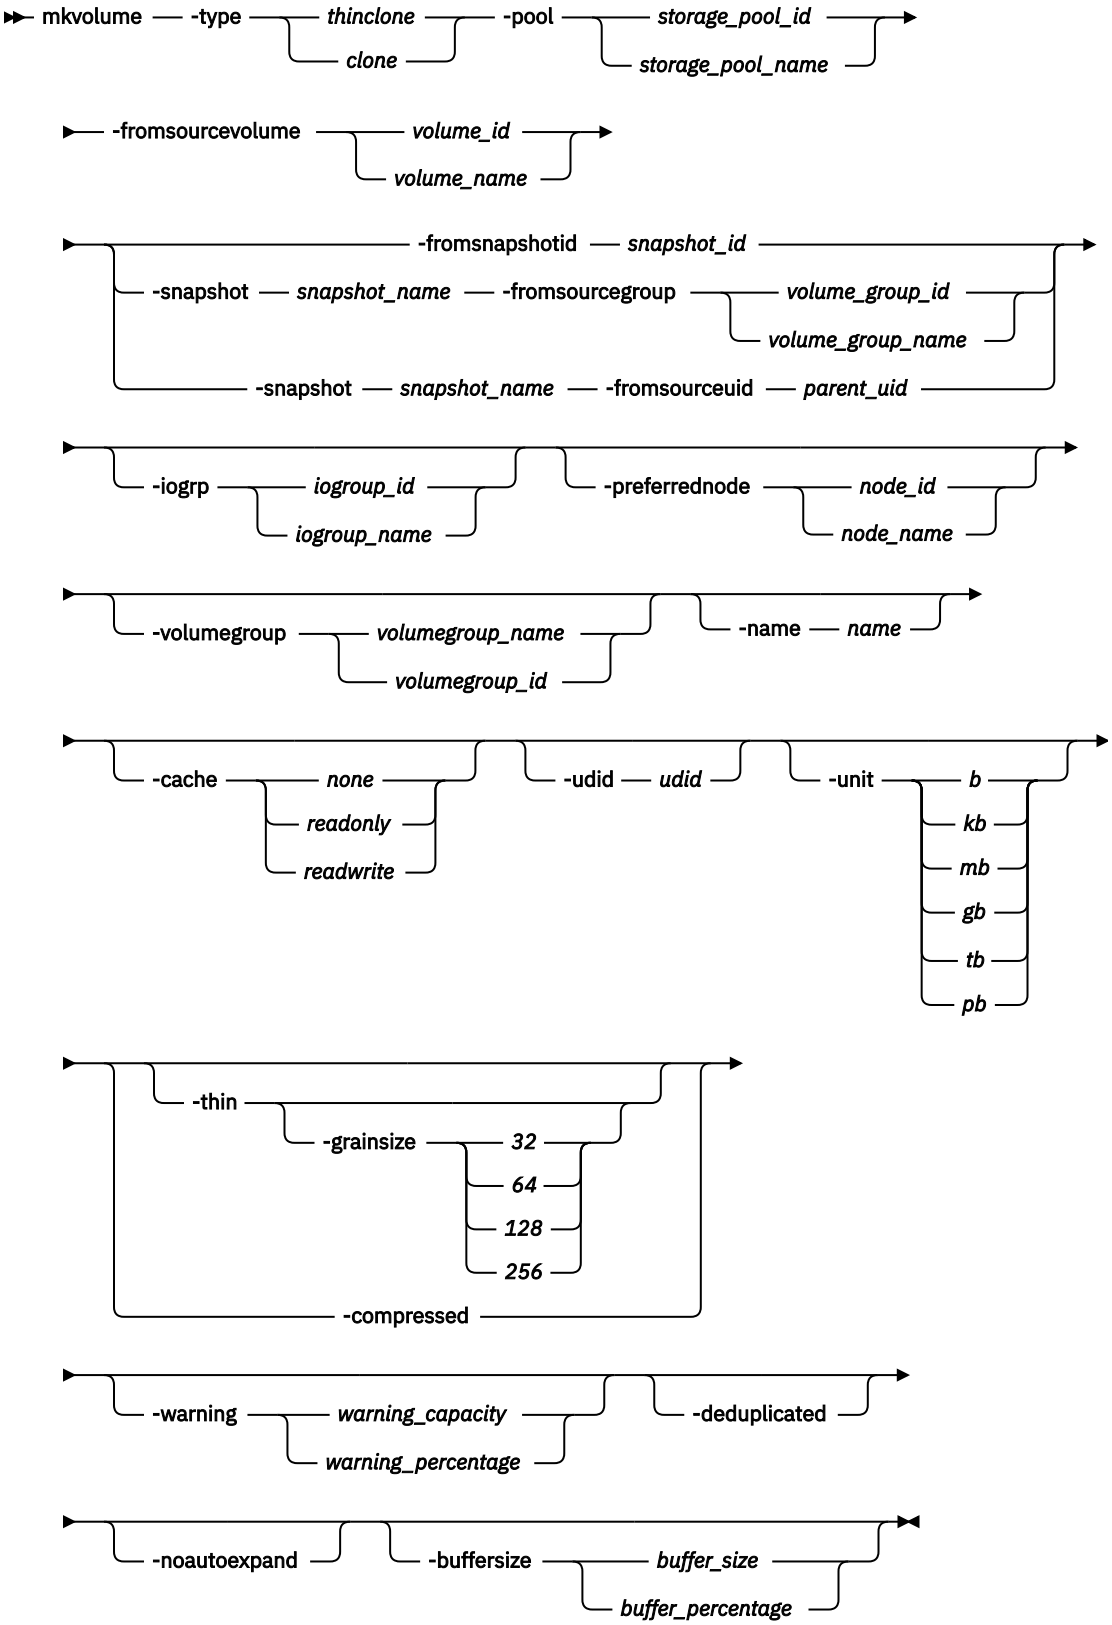
Syntax with a provisioning policy



Syntax without a provisioning policy



Syntax using existing snapshot



Parameters

Parameters with a provisioning policy

-name *name*
(Optional) Specifies the name that is used for the volume that is created. This value must be an alphanumeric string 1 - 63 characters.

Remember: If you do not specify **-name**, a unique default name such as `volume1` is used.

-size *disk_size*
(Required) Specifies the capacity of the volume, which is used with the value of the unit. The default capacity is in MB. When the unit of bytes is used, all capacities must be in multiples of 512 bytes. An entire extent is reserved even if it is only partially used.

-unit *b | kb | mb | gb | tb | pb*
(Optional) Specifies the data units to use with the capacity that is specified by the **-size** parameter. The default unit type is *mb*.

-iogrp *iogroup_id | iogroup_name*
(Optional) Specifies the I/O group that the new volume is cached in. If no value is specified, the caching I/O group is selected by the system based on the storage pool site.

-iogrp *iogroup_id_list | iogroup_name_list*
(Optional) Specifies the I/O group that the new volume is cached in. The value can be a colon-separated list of up to two I/O group IDs or names. If no value is specified, the caching I/O group is selected based on the storage pool site. If you do not specify the **-iogrp** parameter, the caching I/O group is selected by the system.

Important: If two I/O groups are specified, they must be in different sites, and the storage pools that are specified must be in different sites. The order of the sites must correspond.

-preferrednode *node_id | node_name*
(Optional) Specifies the preferred node that is used to access the volume.

Note: The -**preferrednode** parameter can only be specified when a single I/O group is specified in the **-iogrp** parameter.

-pool *storage_pool_id_list | storage_pool_name_list*
(Required) Specifies the storage pool in which to create the new volume. The value must be a colon-separated list of up to two storage pool IDs or names.

Note: If one storage pool is specified, a basic volume is created with one copy.

On systems with `standard` topology, a mirrored volume can be created by specifying two storage pools.

-volumegroup *volumegroup_name | volumegroup_id*
(Optional) Specifies the volume group that the volume should be added to.

Parameters without a provisioning policy

-name *name*
(Optional) Specifies the name that is used for the volume that is created. This value must be an alphanumeric string 1 - 63 characters.

Remember: If you do not specify **-name**, a unique default name such as `volume1` is used.

-size *disk_size*
(Required) Specifies the capacity of the volume, which is used with the value of the unit. The default capacity is in MB. When the unit of bytes is used, all capacities must be in multiples of 512 bytes. An entire extent is reserved even if it is only partially used.

-unit *b | kb | mb | gb | tb | pb*
(Optional) Specifies the data units to use with the capacity that is specified by the **-size** parameter. The default unit type is *mb*.

-iogrp *iogroup_id | iogroup_name*
(Optional) Specifies the I/O group that the new volume is cached in. If no value is specified, the caching I/O group is selected by the system based on the storage pool site.

-iogrp *iogroup_id_list | iogroup_name_list*
(Optional) Specifies the I/O group that the new volume is cached in. The value can be a colon-separated list of up to two I/O group IDs or names. If no value is specified, the caching I/O group is selected based on the storage pool site. If you do not specify the **-iogrp** parameter, the caching I/O group is selected by the system.

-preferrednode *node_id | node_name*
(Optional) Specifies the preferred node that is used to access the volume.

Important: If two I/O groups are specified, they must be in different sites, and the storage pools that are specified must be in different sites. The order of the sites must correspond.

-pool *storage_pool_id_list | storage_pool_name_list*
(Required) Specifies the storage pool in which to create the new volume. The value must be a colon-separated list of up to two storage pool IDs or names.

Note: If one storage pool is specified, a basic volume is created with one copy.

On systems with `standard` topology, a mirrored volume can be created by specifying two storage pools.

-cache *none | readonly | readwrite*
(Optional) Specifies the caching options for the volume. Use one of the following valid entries:

- `readwrite` enables the cache for the volume (default).
- `readonly` disables write caching but allows read caching for a volume.
- `none` disables the cache mode for the volume.

-thin
(Optional) Specifies that the volume is to be created with thin provisioning. You cannot specify this parameter with **-compressed**. If you do not specify **-thin** or **-compressed**, the volume that is created is standard-provisioned.

-compressed
(Optional) Specifies that the volume is to be created compressed. If the `-iogrp` parameter is not specified, the least used I/O group is used for compressed copies (considering the subset of I/O groups that support compression).

Remember: This command fails if no I/O groups support compression.

You cannot specify this parameter with **-thin**. If you do not specify **-thin** or **-compressed**, the volume that is created is standard-provisioned.

-deduplicated
(Optional) Creates a deduplicated volume. If you specify **-deduplicated**, you must also specify **-thin** or **-compressed** because it applies only to thin-provisioned or compressed volumes.

Note: Data deduplication works only with data reduction storage pools. You can only create deduplicated volumes and volume copies in an I/O group if there are no compressed volumes or volume copies in regular storage pools.

-buffersize *buffer_size | buffer_percentage*
(Optional) Specifies the pool capacity the volume attempts to reserve as a buffer for thin-provisioned and compressed volumes. You must specify either **-thin** or **-compressed** with this parameter. The default value is 2%. If you are adding a volume to a pool that uses a provisioning policy, this parameter uses the default of 2% .

Note: You cannot specify a buffer size for thin-provisioned or compressed volumes that are in data reduction pools.

-warning *warning_capacity | warning_percentage*
(Optional) Specifies a threshold at which a warning error log is generated for volumes. A warning is generated when the used disk capacity on the thin-provisioned volume exceeds the specified threshold. You must specify either **-thin** or **-compressed** with this parameter. The default value is 80%.

-noautoexpand
(Optional) Specifies that the volume will not automatically expand as it is written to. The available buffer capacity decreases as the used capacity increases. The volume copy goes offline if the buffer capacity is fully used. The buffer capacity can be increased by specifying `expandvdisksize -rsize`. You must specify either **-thin** or **-compressed** with this parameter. If you do not specify **-noautoexpand**, the volume automatically expands as it is written to.

-grainsize *32 | 64 | 128 | 256*
(Optional) Sets the grain size (KB) for a thin-provisioned volume. If you are using the thin-provisioned volume in a FlashCopy map, use the same grain size as the map grain size for best performance. If you are using the thin-provisioned volume directly with a host system, use a small grain size. The grain size value must be 32, 64, 128, or 256 KB. The default is 256 KB.

-udid *udid*
(Optional) Deprecated.

-volumegroup *volumegroup_name | volumegroup_id*
(Optional) Specifies the volume group that a volume belongs to. The value must be either an alphanumeric string for the volume group name or a number for the volume group ID.

Parameters using existing snapshot

-type *thinclone | clone*
(Required) Specifies the type of volume to create. Volume can be `thinclone` or `clone` type.

-pool *storage_pool_id_list | storage_pool_name_list*
(Required) Specifies the storage pool in which to create the new volume. The value must be a colon-separated list of up to two storage pool IDs or names.

Note: If one storage pool is specified, a basic volume is created with one copy.

On systems with `standard` topology, a mirrored volume can be created by specifying two storage pools.

-fromsourcevolume *volume_id | volume_name*
(Required) Specifies the image of the volume ID or name in the snapshot used to pre-populate the new volume.

-fromsnapshotid *snapshotid*
(Required) Specifies the ID for the snapshot used to pre-populate the new volume.

-snapshot *snapshot_name*
(Optional) Specifies the snapshot name used to pre-populate the new volume.

-fromsourcegroup *volume_group_id_or_name*
(Optional) Specifies the parent volume group ID or name of the snapshot to use to pre-populate the new volume.

-fromsourceuid *parent_uid*

(Optional) Specifies the UID used to identify the parent volumes of the snapshot to use to pre-populate the new volume.

Common parameters

-iogrp *iogroup_id | iogroup_name*
(Optional) Specifies the I/O group that the new volume is cached in. If no value is specified, the caching I/O group is selected by the system based on the storage pool site.

-iogrp *iogroup_id_list | iogroup_name_list*
(Optional) Specifies the I/O group that the new volume is cached in. The value can be a colon-separated list of up to two I/O group IDs or names. If no value is specified, the caching I/O group is selected based on the storage pool site. If you do not specify the **-iogrp** parameter, the caching I/O group is selected by the system.

Important: If two I/O groups are specified, they must be in different sites, and the storage pools that are specified must be in different sites. The order of the sites must correspond.

-preferrednode *node_id | node_name*
(Optional) Specifies the preferred node that is used to access the volume.

Note: The -**preferrednode** parameter can only be specified when a single I/O group is specified in the **-iogrp** parameter.

-volumegroup *volumegroup_name | volumegroup_id*
(Optional) Specifies the volume group that a volume belongs to. The value must be either an alphanumeric string for the volume group name or a number for the volume group ID.

-name *name*
(Optional) Specifies the name that is used for the volume that is created. This value must be an alphanumeric string 1 - 63 characters.

Remember: If you do not specify **-name**, a unique default name such as `volume1` is used.

-cache *none | readonly | readwrite*
(Optional) Specifies the caching options for the volume. Use one of the following valid entries:

- `readwrite` enables the cache for the volume (default).
- `readonly` disables write caching but allows read caching for a volume.
- `none` disables the cache mode for the volume.

-udid *udid*
(Optional) Specifies the unit number `udid` for the volume.

Important: The `udid` is an identifier that is required to support OpenVMS hosts (no other systems use this parameter).

Valid options are a decimal number from `0` through `32767`, or a hexadecimal number from `0` through `0x7FFF`. A hexadecimal number must be preceded by `0x` (for example, `0x1234`).

-unit *b | kb | mb | gb | tb | pb*
(Optional) Specifies the data units to use with the capacity that is specified by the **-size** parameter. The default unit type is *mb*.

-thin
(Optional) Specifies that the volume is to be created with thin provisioning. You cannot specify this parameter with **-compressed**. If you do not specify **-thin** or **-compressed**, the volume that is created is standard-provisioned.

-grainsize *32 | 64 | 128 | 256*
(Optional) Sets the grain size (KB) for a thin-provisioned volume. If you are using the thin-provisioned volume in a FlashCopy map, use the same grain size as the map grain size for best performance. If you are using the thin-provisioned volume directly with a host system, use a small grain size. The grain size value must be 32, 64, 128, or 256 KB. The default is 256 KB.

-compressed
(Optional) Specifies that the volume is to be created compressed. If the `-iogrp` parameter is not specified, the least used I/O group is used for compressed copies (considering the subset of I/O groups that support compression).

Remember: This command fails if no I/O groups support compression. If two sites exist, both sites must have at least one I/O group that supports compression.

You cannot specify this parameter with **-thin**. If you do not specify **-thin** or **-compressed**, the volume that is created is standard-provisioned.

-warning *warning_capacity | warning_percentage*
(Optional) Specifies a threshold at which a warning error log is generated for volumes. A warning is generated when the used disk capacity on the thin-provisioned volume exceeds the specified threshold. You must specify either **-thin** or **-compressed** with this parameter. The default value is 80%.

-deduplicated
(Optional) Creates a deduplicated volume. If you specify **-deduplicated**, you must also specify **-thin** or **-compressed** because it applies only to thin-provisioned or compressed volumes.

Note: Data deduplication works only with data reduction storage pools. You can only create deduplicated volumes and volume copies in an I/O group if there are no compressed volumes or volume copies in regular storage pools.

-noautoexpand
(Optional) Specifies that the volume will not automatically expand as it is written to. The available buffer capacity decreases as the used capacity increases. The volume copy goes offline if the buffer capacity is fully used. The buffer capacity can be increased by specifying `expandvdisksize -rsize`. You must specify either **-thin** or **-compressed** with this parameter. If you do not specify **-noautoexpand**, the volume automatically expands as it is written to.

-buffersize *buffer_size | buffer_percentage*
(Optional) Specifies the pool capacity the volume attempts to reserve as a buffer for thin-provisioned and compressed volumes. You must specify either **-thin** or **-compressed** with this parameter. The default value is 2%. If you are adding a volume to a pool that uses a provisioning policy, this parameter uses the default of 2% .

Note: You cannot specify a buffer size for thin-provisioned or compressed volumes that are in data reduction pools.

Description

This command creates an empty volume, which is a formatted (zeroed) volume, by using storage from existing storage pools. If the storage pool uses a provisioning policy, the volume that is created uses the capacity saving methods and default values that are defined in the provisioning policy. You do not have to define capacity savings at the volume level, making the process of creating volumes more efficient and faster. If you specify values when the pool is assigned to a provisioning policy, the mkvolume command fails. However, if you are using pools without provisioning policies, you still create volumes by specifying these values in the command. Provisioning policies are supported on standard and data reduction pools and child pools adopt the provisioning policy of the parent pool. Data reduction pools can be created with or without a provisioning policy. Data reduction pools use built-in capacity savings methods. For data reduction pools with or without provisioning policies, use the following information:

- If you create a thin-provisioned or compressed volume from a data reduction storage pool, the properties of the storage pool are used for the new volume. You can create standard-provisioned volumes from data reduction storage pools, but those volumes do not use the storage pool properties.
- On some node types, you can create a compressed volume copy in a data reduction storage pool for an I/O group. You can create thin-provisioned volume copies on any node type. Volumes can also have standard-provisioned volume copies in data reduction storage pools. A compressed volume copy in a data reduction pool can only be created in an I/O group with sufficient memory.
- You cannot specify **-noautoexpand** when you create thin-provisioned or compressed volume copies from a data reduction storage pool.
- You cannot create a volume copy that is a thin-provisioned or compressed volume in a data reduction storage pool, and the volume caching mode is `none` or `readonly`. You must specify **chvdisk** to change the volume caching mode to `readwrite`.
- You cannot specify **-warning** for a thin-provisioned or compressed volume copy in a data reduction storage pool.

- You cannot specify **-grainsize** for thin-provisioned and compressed volume copies in data reduction storage pools. This type of volume copy is created with a size of 8 KB.
- Thin-provisioned or compressed volume copies in data reduction pools cannot be created if the data reduction storage pool is offline and requires recovery. If the recovery is still in progress, you must wait until the recovery is complete and the pool is in online state.

Note: The **-volumegroup** parameter is not allowed if the volume group is associated with a snapshot policy and the volume has the following properties:

- The volume has zero size.
- The ownership group of the storage pool for the volume does not match the ownership group for the volume group.

Use the **mkimagevolume** command to create a new volume by importing existing data on a managed disk.

If a volume type is specified, the user must also provide a source volume with **-fromsourcevolume** parameter. They must also provide a snapshot using either the **-fromsnapshotid** parameter or the **-snapshot** parameter along with a **-fromsourcegroup** or **-fromsourceuid** parameter.

The volume properties are partially defined by clone or a thin-clone copy type. The pool attribute must be specified and it is possible to create a mirrored volume by specifying a colon separated list of pools. In addition the user can specifies the following parameters:

- **-iogrp**
- **-preferrednode**
- **-volumegroup**
- **-name**
- **-cache**
- **-buffersize**
- **-warning**
- **-noautoexpand**
- **-grainsize**
- **-thin**
- **-compressed**
- **-deduplicated**
- **-unit**

If the **-name** parameter is not specified, the new volume will be assigned a name based on the source volume, with the addition of a numerical suffix to align with the naming pattern established by the **mkvolumegroup** command. Specifying more than one **iogroup** is not allowed when using the **-fromsourcevolume** parameter. If the **-fromsourcevolume** parameter is specified, the following options will be rejected:

- **-size**
- **-skipinitialization**
- **-remotevolume1**

Creating a volume in a volume group associated with a storage partition

Note: Additional restrictions are applied if a volume group associated with a storage partition is specified. The **mkvolume** command will not run under the following conditions:

- cache disabled volume being created
- zero sized volume
- volume with a copy in a standard pool and a copy in a data reduction pool

When a storage partition has an HA replication policy, it is mandatory for the pool of each volume copy to be linked to all the remote systems specified in the replication policy.

Volumes can be created while the remote system in the replication policy is connected, or under the following situations while disconnected:

- If the partition has HA established and there is management access to the remote system, the command will succeed and the change will be applied to both systems.
- If the partition has HA established and there is no management access to the remote system then the command will fail.
- If the partition does not have HA established then the local volume create will succeed and the configuration change will be applied to the remote system asynchronously as part of re-establishing HA.

Creating a volume in a volume group associated with a storage partition with an HA replication policy is only permitted for the user on the active management system.

Invocation examples with a provisioning policy

An invocation example of any volume

```
mkvolume -pool 0 -size 1000
```

The detailed resulting output:

```
Volume, id [1], successfully created.
```

Invocation examples without a provisioning policy

An invocation example to create a volume in storage pool 0

```
mkvolume -pool 0 -size 1000
```

The detailed resulting output:

```
Volume, id [0], successfully created.
```

An invocation example to create a thin-provisioned volume from a data reduction storage pool

```
mkvolume -pool datareductionpool2 -size 10 -unit gb -thin
```

The detailed resulting output:

```
Volume, id [6], successfully created.
```

An invocation example

```
svctask mkvolume -type thinclone -fromsourcevolume sourceVolA -fromsourcegroup sourceGroupA
-snapshot noon
```

The detailed resulting output:

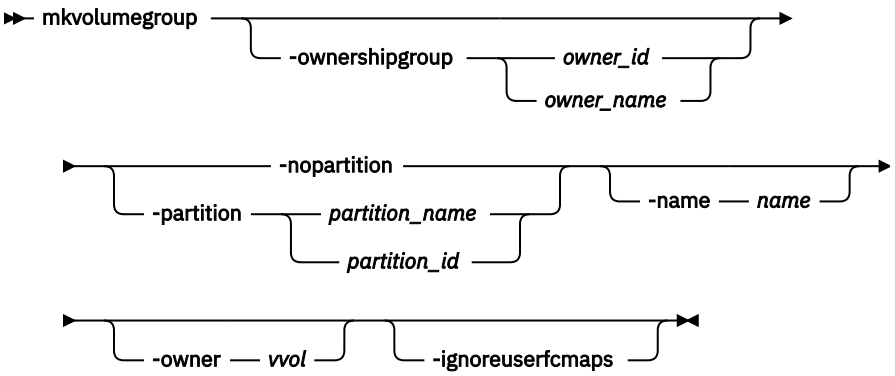
```
Volume, id [x], successfully created.
```

mkvolumegroup

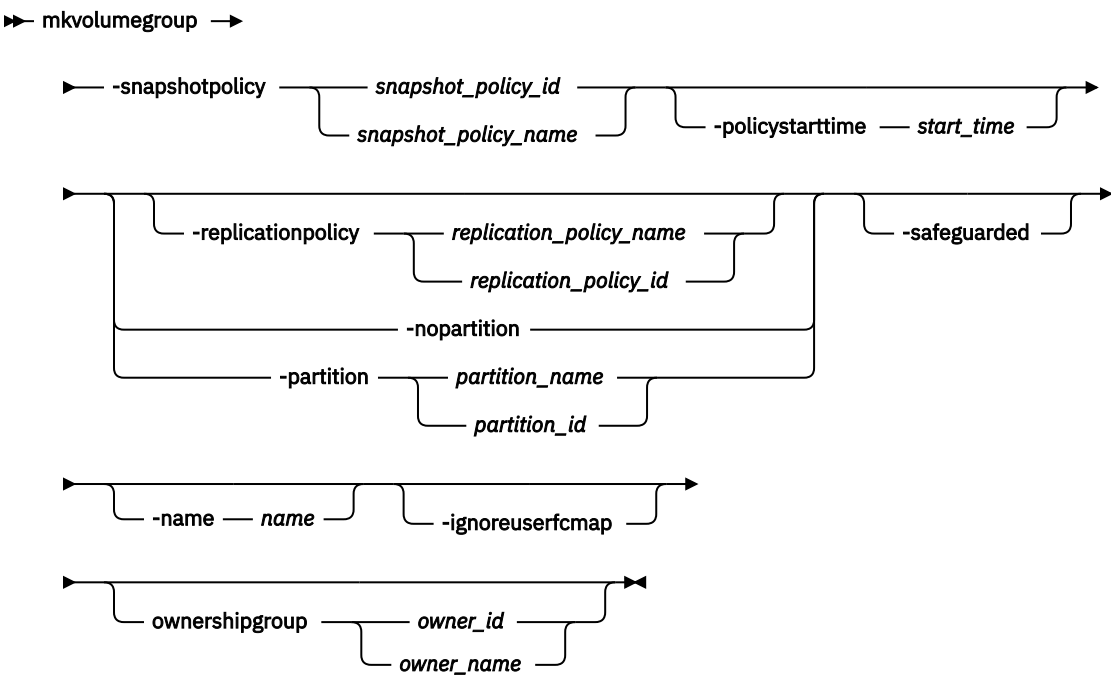
Use the **mkvolumegroup** command to create and configure a new volume group.

Syntax

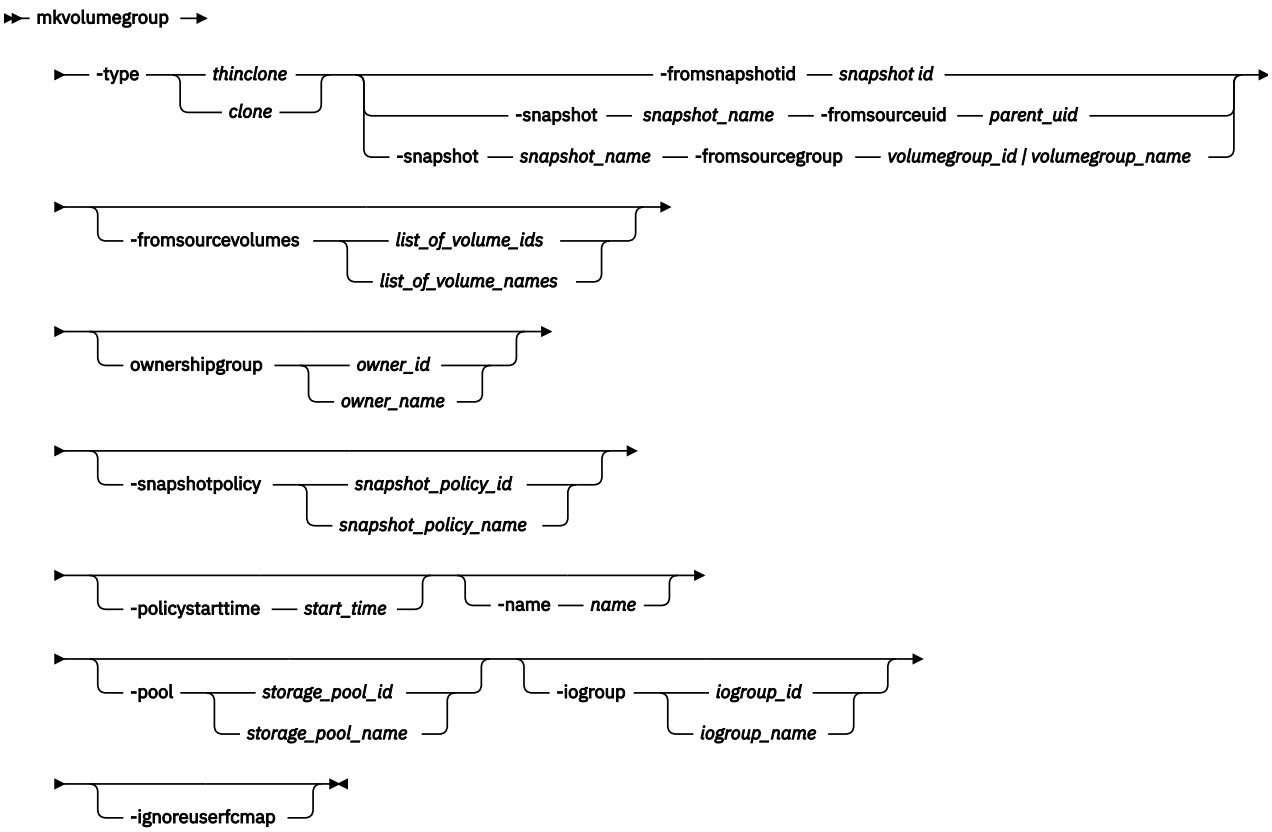
Make a basic volume group



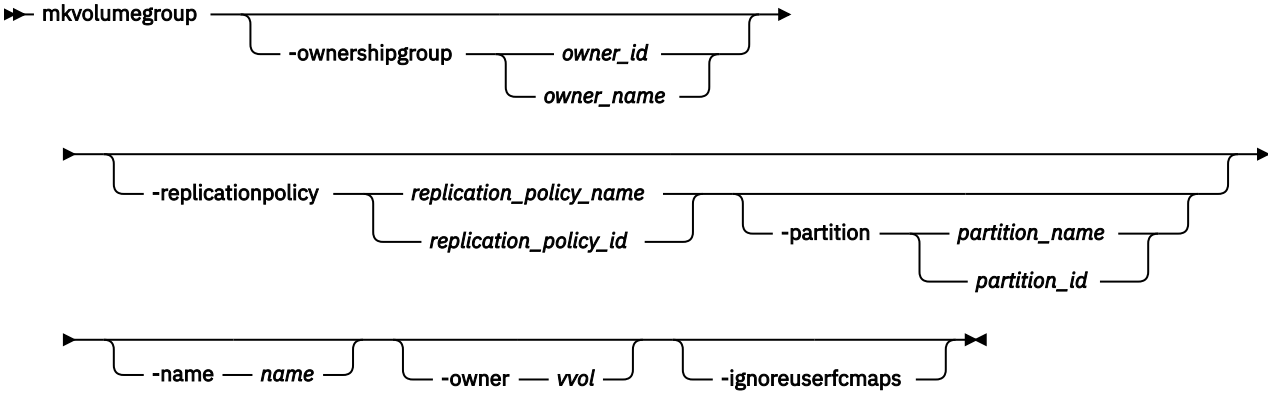
Make a volume group with a snapshot policy



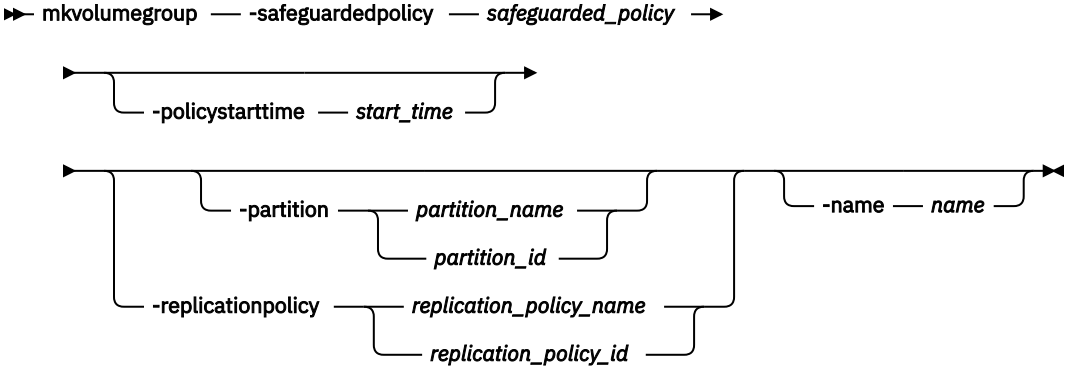
Make a volume group from a snapshot



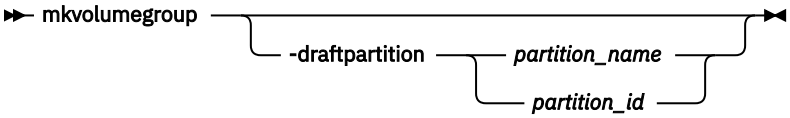
Make a volume group with a replication policy



Make a safeguarded volume group (this does not apply to safeguarded copy with volume group snapshot)



Make a volume group in a draft partition



Parameters

-name *volumegroup_name*
(Optional) Specifies a volume group name. The value must be an alphanumeric value. If you do not specify a volume group name, one is automatically created and assigned to the volume group.

-ownershipgroup *owner_id | owner_name*
(Optional) Specifies the name or ID of the ownership group to which the object is added.

-owner *vvol*
(Optional) Specifies the type of ownership for the volume group. The value must be **vvol** for the volume group to be used for VMware Virtual Volumes.

-snapshotpolicy *snapshot_policy_id | snapshot_policy_name*
(Optional) Specifies the snapshot policy to assign to the volume group.

-type *thinclone | clone*
(Optional) Specifies the type of volume to create.

-fromsourceuid *parent_uid*
(Optional) Specifies the UID used to identify the parent volumes of the snapshot to pre-populate the new volume in the new volume group.

-fromsourcegroup *volume_group_id | volume_group_name*
(Optional) Specifies the name or ID of the parent volume group of the snapshot to pre-populate the new volume in the new volume group.

-snapshot *snapshot_name*
(Optional) Specifies the name of the snapshot that is used to pre-populate the new volumes in the new volume group.

-fromsnapshotid *snapshotid*
(Optional) Specifies the ID of the snapshot used to pre-populate the new volumes in the new volume group.

-fromsourcevolumes *list_of_volume_ids | list_of_volume_names*
(Optional) Specifies the IDs or name of the parent volumes.

Note: When combined with the **-type** and **-snapshot** parameters, you can create a volume group with a subset of those volumes whose image is present in a snapshot.

-pool *pool_name_or_id*
(Optional) Specifies the name or ID of the pool to contain the volumes.

-iogroup *io_group_name_or_id*
(Optional) Specifies the name or ID of the I/O group for the new volume.

-replicationpolicy *replication_policy_name_or_id*
(Optional) Specifies the replication policy to assign to the volume group.

If the replication policy has an async-dr topology the volume group must be created in the storage partition associated with the async-dr replication policy. A replication policy with an async-dr topology is associated with only one storage partition.

If the associated storage partition is a 2-site-ha storage partition the volumegroup will be HA with async DR replication configured, otherwise it will create a volumegroup with DR replication. An async-dr replication policy will have DR linked storage partitions.

Note: This command can succeed whether the remote system in the replication policy is connected or disconnected. When connectivity is restored, the volume group is created on the remote system.

Note: You cannot assign a replication policy to the volume group if the group contains any volumes with the following characteristics:

- Volumes are in a different I/O group than the I/O group that is specified in the replication policy.
- Volumes have disabled cache.
- Volumes are image mode volumes.
- Volumes are used in Transparent Cloud Tiering (TCT).

-safeguarded
(Optional) Specifies that the snapshot policy creates Safeguarded snapshots of source volumes in the volume group.

Note: The **-safeguarded** parameter can only be specified with the **-snapshotpolicy** parameter.

-ignoreuserfcmaps
(Optional) Indicates that snapshots can be created with the scheduler or with the **addsnapshot** command, if any volumes in the volume group are used as a source volume in legacy FlashCopy mapping.

-partition ***partition_name | partition_id***
(Optional) Specifies the storage partition to assign to the volume group. This operation can be performed whether the remote systems in the partition are connected or disconnected. The volume group is created on the remote systems when connectivity allows. The command must be run on the active management system for the specified partition.

Note: If a storage partition is assigned to a volume group, that storage partition cannot be removed from the volume group and that volume group's replication policy can not be changed. A storage partition can only be removed from a volume group by removing the volume group using **rmvolumegroup**.

Note: This parameter is mutually exclusive with the **-type** parameter and the **-replicationpolicy** parameter. The **-ownershipgroup** parameter is not supported when this parameter is specified.

-draftpartition ***partition_name | partition_id***
(Optional) Specifies the draft partition to assign to the volume group.

-nopartition
(Optional) Creates a clone or a thinclone volume group outside of the partition that the volume group or a snapshot was in originally.

Note: The -nopartition parameter allows the restored snapshots to be created outside the original partition.

While creating a thinclone volume group from a particular source volume group or a snapshot, if the original source volume group or a snapshot is in a partition and you do not specify **-nopartition** parameter, then the new thinclone volume group is created in the same partition. However, if the original source volume group or a snapshot exists in an HA partition, then the **mkvolumegroup** command fails with the *CMMVC1243E* error.

-setpartitiondefault
(Optional) Mark the volume group as default volume group. When added to a partition, the volume group will act as default volume group for the partition.

Description

This command creates and configures a new volume group. The system supports snapshot and replication policies.

An invocation example

```
mkvolumegroup
```

The resulting output:

```
Volume Group, id [0], successfully created.
```

An invocation example

```
mkvolumegroup -name Sunday
```

The resulting output:

```
Volume Group, id [1], successfully created
```

An invocation example creates a volume group with an association to a snapshot policy:

```
mkvolumegroup  -snapshotpolicy 2
```

The resulting output:

```
Volume Group, id [4], successfully created.
```

An invocation example to create a volume group with an association to a snapshot policy and specifies a start time to use with this volume group:

```
mkvolumegroup  -snapshotpolicy 2 -policystarttime 2203011800
```

The resulting output:

```
Volume Group, id [5], successfully created.
```

An invocation example

```
mkvolumegroup -type thinclone -fromsourcegroup sourceGroupA -snapshot noon
```

The resulting output:

```
Volume Group, id [6], successfully created.
```

An invocation example

The following example creates a volume group and associates it with a storage partition:

```
mkvolumegroup -partition partition0
```

The resulting output:

```
Volume Group, id [0], successfully created
```

An invocation example

The following example creates a volume group and associates it with a storage partition:

```
mkvolumegroup -partition partition0
```

The resulting output:

```
Volume Group, id [0], successfully created
```

An invocation example

```
svctask mkvolumegroup -type thinclone -fromsourcegroup sourceGroupA -snapshot noon
-fromsourcevolumes 4:5
```

The resulting output:

```
Volume Group, id [0], successfully created
```

An invocation example

```
 svctask mkvolumegroup -type clone -fromsnapshotid 0 -name cloneVolumeGroupB -pool clonesPool
-fromsourcevolumes volA:volB:volC
```

The resulting output:

```
Volume Group, id [0], successfully created
```

An invocation example

The following example creates a volume group with **-nopartition**:

```
 mkvolumegroup -type clone -fromsourcegroup 0 -snapshot snapshot0 -nopartition
```

The resulting output:

```
Volume Group, id [0], successfully created
```

An invocation example

```

mkvolumegroup -type clone -fromsourcegroup 0 -snapshot ha_snapshot0
CMMVC1243E The command failed because the volume group is part of a storage partition that has
a replication policy.
```

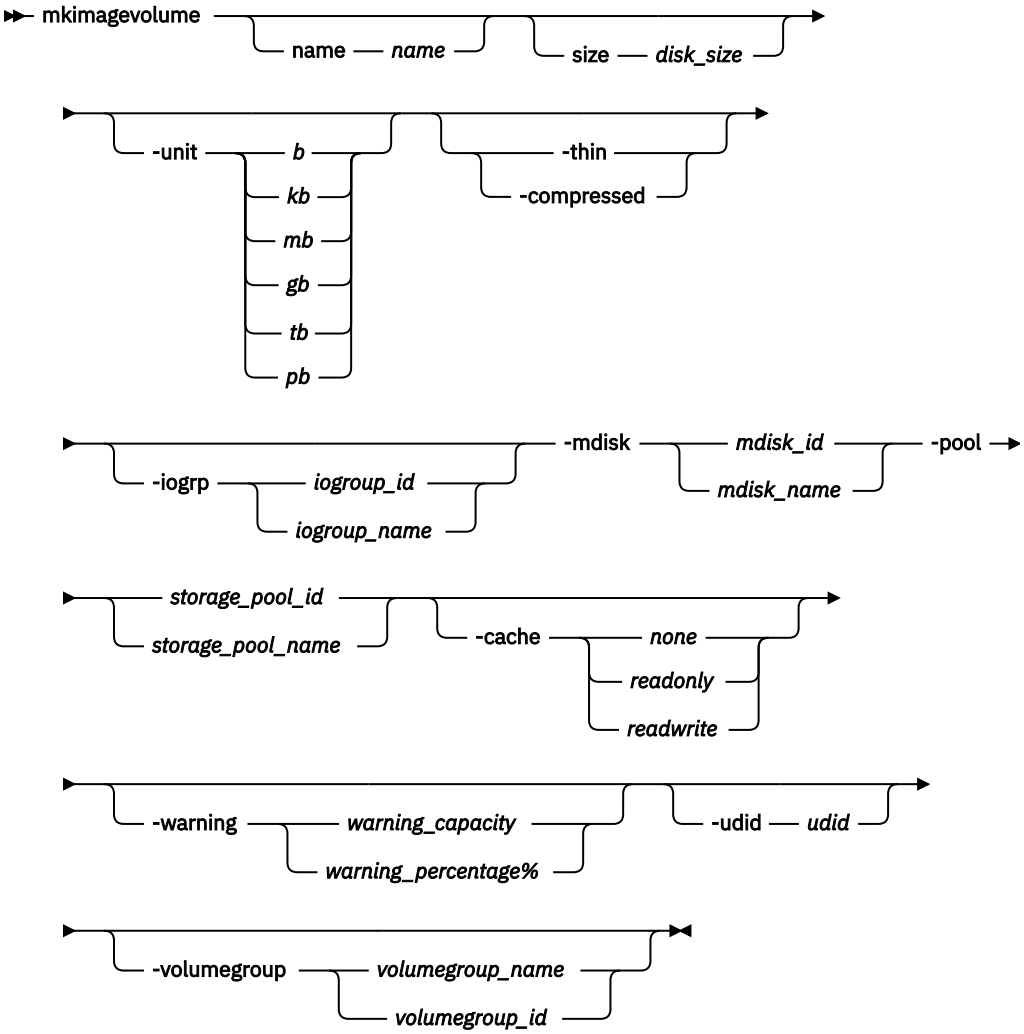
The resulting output:

```
Volume Group, id [0], successfully created
```

mkimagevolume

Use the **mkimagevolume** command to create an image mode volume by importing (preserving) data on a managed disk from another storage system.

Syntax



Parameters

-name *name*
(Optional) Specifies the name that is used for the volume that is created. This value must be an alphanumeric string 1 - 63 characters.

-size *disk_size*
(Required if **-thin** or **-compressed** is specified) Specifies the capacity of the volume, which is used with the value of the unit. The default capacity is in MB. When the unit of bytes is used, all capacities must be in multiples of 512 bytes. An entire extent is reserved even if it is only partially used.

Note: This parameter is optional if **-thin** or **-compressed** are not specified. If this parameter is not specified, the volume is full-provisioned.

For thin and compressed volumes, the real capacity is set from the MDisk size.

-unit *b | kb | mb | gb | tb | pb*
(Optional) Specifies the data units to use with the capacity that is specified by the **-size** parameter. The default unit type is *mb*.

-thin
(Optional) Specifies that the volume is to be created with thin provisioning. You cannot specify this parameter with **-compressed**. If you do not specify **-thin** or **-compressed**, the volume that is created is full-provisioned.

-compressed
(Optional) Specifies that the volume is to be created compressed. If the `-iogrp` parameter is not specified, the least used I/O group is used for compressed copies (considering the subset of I/O groups that support compression).

Remember: This command fails if no I/O groups support compression.

You cannot specify this parameter with **-thin**. If you do not specify **-thin** or **-compressed**, the volume that is created is fully-provisioned.

-iogrp ***iogroup_id | iogroup_name***
(Optional) Specifies the I/O group that the new volume is cached in.

-mdisk ***mdisk_idmdisk_name***
(Required) Specifies which currently unused MDisk to use to create the image mode volume.

-pool ***storage_pool_id | storage_pool_name***
(Required) Specifies the storage pool in which to create the new volume.

-cache ***none | readonly | readwrite***
(Optional) Specifies the caching options for the volume. Valid entries are:

- `readwrite` enables the cache for the volume (default)
- `readonly` disables write caching but allows read caching for the volume
- none disables caching for the volume

-warning ***warning_capacity | warning_percentage***
(Optional) Specifies a threshold at which a warning error log is generated for volume copies. A warning is generated when the used disk capacity on the thin-provisioned volume exceeds the specified threshold. You must specify either **-thin** or **-compressed** with this parameter. The default value is 80%.

-volumegroup ***volumegroup_name | volumegroup_id***
(Optional) Specifies the volume group that the volume should be added to.

Note: This command is not permitted for a volume group with a replication policy.

Description

Use the **mkimagevolume** command to create a new image mode volume. This command is used to import a volume, preserving existing data.

Note: A volume copy in a data reduction pool cannot be an image mode volume copy. The **mkimagevolume** command is not permitted if the pool specified is a Safeguarded backup location.

Create a fully-provisioned image mode volume in storage pool 0 with mdisk ID 2

```
mkimagevolume -mdisk 2 -pool 0
```

The resulting output:

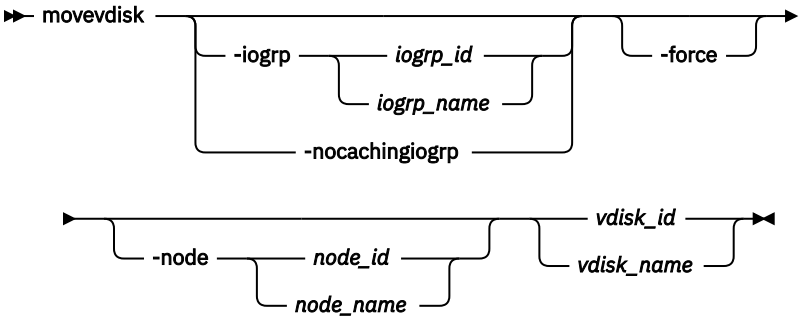
```
Volume, id [0], successfully created.
```

movevdisk

Use the **movevdisk** command to update the preferred node of a volume either within the same caching I/O group or to another caching I/O group.

Note: Multiple I/O groups have been replaced by FlashSystem grid as the method to scale out the storage system.

Syntax



Parameters

-iogrp ***iogrp_id | iogrp_name*** **IBM SAN Volume Controller only**
(Optional) Specifies the I/O group to move the volume to.

-nocachingiogrp
(Optional) If you specify this parameter, the volume has no caching I/O group. This option should be used under the direction of IBM Support only. The **-nocachingiogrp** parameter is mutually exclusive with the **-iogrp** and **-node** parameters.

-force
(Optional) If you specify the **-force** parameter, the contents of the cache are discarded and the volume might be corrupted. Use with caution.

-node ***node_id | node_name***
(Optional) Specifies the node ID or name that is assigned as the preferred node.

vdisk_id | vdisk_name
(Required) Specifies the volume to move.

Description

Use the **movevdisk** command to update the preferred node of a volume.

This command is not supported to change the I/O group if either copy is thin-provisioned or compressed and within a data reduction pool. The preferred node can be changed for volumes that are in data reduction pools.

A compressed volume can also be moved, and you can specify the preferred node in the new I/O group. You can move a volume that is in a FlashCopy mapping, but the FlashCopy bitmaps remain in the original I/O group. You cannot move volumes when the FlashCopy mapping is in `preparing` or `prepared` state. Additionally, a volume can be moved if it is the target of a FlashCopy mapping that is in `stopping` state.

If the volume is offline, use one of the **recovervdisk** commands to recover the volume and bring it back online. To specify a preferred node for the volume, use the **-node node_id | node_name** parameter with the **movevdisk** command. Use the **movevdisk** command to change the I/O group with which this volume is associated.

Important: Do not move:

- A volume to an offline I/O group under any circumstance.

Remember: To avoid data loss, make sure that the I/O group is online before you move the volume.

- An offline volume to the recovery I/O group.

Important: Do not move an offline volume to the recovery I/O group.

Remember: To avoid data loss, make sure that the I/O group is online before you move the volume.

You can migrate a volume to a new I/O group to manually balance the workload across the nodes in the clustered system. You might end up with a pair of nodes that are overworked and another pair of nodes that are underworked.

Remember: You cannot move a volume if that volume is being formatted.

For systems that support multiple I/O groups, if the volume is a target of a FlashCopy mapping with a source volume in an `active-active` relationship, then the new I/O group must be in the same site as the source volume. The system allows moving a volume in a remote copy relationship if the move does not change the I/O group (it changes the preferred node). If the volume is in an `active-active` relationship, the new I/O group must be in the same site as the source I/O group.

A volume in a volume group with a replication policy cannot be moved to a different I/O group. The preferred node for the volume can be changed if the new preferred node is in the same I/O group. When the preferred node is changed, the preferred node for the change volume will also be changed.

Modifying the I/O group that services the volume can be done concurrently with I/O operations if the host supports non- disruptive volume move. It also requires a rescan at the host level to ensure that the multipathing driver is notified that the allocation of the preferred node has changed and the ports by which the volume is accessed has changed. This can be done in the situation where one pair of nodes becomes over used.

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

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

The volume between I/O groups can be moved by using the **addvdiskaccess** command.

The following example move the DB_Volume to I/O group 2

```
movevdisk -iogrp 2 DB_Volume
```

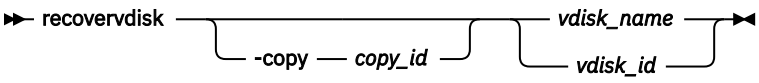
The resulting output:

```
No feedback
```

recovervdisk

Use the **recovervdisk** command to acknowledge volume data loss and brings the volume back online.

Syntax



Parameters

vdisk_name | vdisk_id
(Required) Specifies the volume to recover.

-copy *copy_id*
(Optional) Specifies the ID of the copy to recover.

Description

The specified volume, and all copies if mirrored, are recovered and brought back online. If the volume is thin-provisioned or has thin-provisioned copies, this command triggers the thin-provisioned repair process. If the volume is mirrored, the **recovervdisk** command triggers a resynchronization from a synchronized copy. The progress of the resynchronization can be monitored by using the **lsvdisksyncprogress** command. The volume remains online during the resynchronization process.

The **recovervdisk** command also starts the repair of any thin-provisioned copies that have a `fast_write_state` of corrupt. The progress of the repair process can be monitored by using the **lsrepairsevdiskcopyprogress** command.

A volume that is still offline because it is being repaired following the **recovervdisk** command has a `fast_write_state` of repairing. The volume is brought online when the repair process is complete.

An invocation example (to recover volume 45)

```
recovervdisk vdisk45
```

The following command is an invocation example to recover copy 0 of volume 45:

```
recovervdisk -copy 0 vdisk45
```

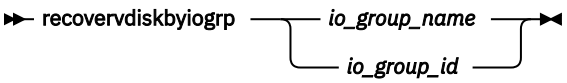
recovervdiskbycluster (Discontinued)

Attention: The **recovervdiskbycluster** command has been discontinued. Use the **recovervdiskbysystem** command instead.

recovervdiskbyiogrp

Use the **recovervdiskbyiogrp** command to acknowledge data loss for all volumes in the specified I/O group with a fast_write_state of corrupt and brings the volumes back online.

Syntax



Parameters

io_group_name | io_group_id
(Required) Specifies the I/O group for volume recovery.

Description

All volumes in the specified I/O group that have a fast_write_state of corrupt; and all copies, if mirrored, are recovered and brought back online. If any of the volumes are thin-provisioned or have thin-provisioned copies, the **recovervdiskbyiogrp** command triggers the thin-provisioned repair process. If volumes are mirrored, the command triggers a resynchronization from a synchronized copy. The progress of the resynchronization can be monitored by using the **lsrepairsevdiskcopyprogress** command. Volumes remain online during the resynchronization process.

If none of the volumes in the specified I/O group have a fast_write_state of corrupt, the **recovervdiskbyiogrp** command still starts the repair process for any corrupted copies of mirrored volumes. The progress of the repair process can be monitored by using the

lsrepairsevdiskcopyprogress command. If no corrupted volumes exist or no repairs to copies are needed, no error is returned.

Volumes that are still offline because they are being repaired following the **recovervdiskbyiogrp** command have a fast_write_state of repairing. Volumes are brought online when the repair process is complete.

An invocation example

```
recovervdiskbyiogrp iogrp2
```

The following output is displayed:

```
No feedback
```

recovervdiskbysystem

Use the **recovervdiskbysystem** command to acknowledge data loss for all volumes in the system with a fast_write_state of corrupt and bring the volumes back online.

Syntax



Parameters

There are no parameters for this command.

Description

Note: Use this command only under the supervision of IBM Support personnel.

All volumes in the system that have a `fast_write_state` of `corrupt`; and all copies, if mirrored, are recovered and brought back online. If any of the volumes are thin-provisioned or have thin-provisioned copies, the **recovervdiskbysystem** command triggers the thin-provisioned repair process. If volumes are mirrored, the command triggers a resynchronization from a synchronized copy. The progress of the resynchronization can be monitored by using the **lsvdisksyncprogress** command. Volumes remain online during the resynchronization process.

If none of the volumes in the system have a `fast_write_state` of `corrupt`, the **recovervdiskbysystem** command still starts the repair process for any corrupt copies of mirrored volumes. The progress of the repair process can be monitored using the **lsrepairsevdiskcopyprogress** command. If there are no corrupt volumes or no repairs to copies are required, no error is returned.

Volumes that are still offline because they are being repaired following the **recovervdiskbysystem** command have a `fast_write_state` of `repairing`. Volumes are brought online when the repair process is complete.

An invocation example

```
recovervdiskbysystem
```

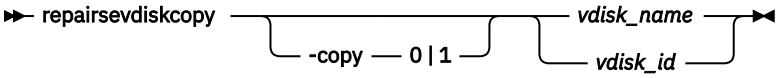
The resulting output:

```
No feedback
```

repairsevdiskcopy

The **repairsevdiskcopy** command repairs the metadata on a thin-provisioned volume copy or a compressed volume copy.

Syntax



Parameters

-copy 0 | 1
(Optional) Specifies the volume copy to repair.

vdisk_name | vdisk_id
(Required) Specifies the volume to repair.

Description

The **repairsevdiskcopy** command repairs the metadata on a thin-provisioned volume or compressed volume copy. Run this command only when you are directed by the fix procedures or by your product support information.

Running the command automatically detects corrupted metadata. The command holds the volume offline during the repair. During this time there are limited operations permissible.

If a repair operation completes successfully and the volume was previously offline because of corrupted metadata, the command brings the volume back online. The only limit on the number of concurrent repair operations is the number of volume copies in the configuration. Once started, a repair operation cannot be paused or canceled; the repair can only be ended by deleting the copy.

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

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

An invocation example

```
repairsevdiskcopy vdisk8
```

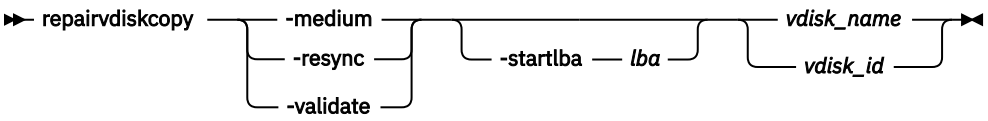
The resulting output:

```
No feedback
```

repairvdiskcopy

Use the **repairvdiskcopy** command on mirrored volumes to detect and (optionally) correct any volume copies that are not identical.

Syntax



Parameters

-medium
(Optional) Converts sectors that contain different readable data into virtual medium errors on the specified volume. It fixes preexisting medium errors found on only one volume copy by replacing them with data from the other volume copy. This parameter cannot be used with the `-validate` and `-resync` parameters. You must specify one of the three parameters.

-resync
(Optional) Corrects sectors that contain different readable data by copying contents from the primary volume copy to other copies on the specified volume. It fixes preexisting medium errors found on only one volume by replacing them with data from the other volume. This parameter cannot be used with the `-medium` and `-validate` parameters. You must specify one of the three parameters.

-validate
(Optional) Reports the first difference in readable data found on synchronized online copies of the specified volume, on or after the specified `-startlba` value. This parameter cannot be used with the `-medium` and `-resync` parameters. You must enter one of the three parameters.

-startlba *lba*
(Optional) Specifies a starting logical block address (LBA) on which to begin the command. The LBA must be specified in hex, with a 0x prefix.

vdisk_name* | *vdisk_id
(Required) Specifies the volume to repair. You must specify this parameter last on the command line.

Description

The **repairvdiskcopy** command detects and optionally, corrects any volume copies that are not identical. For the purposes of comparison, preexisting medium errors found on only one volume are ignored and fixed by replacing them with data from the other volume copy. The results are logged to the error log. The **-validate** parameter compares synchronized online copies of the specified volume. The **-medium** parameter changes any sectors that are not identical into virtual medium errors. The **-resync** parameter copies any sectors that are not identical to the other volume copies. You cannot use this command with a volume that is fast formatting.

You must specify **-validate**, **-medium**, or **-resync**.



Attention:

1. Before you run the **repairvdiskcopy** command, ensure that all volume copies are synchronized.
2. Only one **repairvdiskcopy** command can run on a volume at a time. You must wait for the **repairvdiskcopy** command to complete processing before running the command again.
3. Once you start the **repairvdiskcopy** command, you cannot use the command to stop processing.
4. The primary copy of a mirrored volume cannot be changed while the **repairvdiskcopy -resync** command is running.

Use the **-startlba** parameter to specify a starting Logical Block Address (LBA). Enter an LBA value from `0` to full disk size minus one. The parameter logs the first error found and then stops the command. By repeating this parameter, you can collect all of the instances where the volume copies are not identical.

During **repairvdiskcopy** command operation, the volume remains online. The I/O and synchronization operations are allowed while the command is in progress.

The rate for the **repairvdiskcopy** command is controlled by the synchronization rate of the volume that is being repaired. To suspend the repair process, set the synchronization rate of the volume to `0` using the **chvdisk** command.

Use the **lsrepairvdiskcopyprogress** command to display the progress of mirrored volume validation and repairs. You can specify a volume copy by using the **-copy** *id* parameter. To display the volume that

has two or more copies with an active task, specify the command with no parameters; it is not possible to have only one volume copy with an active task.

An invocation example

```
repairvdiskcopy -resync -startlba 0x0 vdisk8
```

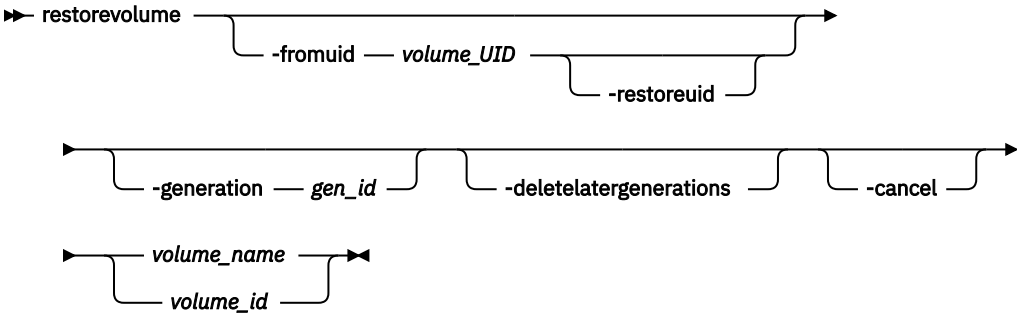
The resulting output:

```
No feedback
```

restorevolume

Use the **restorevolume** command to restore a volume for cloud tiering.

Syntax



Parameters

-fromuid ***volume_UID***
(Optional) Specifies the volume snapshot to restore (specified by volume UID). The value must be a hexadecimal number.

Use this parameter to restore a snapshot from a different volume. This means that the specified UID must be different from the UID of the volume being restored (when you specify *volume_name* or *volume_id*).

Note: The volume being restored to cannot have cloud snapshot enabled if you specify this parameter.

-restoreuid
(Optional) Specifies the UID of the restored volume should be set to the UID of the volume snapshot that is being restored. You must specify **-fromuid** with this parameter. The **-restoreuid** parameter is not supported if cloud account is in import mode.

-generation ***gen_id***
(Optional) Specifies the snapshot generation to restore. The value must be a number.

-deletelatergenerations
(Optional) Specifies that all snapshot generations should be deleted after the generation is restored.

Note: This parameter is required when the volume has cloud snapshot enabled and the generation being restored is not the most recent snapshot of the volume.

-cancel
(Optional) Cancels the restore operation.

volume_name | volume_id
(Required) Specifies the volume name or ID to restore. The value for the volume ID must be a number and the value for the volume name must be an alphanumeric string.

Note: Restoring a cloud backup to a volume other than the backup source volume is not allowed if the target volume has snapshots.

Description

This command restores a volume from a snapshot generation.

The restore operation is performed directly on the volume specified by volume name or volume ID (without use of a temporary volume). The volume is `offline` while the restore operation is in progress. If the restore process is canceled before it completes the data on the volume is inconsistent and not usable.

An invocation example

To restore an earlier generation (generation 3) for volume `volume7`:

```
restorevolume -generation 3 -deletelatergenerations volume7
```

The resulting output:

```
No feedback
```

rmmetadatavdisk

Use the **rmmetadatavdisk** command to remove a metadata volume.

Syntax



Description

This command removes the metadata volume from the system.

An invocation example

```
rmmetadatavdisk
```

The resulting output:

```
No feedback.
```

rmvdisk

Use the **rmvdisk** command to delete a volume. You cannot use this command for volumes that are replicated or in storage partitions.

Syntax



Parameters

-force
(Optional) The specified volume is to be deleted, even if mappings still exist between this volume and one or more hosts. Host-to-volume mappings and any FlashCopy mappings that exist for this volume are deleted.

Important: If you stop a FlashCopy mapping that has dependent FlashCopy mappings, the dependent mapping target volumes might become unusable.

To determine your dependent FlashCopy mappings before you use the **-force** parameter, run the **lsfcmapdependentmaps** command.

Important: To prevent an active volume from being deleted unintentionally, you can use a global system setting to enable volume protection (see parameters **-vdiskprotectionenabled** and **-vdiskprotectiontime** in the **chsystem** command). You can specify a time period for which the volume must be idle before you can delete it. If volume protection is enabled and the time period is not expired, the volume deletion fails even if the **-force** or **-removehostmappings** parameters are used.

If the **-force** deletion of a volume causes dependent mappings to be stopped, any target volumes for those mappings that are in relationships are also stopped. The dependent mappings can be identified by using the **lsvdiskdependentmaps** command on the volume that you want to delete.

Note: Using the **-force** parameter might result in a data loss. Use it only under the direction of your product support information, or if you are willing to accept the risk of volume data loss.

If you do not specify this parameter, a volume cannot be deleted while a backup operation is in progress. Additionally, a volume that contains image mode copies cannot be deleted while a restore operation is in progress (if the volume contains inconsistent data).

-removehostmappings
(Optional) Removes all host mappings for the specified volume before the volume is deleted.

Note: Using the **-removehostmappings** parameter might result in a data loss. Use it only under the direction of your product support information, or if you are willing to accept the risk of volume data loss.

Important: To prevent an active volume from being deleted unintentionally, you can use a global system setting to enable volume protection (see parameters **-vdiskprotectionenabled** and **-vdiskprotectiontime** in the **chsystem** command). You can specify a time period for which the volume must be idle before you can delete it. If volume protection is enabled and the time period is not expired, the volume deletion fails even if the **-force** or **-removehostmappings** parameters are used.

vdisk_id | vdisk_name
Specifies the name of the volume to delete, either by ID or by name.

Description

This command deletes an existing managed mode volume or an existing image mode volume. The extents that made up this volume are returned to the pool of free extents that are available on the storage pool, if the volume is in managed mode.

Remember: If you run this command, any data that was on the volume is lost. Before you run this command, ensure that the volume (and any data that resides on it) is no longer required.

This command is unsuccessful if:

- Volume protection is enabled (by using the command **chsystem**) and the volume that is being removed received I/O within the defined volume protection time period.
- The data reduction pool is corrupted.

Deleting a managed mode volume

When you use this command to delete a managed mode volume, all the data on the volume is deleted. The extents that make up the volume are returned to the pool of free extents that are available in the storage pool.

If host mappings exist for the volume, or if any FlashCopy mappings would be affected, the deletion fails. You can use the **-force** parameter to force the deletion. If you use the **-force** parameter, mappings

that have the volume as source or target are deleted, other mappings in a cascade might be stopped, and then the volume is deleted. The **-force** parameter also deletes relationships that exist for the specified volume (and any information that is not staged in the fast write cache).

If the volume is in the process of migrating to an image mode volume (by using the **migratetoimage** command), the deletion fails unless you use the **-force** parameter. If you use the **-force** parameter, the migration is halted and then the volume is deleted. Before you run this command, ensure that the volume (and any data that resides on it) is no longer required.

Deleting an image mode volume

If the volume is mirrored and one or both copies is in image mode, you must first wait for all fast-write data to be moved to the controller logical unit. This pause ensures that the data on the controller is consistent with the data on the image mode volume before the volume is deleted. This process can take several minutes to complete, and is indicated by the fast_write_state of the volume, which is empty. If the **-force** parameter is specified, the fast-write data is discarded and the volume is deleted immediately; the data on the controller logical unit is left inconsistent and unusable. If the copies are not synchronized, you must use the **-force** parameter.

If you run the command while data is in the cache, the system attempts to move the data out of the cache; this process can time out, however.

If any virtual medium errors exist on the volume, the command fails. You can force the deletion by using the **-force** parameter; however, using **-force** 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 a medium error was found. At that moment, you must have two identical copies of data and you must then simulate a medium error on the target disk. You can simulate a medium error on the target disk by creating a virtual medium error on the target disk.

If FlashCopy mappings or host mappings exist for the volume, the deletion fails unless you use the **-force** parameter. If you use the **-force** parameter, mappings are deleted and the volume is deleted. If any data is not staged in the fast write cache for this volume, the deletion of the volume fails. When the **-force** parameter is specified, any data that is not staged in the fast write cache is deleted. Deleting an image mode volume causes the managed disk that is associated with the volume to be removed from the storage pool. The mode of the managed disk is returned to unmanaged.

An invocation example

```
rmvdisk -force vdisk5
```

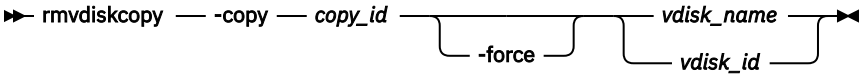
The resulting output:

```
No feedback
```

rmvdiskcopy

Use the **rmvdiskcopy** command to remove a volume copy from a mirrored volume. Use the **rmvolumecopy** command for high availability volumes.

Syntax



Parameters

-copy *copy_id*
(Required) Specifies the ID of the copy to delete.

-force
(Optional) Forces the deletion of the last synchronized copy of a volume, which deletes the entire volume.

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

The parameter also forces the deletion of a nonmirrored volume, a copy that is migrating to image mode, or an image-mode copy that has virtual medium errors.

Important: To prevent an active volume from being deleted unintentionally, you can use a global system setting to enable volume protection (see the **chsystem** command). You can specify a time period for which the volume must be idle before you can delete it. If volume protection is enabled and the time period has not expired, the volume deletion fails even if the **-force** parameter is used.

Note: When you use the **-force** parameter to delete a synchronized copy from a mirrored volume, where the synchronized copy is in a storage pool with volume protection inactive, but it has an unsynchronized copy in a different storage pool with volume protection active, then volume protection applies to the volume. It is because the unsynchronized copy is also deleted and it is in a storage pool with volume protection active, as the **-force** is deleting both copies in a mirrored volume the command fails if host I/O is being sent to the volume in the defined volume protection time period.

You cannot use this parameter with volumes that are used in policy-based replication.

vdisk_name | vdisk_id
(Required) Specifies the volume to delete the copy from. You must specify this parameter last on the command line.

Description

The **rmvdiskcopy** command deletes the specified copy from the specified volume. The command fails if all other copies of the volume are not synchronized; in this case, you must specify the **-force** parameter, delete the volume or more, or wait until the copies are synchronized.

Remember: This command is unsuccessful if:

- Volume protection is enabled.
- The last volume copy being deleted has received I/O within the defined volume protection time period.

These changes apply to this command only when deleting the last synchronized copy of a volume or removing the entire volume.

The volume copy can be removed from the volume that contains snapshots only if the parent pool of the remaining volume copy is the same as the parent volume.

An invocation example

This example forces a deletion.

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

```
rmvdiskcopy -copy 0 -force 134
```

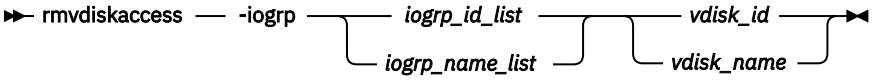
The resulting output:

```
No feedback
```

rmvdiskaccess

Use the **rmvdiskaccess** command to delete one or more I/O groups from the set of I/O groups in which a volume can be made accessible to hosts.

Syntax



Parameters

-iogrp *iogrp_id_list | iogrp_name_list*
(Required) Specifies a list of I/O groups to remove from the I/O group access set of the volume.

vdisk_id | vdisk_name
(Required) Specifies the volume from which to remove access I/O groups.

Description

The **rmvdiskaccess** command removes I/O groups from the volume access set. However, it cannot remove all I/O groups from the access set; a volume must have at least one I/O group in an access set. When an I/O group is removed from the access set, all host mappings created through that I/O group (for the volume) are deleted. Consequently, you cannot access the volume through any related I/O group nodes.

Remember: If an I/O group in the list is not in the access set, no error is generated, but no action is taken for that I/O group.

An invocation example to remove I/O groups 2 and 3 from the volume access set for volume ID 3

```
rmvdiskaccess -iogrp 2:3 3
```

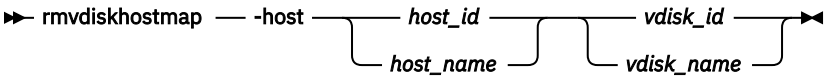
The resulting output:

```
No feedback
```

rmvdiskhostmap

Use the **rmvdiskhostmap** command to delete an existing host mapping the volume is no longer accessible for input/output (I/O) transactions on the given host.

Syntax



Parameters

-host *host_id | host_name*
(Required) Specifies the host that you want to remove from the map with the volume, either by ID or by name.

vdisk_id | vdisk_name
(Required) Specifies the name of the volume that you want to remove from the host mapping, either by ID or by name.

Description

This command deletes an existing mapping between the specified volume and the host. This effectively stops the volume from being available for I/O transactions on the given host.

This command also deletes a Small Computer System Interface (SCSI) or persistent reservation that a host has on a volume. Once the reservation is removed, a new host is allowed to access the volume in the future because the original host no longer has access.

Note: The **rmvdiskhostmap** command deletes the host mapping for all I/O groups in the access I/O group set of the volume.

Use caution when you process this command because to the host, it seems as if the volume has been deleted or is offline.

Remember: This command is unsuccessful if:

- Volume protection is enabled
- The host mapping being deleted is mapped to any volume that has received I/O within the defined volume protection time period

An invocation example

```
rmvdiskhostmap -host host1 vdisk8
```

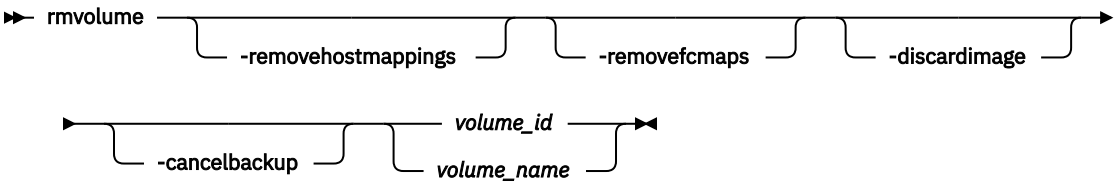
The resulting output:

```
No feedback
```

rmvolume

Use the **rmvolume** command to remove a volume.

Syntax



Parameters

-removehostmappings
(Optional) Allows a volume to be deleted even if host mappings are removed and this volume is deleted.

-removefcmaps
(Optional) Allows a volume to be deleted even if it is a part of a FlashCopy mapping and regardless of the state of the mappings.

-discardimage
(Optional) Allows a volume to be deleted even if data on an image mode copy cannot be made consistent.

Important: Using this parameter might result in data loss. Use it only under the direction of your product support information, or if you are willing to accept the risk of data loss on the volume.

-cancelbackup
(Optional) Allows a volume to be deleted even if a backup operation is in progress.

Important: Using this parameter might result in data loss. Use it only under the direction of your product support information, or if you are willing to accept the risk of data loss on the volume.

volume_id | volume_name
(Required) Specifies the volume to be removed.

Description

Use the **rmvolume** command to remove a volume.

Important: Using this command deletes all copies of a volume and it cannot be undone. If the volume is replicated, the recovery copy of the volume will also be removed unless the recovery volume group has independent access enabled.

An invocation example to remove a volume

```
rmvolume 0
```

The detailed resulting output:

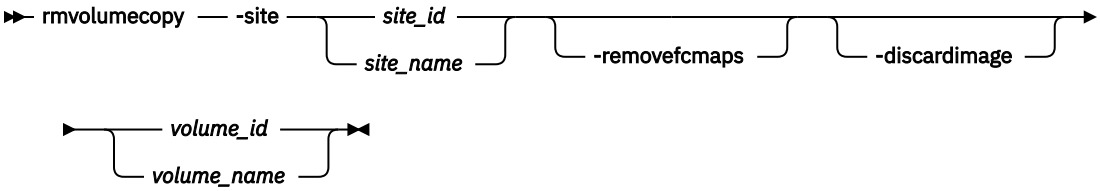
```
No feedback
```

rmvolumecopy

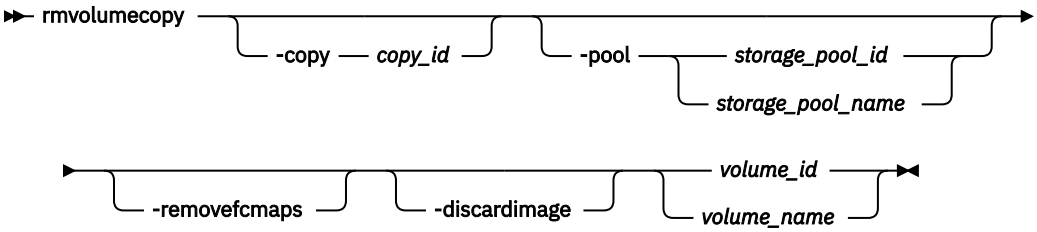
Use the **rmvolumecopy** command to remove a volume copy from a volume.

Syntax

This syntax diagram specifies a volume copy by site.



This syntax diagram specifies a volume copy by copy ID or storage pool. You must specify either **-copy** or **-pool**.



Parameters

-site ***site_id | site_name***
(Optional) Specifies the site that the volume copy is removed from. You cannot specify this parameter if you specify **-copy** or **-pool**.

-pool ***storage_pool_id | storage_pool_name***
(Optional) Specifies the storage pool that the volume copy is removed from.

-copy *copy_id*
(Optional) Specifies the copy ID for the volume copy to be deleted. The value is 0 or 1. This keyword cannot be specified when **-site** is specified.

Important: If the volume has copies in multiple sites you must specify **-pool** to identify the volume copy to remove.

.

-removefcmaps
(Optional) Allows a volume copy to be deleted even if it is part of a FlashCopy mapping.

-discardimage
(Optional) Allows a volume copy to be deleted even if data on an image mode copy cannot be made consistent.

Important: Using this parameter might result in data loss. Use it only under the direction of your product support information, or if you are willing to accept the risk of data loss on the volume.

volume_id | volume_name
(Required) Specifies the volume ID or name for the volume copy to remove.

Description

Use the **rmvolumecopy** command to remove a copy of a volume.

This command fails if a backup operation is in progress.

An invocation example to remove a volume copy when there are two copies in the same storage pool.

```
rmvolumecopy -pool 5 -copy 1 volume5
```

The detailed resulting output:

```
No feedback
```

An invocation example to remove a volume copy with FlashCopy mappings.

```
rmvolumecopy -site 1 -removefcmaps 1
```

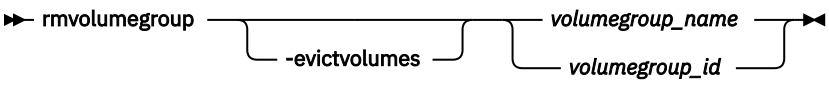
The detailed resulting output:

```
No feedback
```

rmvolumegroup

Use the **rmvolumegroup** command to remove a volume group from a system.

Syntax



Parameters

-evictvolumes
(Optional) This allows the user to remove all volumes from a group before deleting the group. This is only permitted for volume groups that do not have a replication policy assigned.

volumegroup_name | volumegroup_id
(Required) Specifies the volume group name or ID for the volume to remove from the system. The value for the volume group ID must be a number and the value for the volume group name must be an alphanumeric string.

Description

This command removes a volume group from the system. If the volume group contains active volumes, it cannot be removed from the system. Only empty volume groups can be deleted from the system. If an empty volume group is used for policy-based replication, you must remove the volume group from the production system and not the recovery system.

Note: This command can run whether the remote systems in the replication policy are connected or disconnected. If the command succeeds, the volume group is removed immediately on the local system. When the connectivity is restored, the volume group is removed asynchronously on the remote systems.

An invocation example

```
rmvolumegroup Vardy1
```

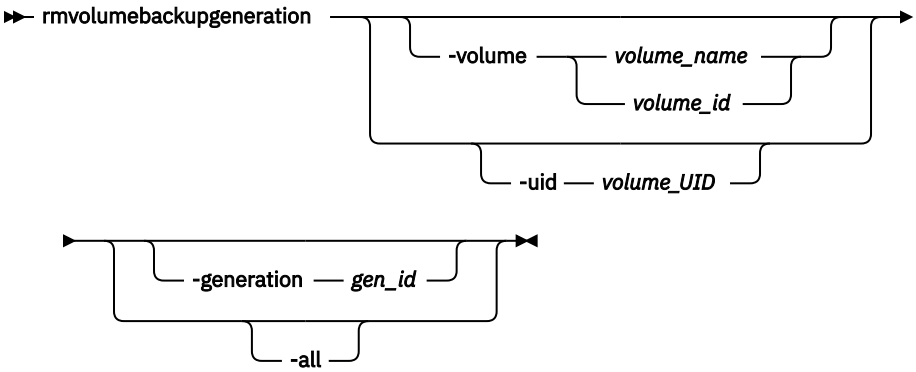
The resulting output

```
No feedback.
```

rmvolumebackupgeneration

Use the **rmvolumebackupgeneration** command to delete a volume backup or cancel a volume snapshot operation that is in progress.

Syntax



Parameters

-volume ***volume_name | volume_id***
(Optional) Specifies the volume snapshot by name or ID.

Note: The volume must exist on the local system.

The value for the volume name must be an alphanumeric string and the value for the volume ID must be a number. This parameter is mutually exclusive with **-uid**.

-uid ***volume_UID***
(Optional) Specifies the volume snapshot UID. The value for a volume UID must be a number 0 - 32. This parameter is mutually exclusive with **-volume**.

-generation *gen_id*
(Optional) Specifies the snapshot generation to be deleted for the volume. Only a single snapshot generation is deleted. If the specified generation is for a snapshot operation that is in progress, that snapshot operation is canceled. The value for the generation ID must be a number. This command is mutually exclusive with **-all**.

-all
(Optional) Specifies deletion for all volume snapshots (which cancels all snapshot generations as well). This command is mutually exclusive with **-generation**.

Description

This command deletes a volume snapshot or cancels a volume snapshot operation that is in progress.

Note: If the:

- Command completes immediately, the delete operation is performed asynchronously
- Volume has more than one snapshot generation, you cannot delete the most recent snapshot generation

An invocation example

To delete snapshot generation 22 for volume with the UID 600507680CA880DF1800000000000007:

```
rmvolumebackupgeneration -uid 600507680CA880DF1800000000000007 -generation 22
```

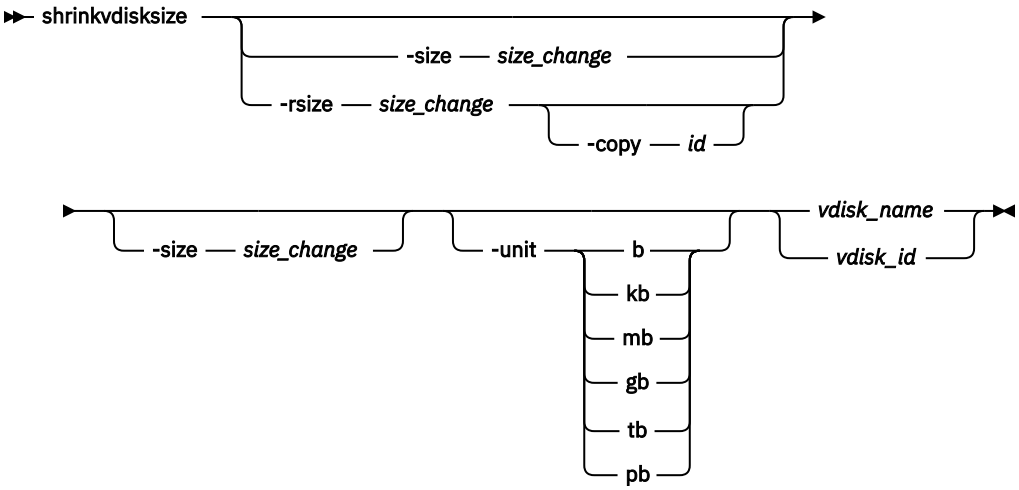
The resulting output:

```
No feedback
```

shrinkvdisksize

Use the **shrinkvdisksize** command to reduce the size of a volume by the specified capacity.

Syntax



Parameters

-size *size_change*
(Optional) Specifies the size reduction (change in size) for the designated volume. The **-size** parameter cannot be used with the **-rsize** parameter. You must specify either **-size** or **-rsize**.

Important: This parameter does reduce the size of a volume (the specified virtual size capacity).

-rsize *size_change*
(Optional) Reduces the real size of a thin-provisioned volume by the specified amount. It indicates the change in size as a result of the reduction. Specify the *size_change* value by using an integer. Specify the units for a *size_change* integer by using the **-unit** parameter; the default is MB. You must specify either **-rsize** or **-size**.

Remember: You cannot use **-rsize** to resize a thin-provisioned or compressed volume copy that is in a data reduction pool.

-copy *id*
(Optional) Specifies the copy to change the real capacity for. You must also specify the **-rsize** parameter. If the **-copy** parameter is not specified, all copies of the volume are reduced. This parameter is required if the volume is mirrored and only one copy is thin-provisioned.

-unit b | kb | mb | gb | tb | pb
(Optional) Specifies the data units to be used along with the value that is specified by the **-size** parameter.

vdisk_name | vdisk_id
(Required) Specifies the volume that you want to modify, either by ID or by name.

Description

The **shrinkvdisksize** command reduces the capacity that is allocated to the particular volume by the amount that you specify. You cannot shrink the real size of a thin-provisioned volume below its used size. All capacities, including changes, must be in multiples of 512 bytes. An entire extent is reserved even if it is only partially used. The default capacity units are MB. You can use **shrinkvdisksize** while the volume is fast formatting and while the volume is synchronizing.You cannot use **shrinkvdisksize** if the volume is fast formatting.

The command can be used to shrink the usable capacity that is allocated to a particular volume by the specified amount. The command can also be used to shrink the provisioned capacity of a thin-provisioned volume without altering the usable capacity that is assigned to the volume. To change the capacity of a non-thin-provisioned disk, use the **-size** parameter. To change the real capacity of a thin-provisioned disk, use the **-rsize** parameter. To change the provisioned capacity of a thin-provisioned disk, use the **-size** parameter.

Volumes can be reduced in size, if required.

When the virtual size of a thin-provisioned volume is changed, the warning threshold is automatically scaled to match. The new threshold is stored as a percentage.



Attention: If the volume contains data that is being used, do not shrink the volume without backing up the data first.

The system reduces the capacity by removing extents that are allocated to the end of the volume. You cannot control how the host used the capacity that was provisioned to it and so you cannot assume that the unused space is all at the end of the volume.

Remember: Before you shrink a volume, validate that the volume is not mapped to any host objects.

You can determine the exact capacity of the source or master volume by issuing the **lsvdisk -bytes** *vdiskname* command. Shrink the volume by the required amount by issuing the **shrinkvdisksize -size** *size_change***-unit** b | kb | mb | gb | tb | pb *vdisk_name | vdisk_id* command.

Remember:

1. You cannot resize (shrink) an image mode volume.
2. You cannot resize (shrink) the disk if the volume contains data.
3. You cannot resize (shrink) a volume if that volume is being formatted.
4. You cannot resize (shrink) a volume that is being migrated.
5. You cannot resize (shrink) a volume if cloud snapshot is enabled on that volume.
6. You cannot resize (shrink) a volume if it is replicated or in a highly-available storage partition.

A volume cannot be shrunk if it is the target of a FlashCopy mapping. If the volume is the a source of a FlashCopy mapping it can be shrunk provided no FlashCopy mapping for which it is the source has a starting capacity that is greater than the proposed size of the volume. The starting capacity can be found in the **lsfcmap** view.

An invocation example to decrease the capacity of vdisk1 by 2 KB

```
shrinkvdisksize -size 2048 -unit b vdisk1
```

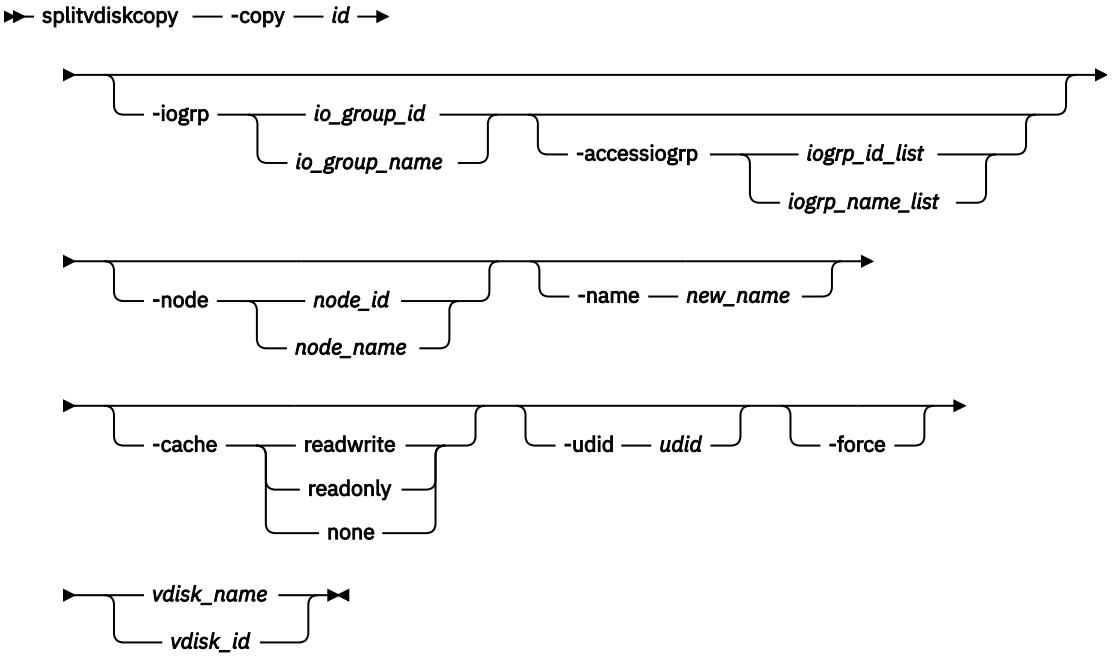
The resulting output:

```
No feedback
```

splitvdiskcopy

Use the **splitvdiskcopy** command to create a separate volume from a synchronized copy of a mirrored volume.

Syntax



Parameters

-copy *id*
(Required) Specifies the ID of the copy to split.

-iogrp *io_group_id | io_group_name*
(Optional) Specifies the I/O group to add the new volume to. The default is the I/O group of the specified volume.

-accessiogrp *iogroup_id_list | iogroup_name_list*
(Optional) Specifies which I/O groups provide access to the volume. If the **-accessiogrp** parameter is used, the **-iogrp** parameter must also be specified. The specified volume access set must include the caching I/O group. If the **-accessiogrp** parameter is not specified, the new volume is assigned its caching I/O group as the only I/O group that provides access, where the caching I/O group is assigned by the system or by using **-iogrp**.

-node *node_id | node_name*
(Optional) Specifies the preferred node ID or the name for I/O operations to this volume. You can use the **-node** parameter to specify the preferred access node.

-name *new_name*
(Optional) Assigns a name to the new volume.

-cache readwrite | readonly | none
(Optional) Specifies the caching options for the new volume. (Optional) Specifies the caching options for the volume. Valid entries are:

- `readwrite` to enable the cache for the volume
- `readonly` to disable write caching while allowing read caching for a volume
- none to disable the cache mode for the volume

The default is `readwrite`.

Remember: If you do not specify the **-cache** parameter, the default value (`readwrite`) is used.

-udid *udid*
(Optional) Specifies the *udid* for the new volume. The *udid* is a required identifier for OpenVMS hosts; no other hosts use this parameter. Supported values are a decimal number `0 - 32767`, or a hexadecimal number `0 - 0x7FFF`. A hexadecimal number must be preceded by `0x`; for example, `0x1234`. The default *udid* value is `0`.

-force
(Optional) Allows the split to proceed even when the specified copy is not synchronized, or even when the cache flush is likely to fail. The newly created volume might not be consistent.

Description

The **splitvdiskcopy** command creates a new volume in the specified I/O Group from a copy of the specified volume.

For thin or compressed volumes that are in data reduction storage pools:

- You cannot specify a cache mode of none or `readonly`.
- You cannot specify an I/O group that is different from the current I/O group that is associated with the volume.

If the copy that you are splitting is not synchronized, you must use the **-force** parameter. The command fails if you are attempting to remove the only synchronized copy. To avoid command failure, wait for the copy to synchronize or split the unsynchronized copy from the volume by using the **-force** parameter. You can run the command when either volume copy is offline.

An invocation example for creating a volume with I/O groups 2 and 3 in its I/O group access set

```
splitvdiskcopy -copy 1 -iogrp 2 -node 7 -accessiogrp 2:3 DB_Disk
```

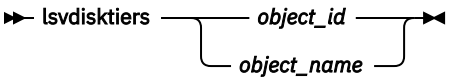
The resulting output:

```
Virtual Disk, id [X], successfully created.
```

lsvdisktiers

Use the **lsvdisktiers** command to list the volume tier capacities for the IBM Storage Insights.

Syntax



Parameters

-delim ***delimiter***
(Optional) Specify a delimiter to separate data in the output.

-nohdr
(Optional) Suppress the headings in the output.

-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.

-filtervalue?
(Optional) Displays a list of valid filter attributes.

object_id | object_name
(Required) Specifies the name or ID of an object.

Description

This table provides the attribute values that can be displayed as output view data.

Table 170. ***lsvdisktiers*** *attribute values*

Attribute	Value
vdisk_id	Indicates the ID of the volume.
vdisk_name	Indicates the name of the volume.
copy_id	Indicates the ID of the volume copy. The value can be 0 or 1.
mdisk_grp_id	Indicates the ID of the storage pool that the volume copy belongs to.
mdisk_grp_name	Indicates the name of the storage pool that the volume copy belongs to.
tier_scm_capacity	Indicates the total MDisk capacity that is assigned to the volume in tier_scm.
tier0_flash_capacity	Indicates the total MDisk capacity that is assigned to the volume in tier0_flash.
tier1_flash_capacity	Indicates the total MDisk capacity that is assigned to the volume in tier1_flash.
tier_enterprise_capacity	Indicates the total MDisk capacity that is assigned to the volume in tier_enterprise.
tier_nearline_capacity	Indicates the total MDisk capacity that is assigned to the volume in tier_nearline.

Chapter 46. Command-line interface messages

The messages that display when you use the command-line interface (CLI) are listed here.

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 an error code is sent to standard error. If the command succeeds, but the cluster is operating near its licensed virtualization limit, the return code is still `0`, and an event is logged in the event log.

When you issue a **`create`** command, the message ID that is assigned to the new object is returned as part of the success message that is sent to standard output. If the **`-quiet`** parameter is used, only the message ID is sent.

CMMVC1001E **The command failed because not enough volume copy IDs are available in the system.**

Explanation

Adding a snapshot of a single parent volume requires a volume copy ID for each specified pool or for all volume copies in distinct pools of the parent volume (if no pools are specified). This problem occurs when not enough copy IDs are available in the system.

User response

Delete resources to free volume copy IDs in the system, and then retry the command.

CMMVC1002E **The command failed because more pools were specified than copies exist in distinct pools for one or more parent volumes.**

Explanation

For a snapshot, you can't specify more pools than the number of copies that exist in distinct pools of the parent volume.

User response

Run the command again with fewer (or no) pools specified.

CMMVC1003E **The command failed as the specified systems are already configured to use an incompatible replication topology.**

Explanation:
A partnership can only be used for either asynchronous or HA replication, not both.

User response:
Either change the topology to match the existing topology in use by the partnership, or delete any replication policies with a conflicting topology before trying again.

CMMVC1004E **The command failed because parameter clusterip/clusterip_6 is mandatory to configure a new system IP.**

Explanation:
New system IP address cannot be configured with IP address.

User response:
Use **`clusterip`** to configure IPv4 system IP address or **`clusterip_6`** to configure IPv6 system IP address.

CMMVC1005E **The command failed as the existing I/O Group is configured for HA Replication.**

Explanation

HA replication can only be configured on single I/O Group systems.

User response

Remove the HA replication, or continue to use a single I/O Group system.

CMMVC1007E **The command failed because a VDisk that is used to support a snapshot cannot have multiple copies in the same pool.**

Explanation

The volume with snapshots cannot have multiple copies in the same parent pool unless the **`-autodelete`** parameter is specified for one copy when synchronization is completed.

User response

Run the command again with a different pool or specify the **`-autodelete`** parameter.

© Copyright IBM Corp. 2003, 2025

CMMVC1008E **The command failed because a volume with snapshots cannot have multiple copies in the same pool.**

Explanation

A volume or volume with snapshots cannot have multiple copies in the same parent pool unless one copy is set to autodelete when synchronization is completed.

User response

Run the command again with a different pool or specify the **-autodelete** parameter.

CMMVC1009E **The command failed because a volume is in a Remote Copy relationship.**

Explanation:
One of more of the volumes to be restored is part of either a hyperswap, a MM / GM, or a 3site relationship. Snapshot restore is not supported for these products.

User response:
In order to restore the volume(s) they will need to be removed from any hyperswap, MM / GM, or 3site relationship.

CMMVC1011E **The task cannot be initiated because the name of this truststore cannot be changed.**

Explanation:
This error can occur when you attempt to change the name of the common CA database truststore.

User response:
No action needed. You cannot change the name of the common CA database truststore.

CMMVC1012E **The task cannot be initiated because this truststore cannot be removed.**

Explanation:
This error can occur when you attempt to remove the common CA database truststore.

User response:
No action needed. You cannot remove the common CA database truststore.

CMMVC1013E **The command failed because an associated storage partition is preventing this replication policy from being removed.**

Explanation

A replication policy can only be removed when it is not assigned to any storage partition.

User response

Remove the replication policy from the storage partitions it is currently assigned to and resubmit the command.

CMMVC1014E **The task cannot be initiated because the specified tags are invalid for this truststore.**

Explanation:
This error can occur when you attempt to set any other tags on the common CA database truststore.

User response:
You cannot set any other tags on the common CA database truststore.

CMMVC1015E **The command failed because the storage partition preferred management system can only be changed for a storage partition that has a replication policy.**

Explanation

It is only possible to change the preferred management system for a storage partition that has a replication policy. This operation will fail for a single-system storage partition with no replication policy assigned.

User response

Only use this command to change the preferred management system for a storage partition that has a replication policy assigned.

CMMVC1016E **The command failed because this type of metadata volume already exists.**

Explanation:
This error occurs when you try to create a new version 1 metadata volume while the system already has a version 1 metadata volume configured.

User response:
No action required.

CMMVC1017E **The action failed as an existing metadata volume already exists and -migrate was not specified.**

Explanation:
This error occurs when you try to create a new version 2 metadata volume while an existing version 1

metadata volume is already configured on the system. The **-migrate** parameter must be specified to create a new version 2 metadata volume and migrate the existing metadata to the new volume.

User response:
Retry the command with the **-migrate** parameter to migrate the existing metadata to the new volume.

CMMVC1018E **The action failed as an existing version 1 metadata volume does not exist.**

Explanation:
This error occurs when you specify that the metadata must be migrated to a new metadata volume while an existing version 1 volume metadata doesn't exist on the system.

User response:
No action required.

CMMVC1019E **The action failed as a version 2 metadata volume already exists.**

Explanation:
This error occurs when you try to create a new metadata volume while there is already a version 2 metadata volume in the system.

User response:
No action required.

CMMVC1020E **The action failed as the specified pool is not a parent pool.**

Explanation:
This error occurs when you try to create a new metadata volume in a child pool. A metadata volume can only be created in a parent pool.

User response:
Retry the command, specifying a parent pool.

CMMVC1021E **The action failed as a migration from the old metadata volume to the new metadata volume is in progress.**

Explanation:
This error occurs when you try to run the **rmmetadatavdisk** command while there is more than one metadata volume configured on the system.

User response

Complete the following steps to fix this error:

- Wait for the migration process to complete.
- If the migration process fails, contact IBM Support Center.

CMMVC1022E **The action failed because there are vVol volumes that depend on the metadata volume and -ignoreexistingvvols was not specified.**

Explanation:
This error occurs when you try to remove a metadata volume on which other volumes are dependent. Removing such metadata volumes impacts other dependent volumes.

User response

Complete the following steps to fix this error:

- Remove the vVol volumes on this system, then retry the **rmmetadatavdisk** command.
- Alternatively, specify the **-ignoreexistingvvols** parameter when you retry this command. The **-ignoreexistingvvols** parameter ensures that the metadata volume is removed but the dependent vVol volumes are retained.

CMMVC1023E **The command failed as the metadata volume is offline.**

Explanation:
This error occurs when you try to create a new VMware Virtual Volume (vVol) while the metadata volume is offline.

User response:
Use the **lsmetadatavdisk** command to display the metadata volume information. Identify and resolve any errors that affect the metadata volume, then retry the command.

CMMVC1024E **The replication policy can not be changed as the volume group is associated with a storage partition.**

Explanation

The volume group is associated with a storage partition. This action is prohibited for volume groups in a storage partition.

User response

Identify the storage partition associated with the volume group and modify the storage partition replication policy.

CMMVC1025E **The command failed because the volume is associated with a storage partition.**

Explanation:
Volumes in Volume Groups with association to a storage partition cannot have the I/O group property changed.

User response:

None, the command cannot be used on objects associated with storage partition.

CMMVC1026E **The command failed because the host cannot have specific I/O group assigned because it is associated with a storage partitions.**

Explanation:
Hosts associated with a storage partition cannot have the I/O group property changed.

User response:
None, hosts associated with storage partitions cannot have I/O groups added or removed.

CMMVC1027E **The command failed because one or more hosts cannot be added to a host cluster because they are associated with a storage partition.**

Explanation:
Host clusters cannot be associated with a storage partition, therefore hosts that are associated with a storage partition cannot be added to host clusters.

User response:
Configure the host cluster with different hosts none of which are associated storage partition.

CMMVC1028E **The command failed because modifications are only allowed at the active management system for objects associated with a storage partition.**

Explanation:
Change to objects that are associated with storage partitions are only permitted at the active management system for the storage partition.

User response:
Attempt the command at the active management system for the associated storage partition.

CMMVC1029E **The command failed because the host cannot have a site assigned and be associated with a storage partition.**

Explanation:
Hosts that are associated with a storage partition cannot have a site defined.

User response:
Attempt the command without the site option, or remove the host from the storage partition.

CMMVC1030E **The command failed because the volume cannot be mapped to a host cluster because the volume is associated with a storage partition.**

Explanation:
Host clusters cannot be associated with storage partitions, therefore no host cluster mappings can be created for volumes that are associated with storage partitions.

User response:
Attempt the command with a volume that is not associated with a storage partition.

CMMVC1031E **The command failed because the host type is not supported for hosts associated with a storage partition.**

Explanation:
Certain host types are not allowed for hosts associated with a storage partition.

User response:
Attempt the command without the -type option, remove the host from the storage partition.

CMMVC1032E **The command failed because the name cannot be changed because it is associated with a storage partition configured for high availability.**

Explanation:
Names cannot be changed while associated with a storage partition that is associated with a high availability policy.

User response:
Remove the name option, remove the host from the storage partition, or remove the replication policy from the storage partition and run the command again.

CMMVC1033E **The command failed because the protocol is not supported for hosts associated with a storage partition.**

Explanation:
The host is using a protocol that is not allowed on a host with a storage partition.

User response:
Change the host protocol, or remove the storage partition from the command.

CMMVC1034E **The command failed because the host and the volume group the volume is part of must be associated with the same storage partition or neither associated with a storage partition.**

Explanation:

The host and the volume group the volume is part of must be associated with the same storage partition, or both the host and volume group the volume is part of must not be associated with a storage partition, or the host is not in a storage partition and the volume is not part of a volume group.

User response:
Add the volume to the volume group that is associated with the storage partition that is associated with the host. Remove the host from the storage partition.

CMMVC1035E **The command failed because a volume has received I/O within the defined volume protection period.**

Explanation:
The restore command was not run because the system determined that a volume received I/O within the defined volume protection time period.

User response:
Wait the time that is specified in the *vdisk_protection_time* field (see lssystem), and retry the command. Alternatively, if you do not want to have volume protection warnings, then you can disable the *vdisk_protection_enabled* field by using the chsystem command.

CMMVC1036E **The command failed because a storage partition can only be changed from the system which is the active management system.**

Explanation

It is only possible to change a storage partition from the system which is the active management system.

User response

Run the command again on the system which is the active management system.

CMMVC1037E **The command failed to remove the storage partition replication policy because the active management system and preferred management system are different.**

Explanation

By default, the replication policy can only be removed from the active management system, when it is also the preferred management system.

User response

On the system which is the active management system, change the preferred management system to be the same as the active management system and run the command again. If the preferred management system cannot be recovered, specify the **-deletepreferredmanagementcopy** flag to override the policing.

CMMVC1038E **The command failed because the replication policy specified is not compatible with the replication policy currently assigned to the storage partition.**

Explanation

A replication policy can only be assigned to a storage partition which currently has no assigned replication policy (a single-system storage partition) or the topology and locations specified in each policy are the same.

User response

Run the command again with a compatible replication policy.

CMMVC1039E **The command failed because the preferred management system specified is not a system in the replication policy assigned to the storage partition.**

Explanation

It is only possible to change the preferred management system for a storage partition to a system in the replication policy assigned to the storage partition.

User response

Only use this command to change the preferred management system for a storage partition to a system in the replication policy assigned to the storage partition.

CMMVC1040E **The command failed because one or more volumes in a volume group associated with the storage partition have a system-owned volume used for replication.**

Explanation

The user has attempted to add a replication policy to a storage partition with an associated volume

group which contains volumes associated with existing system-owned volumes used for replication. When a storage partition has a replication policy applied all volumes in the associated volume groups must have no system-owned resources for replication. This scenario can occur if a replication policy has recently been removed from the storage partition and the associated resources for replication are still in the process of being removed.

User response

Wait until the system has removed all the previous resources used for replication from all volumes in the volume groups associated with the storage partition, then retry the command.

CMMVC1041E The command failed because TCT can not be enabled while a restore is in progress.

Explanation:
An attempt was made to enable TCT on a volume while a snapshot restore was in progress on that volume.

User response:
Wait for the restore to complete and retry the command.

CMMVC1042E The command failed to remove the storage partition replication policy using `-deletepreferredmanagementcopy` because the active management system and preferred management system are the same.

Explanation

The command failed because it is not permitted to remove a replication policy from a storage partition using the **`-deletepreferredmanagementcopy`** flag when the active management system and preferred management system are the same. The replication policy may only be removed using the flag **`-deletepreferredmanagementcopy`** when the management systems are different.

User response

The **`-deletepreferredmanagementcopy`** parameter is required when the active management system and preferred management system are different. Using this flag will remove the storage partition replication policy on the active management system and remove the storage partition and all associated objects on the preferred management system. Specifically when the management systems are different and the preferred management system cannot be recovered. If the management systems are the same the flag **`-deletepreferredmanagementcopy`** is not required. Rerun the command on the active management system without this flag, which will remove the storage partition replication policy on the active management system, and then the storage partition and associated objects on the non-active management system will be removed.

CMMVC1043E The command failed as the local location status is not suspended.

Explanation

A stale copy of data is indicated by the location status being suspended. This is only possible on the non-active-management-system.

User response

Rerun the command on the system which is the non-active management system where the location status is suspended.

CMMVC1044E The command failed because the `fctargetportmode` cannot be changed from enabled.

Explanation:
The `fctargetportmode` is required to be enabled in this release. The enabled setting provides dedicated virtual Fibre Channel target ports using N_Port ID virtualization to be used for host access and prevents host access using the physical Fibre Channel WWPN that is used for clustering, replication, and backend storage.

User response:
None, `fctargetportmode` cannot be changed from enabled.

CMMVC1045E The command failed because the host is not configured in all I/O groups.

Explanation:
The host is not configured for all I/O groups.

User response:
Add the host to all I/O groups.

CMMVC1046E The command failed as one or more of the parameters specified are not permitted for a metadata volume.

Explanation:
You attempted to add a volume copy to a metadata volume by specifying parameters that are not permitted for this type of volume.

User response

Retry the command by specifying only permitted parameters.

- To add a volume copy to a metadata volume by using the **addvdiskcopy** command, the **-mdiskgrp** parameter must be specified.
- To add a volume copy to a metadata volume by using the **addvolumecopy** command, the **-pool** parameter must be specified.
- When you are adding a volume copy by using either command, the **-autodelete** parameter can optionally be specified.

No other parameters are permitted.

CMMVC1047E **The command failed because a volume cannot be created in a volume group associated with a storage partition from the non-active management system.**

Explanation

The user has attempted to create a new volume in a volume group that is associated with a storage partition with a replication policy on the non-active management system. New volumes can only be created in a volume group associated with a storage partition from the active management system.

User response

Identify the active management system of the storage partition associated with the volume group and create the volume on that system.

CMMVC1048E **The command failed because a cache disabled volume cannot be created in a volume group associated with a storage partition.**

Explanation

A cache disabled volume cannot be created in a volume group associated with a storage partition.

User response

Create a cache disabled volume without specifying a volume group, specify a volume group that is not associated with a storage partition or create the volume with caching enabled.

CMMVC1049E **The command failed because a volume with a copy in a standard pool and a copy in a data reduction pool cannot be created in a volume group associated with a storage partition.**

Explanation

A storage partition can only be associated with a mirrored volume when both copies are in the same type of storage pool. Either both copies must be in a data reduction pool or both copies must be in a standard pool.

User response

Create a volume with all copies in the same type of storage pool.

CMMVC1050E **The command failed because an ownership group can not be assigned to an object associated with a storage partition.**

Explanation

The command assigns an ownership group to an object associated with a storage partition, because the **-ownershipgroup** parameter was specified or the user is in an ownership group. An object associated with a storage partition can not be in an ownership group.

User response

If the user is not in an ownership group, run the command again without specifying the **-ownershipgroup** parameter.

CMMVC1051E **The command failed because the maximum number of TCP NQNs for the system would be exceeded.**

Explanation:
You ran `mkhost` command for a TCP/NVMe host and exceeded the maximum hosts limit per cluster.

User response:
Remove the existing hosts that are not needed.

CMMVC1053E **The command failed because the storage partition has a replication policy and remote objects are associated with the storage partition.**

Explanation

A storage partition can only be removed when it is does not have a replication policy, unless the location where objects should be deleted is specified using one of the two optional

flags **-deletenonpreferredmanagementobjects** or **-deletepreferredmanagementobjects**.

User response

Either remove the replication policy from the storage partition and resubmit the command, or resubmit the command with one of the two optional flags **-deletenonpreferredmanagementobjects** or **-deletepreferredmanagementobjects** to dictate where objects should be deleted.

CMMVC1054E **The command failed because the storage partition is a single-system storage partition and has no remote objects.**

Explanation

A storage partition can only be removed using one of the two optional flags **-deletenonpreferredmanagementobjects** or **-deletepreferredmanagementobjects** when it has a replication policy. The storage partition being deleted is a single-system storage partition with no remote objects.

User response

Resubmit the command without any additional flags.

CMMVC1055E **The command failed because the storage partition has a replication policy and the current state of the system does not allow deletion of the required objects in that location.**

Explanation

A storage partition can only be deleted with a replication policy when the user indicates where the objects that depend on it will be deleted. This is to protect against accidental deletion in the wrong location due to a change of management system. The parameter used to specify the location to delete objects did not match the expected state of the system.

User response

Either remove the replication policy from the storage partition and resubmit the command, or check the state of the storage partition and resubmit the command with **-deletepreferredmanagementobjects** or **-deletenonpreferredmanagementobjects** depending on desired outcome.

CMMVC1056E **The command failed because Global Mirror is either not supported on this system or not supported on the remote system.**

Explanation

Global Mirror is not available in the code version on either this system or the partnered system as it is replaced by asynchronous policy-based replication.

User response

If you need to continue the use of Remote Copy, then convert the relationship or consistency group to Metro Mirror to provide synchronous replication. If you need asynchronous replication, then configure the asynchronous policy-based replication instead.

CMMVC1057E **The command failed because the volume is in a volume group which is associated with a storage partition and the command specified is prohibited for these volumes.**

Explanation

The volume is in a volume group that is associated with a single-system storage partition or a storage partition with a HA replication policy, or it is replicated using a replication policy associated with a storage partition. The command specified is prohibited for these volumes.

User response

Perform the action on a volume that is not in a volume group associated with a storage partition.

CMMVC1058E **The command failed because the volume is in a volume group which is associated with a storage partition with a HA replication policy and the command specified is prohibited for these volumes.**

Explanation

The volume is in a volume group that is associated with a storage partition with a HA replication policy. The command specified is prohibited for these volumes.

User response

Perform the action on a volume that is not in a volume group associated with a storage partition with a replication policy.

CMMVC1059E **The command failed because an image mode volume or copy is not supported in a volume group associated with a storage partition.**

Explanation

An image mode volume or copy is not supported in a volume group associated with a storage partition. The command specified is prohibited as it would create an image mode volume or copy in a volume group associated with a storage partition.

User response

Perform the action on a volume or volume group not associated with a storage partition.

CMMVC1060E **The command failed because the specified volume group is associated with a storage partition.**

Explanation

The command specified is prohibited for a volume group associated with a storage partition.

User response

Retry the command without specifying a volume group, or specifying a volume group that is not associated with a storage partition.

CMMVC1061E **The command failed because the volume is in a volume group associated with a storage partition. Use the `rmvolume` command to remove the volume.**

Explanation

The **`rmvdisk`** and **`rmvdiskcopy`** commands cannot be used to remove a volume from a volume group associated with a storage partition.

User response

Use the **`rmvolume`** command to remove the volume.

CMMVC1062E **The command failed because the volume is in a volume group which is associated with a storage partition and cannot be zero size.**

Explanation

The volume is in a volume group that is associated with a storage partition. A volume associated with a storage partition cannot be zero size. The command specified is prohibited as it would result in a zero sized volume.

User response

Perform the action on a volume that is not in a volume group associated with a storage partition.

CMMVC1063E **The command failed because this system is not the active management system for the storage partition.**

Explanation

The user has attempted to modify or create an object associated with a storage partition that has a replication policy on the non-active management system. Objects associated with a storage partition that has a replication policy can only be modified or created from the partition active management system.

User response

Identify the active management system of the storage partition using **`lspartition`** command and resubmit the command on that system.

CMMVC1064E **The command failed because MTU cannot be changed on a port that has an IP address configured.**

Explanation

Changing the MTU value of an ethernet port, which is already configured with IP, is not allowed.

User response

If you want to change the MTU of the ethernet ports, do it before setting the IP address.

CMMVC1065E **There is insufficient capacity to complete the vdisk operation because the system is still formatting previously freed capacity. Wait until more capacity is available.**

Explanation

There is enough free capacity to complete the vdisk operation, but not all of it is formatted yet.

User response:

Wait until more capacity is formatted and becomes available to complete the vdisk operation.

CMMVC1066E **The command failed because the system has too many I/O Groups for a storage partition to be configured.**

Explanation

A storage partition can only be created on single I/O Group systems.

User response

Remove additional I/O groups before running the command again.

CMMVC1067E **The command failed because the system topology is stretched or Hyperswap and a storage partition can only be created on a system with standard topology.**

Explanation

A storage partition can only be created on a system with a standard topology. The system has either stretched or Hyperswap topology and a storage partition cannot be created on a system with this topology.

User response

Only create a storage partition on a system with standard topology.

CMMVC1068E **The command failed because the system topology is stretched or Hyperswap and a HA replication policy can only be created on a system with standard topology.**

Explanation:
A HA replication policy can only be created on a system with a standard topology. The system has either stretched or Hyperswap topology. A HA replication policy can not be created on a system with this topology.

User response:
Only create a HA replication policy on a system with standard topology. .

CMMVC1069E **The command failed because the volume group is associated with a single-system storage partition and changing the mode is only valid for a volume group with an asynchronous replication policy.**

Explanation

The volume group has no replication policy assigned as it is associated with a single-system storage partition. Changing the mode is only permitted for volume groups with an asynchronous replication policy.

User response

Run the command on a volume group with an asynchronous replication policy.

CMMVC1070E **The command failed because the volume group is associated with a replication policy that does not support restore in place.**

Explanation:
We do not support snapshot restores of volumes in volume groups with a replication policy.

User response:
Remove the replication policy associated with the volume group that is being restored or remove the volume from the volume group and specify the **-volumes** parameter to restore a subset of volumes that are not in a volume group.

CMMVC1071E **The command failed because one or more volumes is of the wrong type to be refreshed.**

Explanation:
Only thinclone volumes can be refreshed from a snapshot.

User response:
Ensure that all volumes in the group or list are of the correct type before retrying the command.

CMMVC1074e **The command failed because the source volume is not present in the specified snapshot for one or more of the volumes to be refreshed.**

Explanation

The snapshot specified does not include an image of one or more of the volumes to be refreshed.

User response

Choose a different snapshot to refresh from or change the list of volumes to refresh to remove those that are missing from the snapshot.

CMMVC1075E **The command failed because one or more of the volumes to refresh has active snapshots.**

Explanation

A volume cannot be refreshed if it has active snapshots itself.

User response

Remove all active snapshots for the volumes to be refreshed and retry the command.

CMMVC1076E **The command failed because one or more of the volumes to refresh is in a group with a snapshot policy defined.**

Explanation:
One or more of the volumes to be refreshed is in a volume group that has a snapshot policy specified. A volume cannot be refreshed if it is in a group with a snapshot policy.

User response:
Remove the snapshot policy from the volume group and retry the command.

CMMVC1078E **The command failed because one or more of the volumes is part of a Remote Copy relationship.**

Explanation

A volume that is participating in a remote copy relationship cannot be refreshed.

User response

Delete the remote copy relationships for all volumes to be refreshed and retry the command.

CMMVC1079E **The command failed because the volume group is in the process of being restored from a snapshot.**

Explanation:
The user has attempted to assign a replication policy to a volume group while that volume group is in the process of being restored from a snapshot. This is an unsupported action.

User response:
Wait until the restore has completed and then reissue the command. The **lsvolumegroup** output includes a field `restore_in_progress`, which indicates whether a restore operation is currently underway or not. Also, **lsvolumegrouppopulation** can be used to view progress information regarding the volume group restore or **lsvolumepopulation** can be used to view the progress of the restore at a volume level.

CMMVC1080E **The command failed because the specified ownership group link UID is already in use on this system.**

Explanation

The replication link UID value must uniquely identify an ownership group on the local system.

User response

Check the ownership group link UID value and retry the command. Alternatively, to move the ownership group link UID to a new ownership group, retry the command specifying the **-moveownershipgrouplink** parameter.

CMMVC1081E **The command failed as the SNMP server specified currently has a trap in progress.**

Explanation

If TLS is enabled then we need to wait for a response back before we can send another test trap.

User response

Wait until the trap is received on the server.

CMMVC1082E **The command failed as one or more SNMP servers currently has a trap in progress.**

Explanation

If TLS is enabled then we need to wait for a response back before we can send another test trap.

User response

Wait until the trap is received on the server.

CMMVC1084E **The command failed as the existing I/O Group has a storage partition configured and additional I/O groups are not supported.**

Explanation

A storage partition is only supported on single I/O group systems. As this feature is in use on this system, the new I/O group cannot be added to this system.

User response

Continue to use a single I/O group system or remove the storage partition. Alternatively, you can utilize the new nodes to set up a completely new system instead of adding another I/O group to the current one.

CMMVC1085E **Test trap failed to send. There is a problem with the SNMP server.**

Explanation

If TLS is enabled then we should report if we failed to establish a secure connection to the SNMP server.

User response

Check that the security settings on the cluster and server side are compatible, that a server certificate is installed in a trust store and that the server certificate is valid.

CMMVC1086E **Test trap failed to send. There is a problem with one or more SNMP servers.**

Explanation

If TLS is enabled then we should report if we failed to establish a secure connection to the SNMP server.

User response

Check that the security settings on the cluster and server side are compatible, that a server certificate is installed in a trust store and that the server certificate is valid.

CMMVC1087I **The test trap was sent successfully. Verify that the server(s) have received the trap.**

Explanation

If TLS is enabled then we should report if we successfully established a secure connection to the SNMP server.

User response

Relax because everything works!

CMMVC1089E **The command failed because an image of one or more of the specified source volumes is not present in the snapshot.**

Explanation:
An image of all source volumes specified in the list must be present in the snapshot.

User response:
Retry the command specifying the correct snapshot and source volumes.

CMMVC1091E **The command failed because the number of volumes on this system exceeds the limit supported by Remote Copy.**

Explanation

The system has a limit on the number of volumes that can be supported with the Remote Copy. If the number of volumes in the system exceed this limit, Remote Copy relationship cannot be added on this system.

User response

Retry creating relationship after volumes on the local system have reduced below the maximum volume limit supported for the Remote Copy.

CMMVC1092E **The command failed because the number of volumes on the remote system exceeds the limit supported by Remote Copy.**

Explanation

The system has a limit on the number of volumes that can be supported with the Remote Copy. If any of the volumes exceed this limit, Remote Copy relationship cannot be added on this system.

User response

Retry creating relationship after volumes on the remote system have reduced below the maximum volume limit supported for the Remote Copy.

CMMVC1094E **The command failed because a recovery key can only be configured when encryption is enabled.**

Explanation:
The user has attempted to configure an encryption recovery key. Encryption is not currently enabled on the system. An encryption recovery key cannot be configured without an existing key management method enabled on the system.

User response:
Enable encryption using USB flash drives or key servers, and then run the command again.

CMMVC1095E **The action failed because specified port does not support short-distance partnerships using RDMA.**

Explanation

Short-distance partnership using RDMA is only supported on RDMA capable adapters. Please check

output of lsportethernet to find out if specified port is RDMA capable or not.

User response:
.Select only RDMA capable Ethernet ports for IP configuration and association with highspeedreplication portsets.

CMMVC1098E **The command failed because the maximum limit of per node IP addresses associated with the highspeedreplication portset is reached.**

Explanation:
Highspeedreplication portset supports max upto 2 IP addresses per node.

User response:
Create a new highspeedreplication portset for IP association or select different node for IP address configuration.

CMMVC1100E **The command failed because portset of type highspeedreplication can be owned and created by global owner only.**

Explanation:
Portsets of type highspeedreplication can only be created by the global owner.

User response:
Contact the system administrator.

CMMVC1101E **The command failed because a recovery key rekey operation is in progress.**

Explanation

The user has attempted to prepare a new encryption key. An encryption recovery key is enabled on the system and is currently in the prepared or confirmed state. Only one rekey operation can be in progress at once.

User response

The rekey operation must either be confirmed and committed, or canceled. Use the **lsencryption** command to check the status of the rekey operation in progress for the recovery key.

If the recovery key is in the prepared state, confirm and commit the rekey operation in progress by entering the following commands:

```
chencryption -recoverykey newkey -key confirm
chencryption -recoverykey newkey -key commit
```

If the recovery key is in the confirmed state, commit the rekey operation in progress by entering the following command:

```
chencryption -recoverykey newkey -key commit
```

If the recovery key is in the confirmed state, cancel the rekey operation in progress by entering the following command:

```
chencryption -recoverykey newkey -key cancel
```

Once the rekey operation is complete, retry the original command.

CMMVC1102E **The command failed because the recovery key has not been confirmed.**

Explanation

The user has attempted to commit an encryption recovery key. An encryption recovery key is enabled on the system and is currently in the prepared state. An encryption recovery key cannot be committed until it has been confirmed by the user. Confirming the recovery key indicates that the key has been correctly copied and stored, and checks that the user's copy of the key is correct.

User response

Confirm the prepared recovery key, by entering the following command:

```
chencryption -recoverykey newkey -key confirm
```

Once the recovery key has been confirmed successfully, retry the original command. Alternatively, cancel the recovery key and prepare a new recovery key.

CMMVC1103E **The command failed because volumes specified for refresh have the same source volume.**

Explanation

When refreshing volumes none of the volumes specified can have the same source volume.

User response

Adjust the specified volumes so that they do not contain volumes with the same source volume.

CMMVC1104E **The command failed because volumes specified belong to a volume group with a replication policy.**

Explanation

Refreshing volumes that are associated with a replication policy is not permitted.

User response

This command can only be used to refresh volumes that are not being replicated. Either remove the replication policy from the volume group containing the volumes and retry the command, or move the volumes out of the volume group and use the **-volumes** parameter to specify the volumes to refresh.

CMMVC1105E **The command failed because it is not supported using the RESTful API.**

Explanation

The user has run a command via the RESTful API that isn't supported.

User response:
Run the command from the Command Line Interface instead.

CMMVC1106E **The command failed because mixing IP address types within a highspeedreplication type portset is not allowed.**

Explanation

All the IP addresses within a highspeedreplication type portset must be of identical type. For example, either all are IPv4 type or all are IPv6 type.

User response:
Select either IPv4 or IPv6 version within a highspeedreplication type portset.

CMMVC1108E **The command failed because the specified portset is used by another partnership.**

Explanation:
Short-distance partnership using RDMA requires a dedicated portset per partnership link.

User response:
Configure dedicated portset per replication link for short-distance partnership using RDMA.

CMMVC1109E **The command failed because mixing of RoCE and iWARP ports within a highspeedreplication type portset is not allowed.**

Explanation:
All the ports within an highspeedreplication type portset must be of identical type. For example, either all are RoCE type or all are iWARP type.

User response:
Select either RoCE or iWARP port types within a highspeedreplication type portset.

CMMVC1110E **The command failed because port sharing within a highspeedreplication type portset is not allowed.**

Explanation

Sharing of highspeedreplication ports with any other functionality is not allowed.

User response:
Select a different highspeedreplication port.

CMMVC1111E **The highspeedreplication portset cannot be created when Ethernet clustering is active.**

Explanation:
The action failed because Ethernet clustering is configured in system.

User response:
Check for the field **eth_clustering** in **lsportethernet** view. Remove Node IP for all port ids, where **eth_clustering** is set by using **satask chnodeip** command.

CMMVC1112E **The command failed because type of specified link portset for partnership is unsupported.**

Explanation

Only a replication or highspeedreplication type portset can be mapped to a partnership.

User response

Select either replication or highspeedreplication type portset.

CMMVC1113E **The command failed because an IP partnership requires at least one link to be specified.**

Explanation

This command is failed because at least one link parameter must be specified in the command of type replication or highspeedreplication.

User response

Provide at least one portset of type replication or highspeedreplication to link1 or link2.

CMMVC1114E The command failed because different link portset types are specified for partnership.

Explanation

Both link1 and link2 must be configured with identical portset type which is either replication or highspeedreplication.

User response

Map link1 and link2 with same portset type either replication or highspeedreplication.

CMMVC1115E The command failed because partnership type of local and remote systems mismatched.

Explanation

This error occurs when the partnership type of local and remote systems does not match. Example: Using highspeedreplication portsets for partnership implies that it is a short-distance partnership using RDMA. Using replication portsets implies that it is a native IP replication partnership.

User response

Use highspeedreplication portsets at both systems if short-distance partnerships using RDMA is desired. Use highspeedreplication portsets at both systems if native IP replication partnership is desired.

CMMVC1116E Secured and compressed flags are not supported for short-distance partnerships using RDMA.

Explanation

Secured and compressed parameters are not supported for short-distance partnerships using RDMA.

User response

Do not use secured and compressed flags for short-distance partnerships using RDMA.

CMMVC1117E Cannot create a short-distance partnerships using RDMA with remote system running software version prior to 8.6.2.0.

Explanation

As older systems do not have short-distance partnerships using RDMA capability, short-distance partnerships using RDMA cannot be built on a system with a software version previous to 8.6.2.0.

User response

To construct a short-distance partnerships using RDMA, upgrade the peer system to a supported version 8.6.2.0.

CMMVC1118E The command failed due to a mismatch in the specified partnership link portset type from earlier settings.

Explanation

It is not allowed to change the link portset type on existing partnerships. For example, a partnership created with replication portsets cannot be modified to have highspeedreplication portsets. Similarly a partnership created with highspeedreplication portsets cannot be modified to have replication portsets.

User response

Specify the identical type of link portsets while modifying partnerships.

CMMVC1119E The command failed because specified feature is not supported on this platform.

Explanation:
The specified action is not allowed on the current platform.

User response:
Contact the system administrator.

CMMVC1120E The command failed because the maximum number of short-distance partnerships using RDMA are exceeded.

Explanation

Maximum number of short-distance partnerships using RDMA are exceeded.

User response

Maximum number of short-distance partnerships using RDMA are exceeded. Please remove the existing short-distance partnerships using RDMA to create a new one.

CMMVC1121E **The command failed because the system cannot increase effective capacity beyond the supported limit.**

Explanation:
The command failed because the addition to the effective capacity takes it beyond the supported limit for the system.

User response:
Create or expand the distributed RAID array with fewer drives (or a smaller stripe width if creating a new distributed RAID array).

CMMVC1122E **The command failed because the volume group is not a production or recovery copy.**

Explanation:
The specified action is only permitted on a production or recovery copy of a volume group.

User response:
None.

CMMVC1123E **The command failed because the recovery key has been entered incorrectly too many times. Wait %1 minute(s) and try again.**

Explanation:
The user has attempted confirm the recovery key three times incorrectly. This message shows when they are temporarily locked out from using the command.

User response:
Wait the specified time or cancel the recovery key and prepare a new one.

CMMVC1124E **The command failed because the recovery key has been entered incorrectly.**

Explanation:
The encryption recovery key entered by the user did not match the one stored on the system.

User response:
Run the command again and enter the recovery key correctly.

CMMVC1125E **The command failed because encryption must be disabled using the technician port while encrypted storage exists.**

Explanation:
As an additional security measure, disabling encryption while there are encrypted objects requires physical access via the technician port.

User response:
Identify which node is the configuration node, connect to the technician port and rerun the command while logged in via the technician port.

CMMVC1126E **The command failed because the maximum number of volumes supported on a system using Remote Copy has been reached.**

Explanation

The system has a limit for maximum number of volumes if a Remote Copy relationship is present. A new volume cannot be added if volume limit for Remote Copy is reached on a system with the existing Remote Copy relationship.

User response

A new volume can be added if number of volumes is reduced, or all Remote Copy relationships are removed.

CMMVC1127I **The new recovery key for the system is: %1.**

Explanation:
A new encryption recovery key has been prepared successfully. The user will need to provide this recovery key to the system in order to bring encrypted storage online when other encryption key methods are unavailable. The current recovery key must also be provided in order to perform a re-key operation of the recovery key.

User response:
Make a note of the encryption recovery key. The key is sensitive, so make sure it has been stored securely. Consider making multiple copies of the key. Run the **`lsencryption`** command to view the **`recovery_key_rekey_uid`**, a unique identifier for the recovery key.

CMMVC1128E **The command failed because an encrypted cloud account cannot be created when the recovery key is the only encryption method that has been enabled.**

Explanation:
An encrypted cloud account cannot be created when the encryption recovery key is the only encryption method that has been enabled.

User response:
Configure encryption using USB flash drives or key servers, then retry the command.

CMMVC1129E **The command failed because safeguarded option cannot be used with cloud backup policy.**

Explanation

The user attempted to use safeguarded option to apply the cloud backup policy to a volume group.

User response

User cannot use safeguarded option to apply the cloud backup policy to a volume group.

CMMVC1130E The snapshot policy cannot be applied as the cloud account is not configured.

Explanation

The user attempted to apply the cloudbackup policy to the volume group but the cloud account does not exist.

User response

Create a cloud account before applying cloudbackup policy to a volume group.

CMMVC1131E The cloud snapshots cannot be enabled because the volume(s) in this volumegroup is part of a FlashCopy mapping.

Explanation

If volumegroup has volumes having legacy FlashCopy mappings with volumes, then the cloudbackup policy cannot be applied to the volume group for enabling cloud backup on those volumes.

User response

User should remove legacy FlashCopy mapping of volumes in the volume group for applying the cloudbackup policy.

CMMVC1132E The command failed because the volume group has a replication policy.

Explanation:
It is not possible to remove a volume group containing volumes if the group has an associated replication policy.

User response:
Remove the replication policy or remove the volumes from the group before retrying the command.

CMMVC1133E The command failed because the volume group has a snapshot policy.

Explanation:
It is not possible to remove a volume group containing volumes if the group has an associated snapshot policy. Also, it is not possible to add volume to a volume group if the volume group has an associated cloud snapshot policy.

User response:
Remove the snapshot policy or remove the volumes from the group before retrying the command.

CMMVC1134E The command failed because the volume group is associated with a storage partition.

Explanation:
It is not possible to remove a volume group containing volumes if the group has an associated storage partition.

User response:
Delete the storage partition before retrying the command.

CMMVC1135E The command failed because the local system is already the active management system.

Explanation:
This error occurs when the **`chpartition -recoverams`** command is attempted on a storage partition that already has the local system set as active management system.

User response:
The system automatically determines the active management system. No further action is required.

CMMVC1136E The command failed because a recovery key has not been configured.

Explanation:
The user has attempted to complete an action that can only be completed if a recovery key has been configured.

User response:
Configure a recovery key and run the command again.

CMMVC1137E The command failed because the recovery key must have been validated within the last %2 minutes.

Explanation:
The user has attempted to prepare a new encryption key. An encryption recovery key is already committed, so the user must prove that they have the existing key before they are allowed to re-key.

User response

Run the following command and enter the existing recovery key when prompted:

```
chencryption -recoverykey validate
```

Then retry the original command within 30 minutes.

CMMVC1138E The command failed because the local system has already been configured for high availability with another system.

Explanation:
The user is trying to configure HA with a remote system, however the local system is already configured for HA with another (third) system.

User response:
Review maximum limits for HA topologies. Either use the existing HA partnership or disable the current HA policies and try again.

CMMVC1139E The command failed because the remote system has already been configured for high availability with another system.

Explanation:
The user is trying to configure HA with a remote system, however the remote system is already configured for HA with another (third) system.

User response:
Review max limits for HA topologies. Either use the existing HA partnership on the remote system or disable the current HA policies on the remote system and try again.

CMMVC1140E The command failed because the existing I/O group is configured for vVol replication.

Explanation:
vVol replication can only be configured on systems with a single I/O group.

User response:
To add an additional I/O group to the system, you must unconfigure vVol replication.

CMMVC1141E The command failed because the volume group is used for vVol replication and is not empty.

Explanation:
When a volume group used for vVol replication contains volumes, a replication policy can only be removed from the volume group by a user with the VasaProvider role.

User response:
Use the VMware vSphere application to remove the volumes from the volume group, then retry the command. Alternatively, if the VMware vSphere application is unavailable, log in as a user with the VasaProvider role, then retry the command.

CMMVC1142E The command failed because the specified volume group does not have a replication policy that supports this function.

Explanation:
The specified action is only permitted if the volume group has a replication policy with a topology that supports this function.

User response:
None.

CMMVC1143E The command failed because the topology of the replication policy is not supported with VMware Virtual Volumes (vVols).

Explanation:
A replication policy with a 2-site-async-dr topology is required to use policy-based replication with vVols.

User response:
Specify a replication policy with a 2-site-async-dr topology to use vVol replication.

CMMVC1144E The command failed because a snapshot policy is not supported with VMware Virtual Volumes (vVols).

Explanation:
Snapshot policies are not supported with vVols.

User response:
Create the volume group without specifying a snapshot policy.

CMMVC1145E The command failed because a Safeguarded policy is not supported with VMware Virtual Volumes (vVols).

Explanation:
Safeguarded policies are not supported with vVols.

User response:
Create the volume group without specifying a Safeguarded policy.

CMMVC1146E The command failed because the ownership group has not been linked to the remote systems in the replication policy.

Explanation:
The user has either attempted to assign a replication policy to a volume group with an ownership group, or assign an ownership group to a volume group with a replication policy, when the ownership group is not linked to a corresponding ownership group on

the remote systems in the replication policy. For a replication policy to be assigned to a volume group with an ownership group, the local ownership group must be linked to an ownership group in each of the remote systems in the replication policy.

User response:
Configure links between the local and remote ownership groups and retry the command.

CMMVC1147E **The command failed because the system has too many I/O groups for vVol replication to be configured.**

Explanation:
vVol replication can only be configured on systems with a single I/O group.

User response:
Remove additional I/O groups before creating volume groups for vVol replication.

CMMVC1148E **The command failed because the volume group is used for vVol replication and has restricted use.**

Explanation:
Some operations for managing volume groups used for vVol replication are restricted and the command specified is prohibited.

User response:
Review the documentation to identify the supported commands and parameters.

CMMVC1149E **The command failed because the recovery copy is a diverged copy as a synchronisation has not completed since the recovery copy of the volume group was made accessible.**

Explanation:
This copy is considered as a diverged copy as the data and configuration of the volume group may have been modified independently of changes made to the production copy. Until the next synchronisation completes, enabling access to the recovery copy is prevented as the volume group may not contain a recovery point that existed on the production system. Once a synchronisation completes, the data and configuration of the volume group will be from the production copy and a recovery point will be available.

User response:
Wait for the synchronisation to complete to establish a recovery point, then retry the command. Alternatively, access can be enabled to the diverged copy by specifying the **-accessdivergedcopy** parameter when changing the mode to independent.

CMMVC1151E **The command failed because the system cannot retrieve the current key from the primary key server or a USB flash drive.**

Explanation:
The user has attempted to either prepare a new recovery key, or disable the existing recovery key for the system. In order to do this, the current encryption key must exist on either the primary key server or a USB flash drive. This error means that the current encryption key could not be fetched from either the primary key server or a USB flash drive.

User response:
Make sure the primary key server containing the current key is online, or that a USB flash drive containing the current key is installed. Then try the command again.

CMMVC1153I **The recovery key has been entered correctly.**

Explanation:
The user has successfully validated their recovery key, which matches the stored HMAC has been used to unlock the system (if required).

User response:
Validate the recovery key by running `chencryption -recoverykey validate` and entering the current recovery key.

CMMVC1160E **The command failed because disablecim/disablesuperusercim is no longer supported.**

Explanation:
disablecim and **disablesuperusercim** are no longer supported parameters and so cannot be used to change whether CIMOM is disabled for user groups or superuser.

User response:
Do not try to enable CIMOM or contact support.

CMMVC1161E **The command failed because last system IP cannot be removed.**

Explanation:
If the last system IP is also removed, then the CLI and GUI becomes inaccessible.

User response:
Configure a new IP address before attempting to remove this one.

CMMVC1162E **The command failed because system_management portset cannot be removed by any ownership group user.**

Explanation:

System defined management port set cannot be removed.

User response:
Invalid use of **chportset** command.

CMMVC1163E **The command failed because Management IP can be configured only on config node.**

Explanation:
Cluster IP or Management IP cannot be configured on a non-config node.

User response:
Use the config node id instead of non-config node id.

CMMVC1166E **This command failed because maximum number of Management IP has been set.**

Explanation

Only two system IP addresses can be configured in any system.

User response

Delete an existing system IP address before adding another IP address.

CMMVC1168E **The command has failed because the maximum number of storage partition migration resources are exhausted.**

Explanation

Storage partition migration resources (for active as well as for queued) are exhausted hence new migration request can not be accepted.

User response

Use the **lspartition** command to list the active storage partition migrations. Retry the command when storage partition migration request are completed.

CMMVC1169E **The command failed as there is a migration in progress for the storage partition.**

Explanation

The user is attempting to apply HA replication policy to a storage partition for which migration is in progress. Storage partition can not have HA replication policy and storage partition migration active at the same time.

User response

N/A.

CMMVC1170E **The command failed as the location cannot be changed as there is a migration in progress for the storage partition.**

Explanation

User is trying to override storage partition migration hence, existing storage partition migration needs to be aborted to start new migration for the provided storage partition.

User response

Use the override flag to override existing storage partition migration location.

CMMVC1171E **The command failed as there is a replication policy assigned to the storage partition.**

Explanation

The user is attempting storage partition migration for which HA replication policy is already assigned. Storage partition can not have HA replication policy and storage partition migration active at the same time.

User response

Use needs to abort HA replication if same storage partition needs to be migrated.

CMMVC1172E **The command failed as the new location can not be accepted at this stage of a storage partition migration.**

Explanation

The user has tried to replace the existing migration location with another location for a storage partition where the ongoing storage partition migration cannot be rolled back.

User response

Retry the operation when existing storage partition migration gets completed.

CMMVC1173E **This command is deprecated, use the chip or rmip command to manage IP addresses on ethernet ports outside 1 and 2.**

Explanation:
chsystemip can be used to change or remove IP only on ethernet port 1 and 2.

User response:
Use **rmip** to remove IP and chip to change IP.

CMMVC1174E **The command failed because NVMeTCP/RDMA host creation is not supported over ports with MTU greater than 1500.**

Explanation:
FlashSystem NVMeTCP/RDMA protocols do not work with Ethernet ports that have MTU greater than 1500. Creating a new NVMeTCP/RDMA host or changing an existing NVMeTCP/RDMA host fails if the portset has Ethernet ports with MTU greater than 1500.

User response:
Use FlashSystem NVMeTCP/RDMA protocols only with Ethernet ports that have MTU of 1500.

CMMVC1175E **It is not possible to add an IP address to a port that has MTU greater than 1500 mapped to an NVMeTCP/RDMA host.**

Explanation:
FlashSystem NVMeTCP/RDMA protocols do not work with Ethernet ports that have MTU greater than 1500. Associating an IP address with a port and adding it to a portset that is linked with a NVMeTCP/RDMA host results in failure.

User response:
Use FlashSystem NVMeTCP/RDMA protocols only with Ethernet ports that have MTU of 1500.

CMMVC1176E **The command failed because last system IP cannot be removed.**

Explanation:
If the last system IP is also removed, then the CLI and GUI becomes inaccessible.

User response:
Configure a new IP address before attempting to remove this one.

CMMVC1177E **Cannot set both system IPs on same ethernet port.**

Explanation:
Both system IP's need to be set on different port for high availability.

User response:
Use port number where system IP is not configured.

CMMVC1181E **The command failed because a replication policy cannot be associated with a draft storage partition.**

Explanation

Draft partitions are not allowed to be associated with a replication policy.

User response

Publish the draft storage partition.

CMMVC1182E **The command failed because the draft storage partition cannot currently be published, because a host has a mapped volume that is not in a volume group.**

Explanation

One or more hosts associated with the draft storage partition have a mapping to a volume that is not part of a volume group.

User response

Examine host mappings for the hosts in the draft storage partition and add mapped volumes to volume groups.

CMMVC1183E **The command failed because the storage partition specified is a draft storage partition.**

Explanation

The partition provided is a draft storage partition but the parameter used is only valid for non-draft storage partitions.

User response

Specify a storage partition that is not a draft partition.

CMMVC1184E **The command failed because the storage partition specified is not a draft storage partition.**

Explanation

The partition provided is not a draft storage partition but the parameter used is only valid for draft storage partitions.

User response

Specify a storage partition that is a draft.

CMMVC1185E **The command failed because the host is mapped to one or more volumes.**

Explanation

The host is mapped to one or more volumes.

User response

Add the volumes to a volume group if they are not part of a volume group, then add the volume group to the draft partition.

CMMVC1187E **Cloud backup cannot be disabled on volume in the volumegroup having cloudbackup policy.**

Explanation

Cloud backup cannot be disabled on volume in the volumegroup having cloudbackup policy.

User response

User should remove volume from the volume group having a cloudbackup policy to disable cloud backup on the volume.

CMMVC1188E **The command failed because schedule cannot be added to a snapshot policy that is already associated with a volume group.**

Explanation:
The user attempted to add a schedule to a snapshot policy that is already associated with a volume group.

User response:
User should disassociate the snapshot policy from the volume group, add a new schedule to it and associate the policy back to the volume group.

CMMVC1189E **Volume could not be added to volumegroup having cloudsnapshot policy because the cloud account configured is in import mode.**

Explanation

Volume could not be added to volumegroup having cloudsnapshot policy because the cloud account configured is in import mode.

User response

User should remove import mode from cloud account and add volume to volumegroup having a cloud snapshot policy.

CMMVC1190E **Cloud snapshot policy could not be applied to the volumegroup because the cloud account configured is in import mode.**

Explanation

The cloud snapshot policy could not be applied to the volumegroup because the cloud account configured is in import mode.

User response

User should remove the import mode from cloud account and apply the cloud snapshot policy to the volumegroup.

CMMVC1191E **The command failed because the specified schedule type already exist.**

Explanation:
The user attempted to add snapshot schedule to an existing snapshot policy that already has that schedule type.

User response:
User should add a different snapshot schedule to the existing snapshot policy.

CMMVC1192E **The command failed because the limit of cloud backup volumes exceeded.**

Explanation

The user attempted to add volume to volume group having a cloudbackup policy that would cause cloud backup enabled volumes to exceed the maximum limit.

User response

User should disable cloud backup on an existing cloud volume and delete its associated cloud snapshots to add volume to a volume group that has a cloudbackup policy.

CMMVC1193E **The command failed because the limit of cloud backup volumes exceeded.**

Explanation

The user attempted to apply cloud backup policy to volume group that would cause cloud backup enabled volumes to exceed the maximum limit.

User response

User must disable cloud backup on an existing cloud volume and delete its associated cloud snapshots to apply cloud backup policy to a volume group.

CMMVC1194E **The command failed because the policy with cloud backup schedule**

cannot be applied on thinclone volume group.

Explanation

Volume group cannot be created with thin cloned volumes using cloud backup policy.

User response

User is not allowed to create volume group with thin cloned volumes using cloud backup policy.

CMMVC1195E **The command failed because the policy with cloud backup schedule cannot be applied while creating clone volume group.**

Explanation

A volume group cannot be created with cloned volumes using cloud backup policy.

User response

User is not allowed to create a volume group with cloned volumes using cloud backup policy.

CMMVC1197E **Volume could not be added to volumegroup having policy with cloud backup schedule because the cloud account is not configured.**

Explanation

Volume could not be added to volumegroup having a policy with cloud backup schedule because the cloud account is not configured.

User response

User should configure a cloud account and add volume to volumegroup having a policy with cloud backup schedule.

CMMVC1199e **The command failed as this feature has been deprecated and compatibility mode is not enabled for this feature.**

Explanation

This feature is deprecated and it will be removed in a future release. More recent code versions include new features and functions that replaced this feature. Support for the feature can be re-enabled by turning on compatibility mode for the feature if required. However, future upgrades may be blocked while the legacy features continue to be used on this system.

User response

Review the objective of what is trying to be achieved by using this command and determine whether newer functions can be used instead. If the newer functionality cannot be used, the feature can be enabled in a compatibility mode.

CMMVC1200e **The command failed as adding another I/O group requires compatibility mode to be enabled.**

Explanation

Adding a second I/O group to the system prevents certain features from being used and may prevent future upgrades. Features in recent code versions enable much of the functionality that was previously available only in systems with multiple I/O groups for systems that are partnered. For example, high-availability and migration.

User response

Review the objective of what is trying to be achieved by adding another I/O group and determine whether newer functions can be used instead, by creating a new system and partnering the two systems. If the newer functionality cannot be used, support for adding more I/O groups can be enabled by enabling the compatibility mode.

CMMVC1201e **The command failed because compatibility mode for the feature cannot be changed as it is currently in use.**

Explanation

To turn off compatibility mode for a feature, it cannot be in use on the system.

To disable Remote Copy, ensure that no Remote Copy relationships or Remote Copy consistency groups are on the system, and the system topology is not *HyperSwap*.

To disable support for multiple I/O groups, ensure that there are nodes only in a single configured I/O group on the system.

User response

Remove all configuration that is related to the feature before trying again.

CMMVC1202E **The command could not be initiated because at least one of the partitions is configured for high availability and the partnership is not connected.**

Explanation

One of the two partitions in the merge operation has an HA replication policy and the partnership is currently not connected.

User response

Fix the connectivity between the two systems and retry.

CMMVC1203E **The command could not be initiated because the replication attributes of the two partitions do not match.**

Explanation

The replication attributes of the partitions need to match exactly. The only fields that are allowed to be different are:

- Name
- Host count
- Host offline count
- Volume group count

User response

Fix any issues blocking HA from operating and wait until both partitions are established.

CMMVC1204E **The command failed because only Management IPs can be changed.**

Explanation:
Only Management IP's can be changed using chip.

User response:
Use **rmip** followed by **mkip** to change an IP other than Management IP.

CMMVC1205E **The command failed because parameter node is required for data IP's.**

Explanation:
Data IP's always require parameter node where it needs to be configured.

User response:
Specify the node where data IP needs to be configured.

CMMVC1206E **The command failed because Management IPs cannot be shared.**

Explanation:
Management portsets cannot share IP addresses.

User response:
Use unique IP addresses for the IP Management portsets.

CMMVC1207E **The command failed because a parent volume of the specified snapshot has pinned data.**

Explanation:
There is pinned data of a parent of the specified snapshot. This data is required for the operation to proceed.

User response:
A parent volume of the specified snapshot has pinned data that must be resolved. This can be caused by the volume being offline or issues with quorum.

CMMVC1211E **The command failed because one or more required hosts is not configured for all I/O groups, the first host id is *host_id*.**

Explanation

The command failed to add a volume group to the partition because a host that is required in the partition does not meet partition membership rules.

User response

- Modify the host to be in all I/O groups, or
- Delete or modify host volume mappings such that the host does not have mappings to volumes in the required volume group(s), or
- Remove the mapped volumes for the host from the volume group(s) that are required as a result of volume mappings.

CMMVC1212E **The command failed because one or more required hosts has a site defined, the first host id is *host_id*.**

Explanation

The command failed to add a volume group to the partition because a host that is required in the partition does not meet partition membership rules.

User response

- Modify the host to not have a site defined, or

- Delete or modify host volume mappings such that the host does not have mappings to volumes in the required volume group(s), or
- Remove the mapped volumes for the host from the volume group(s) that are required as a result of volume mappings.

CMMVC1213E **The command failed because one or more required hosts has a protocol defined that is not supported for use in storage partitions, the first host id is *host_id*.**

Explanation

The command failed to add a volume group to the partition because a host that is required in the partition does not meet partition membership rules.

User response

- Modify the host to not have a supported protocol defined, or
- Delete or modify host volume mappings such that the host does not have mappings to volumes in the required volume group(s), or
- Remove the mapped volumes for the host from the volume group(s) that are required as a result of volume mappings.

CMMVC1214E **The command failed because one or more required hosts is part of a host cluster, which is not supported for hosts in storage partitions, the first host id is *host_id*.**

Explanation

The command failed to add a volume group to the partition because a host that is required in the partition does not meet partition membership rules.

User response

- Modify the host to not be part of a host cluster, or
- Delete or modify host volume mappings such that the host does not have mappings to volumes in the required volume group(s), or
- Remove the mapped volumes for the host from the volume group(s) that are required as a result of volume mappings.

CMMVC1215E **The command failed because one or more required hosts has a host type that is not supported for use in storage partitions, the first host id is *host_id*.**

Explanation

The command failed to add a volume group to the partition because a host that is required in the partition does not meet partition membership rules.

User response

- Modify the host type to generic, or
- Delete or modify host volume mappings such that the host does not have mappings to volumes in the required volume group(s), or
- Remove the mapped volumes for the host from the volume group(s) that are required as a result of volume mappings.

CMMVC1216E **The command failed because one or more required hosts has an ownership group defined, this is not supported for hosts in storage partitions, the first host id is *host_id*.**

Explanation

The command failed to add a volume group to the partition because a host that is required in the partition does not meet partition membership rules.

User response

- Remove the host from the ownership group, or
- Delete or modify host volume mappings such that the host does not have mappings to volumes in the required volume group(s), or
- Remove the mapped volumes for the host from the volume group(s) that are required as a result of volume mappings.

CMMVC1217E **The command failed because one or more required volume groups does not meet the prerequisites for use in a storage partition the first volume group id is *volume_group_id*.**

Explanation

The command failed to add a volume group to the storage partition because a volume group that is required in the storage partition does not meet storage partition membership rules.

User response

- Examine the volumes and volume group identified in the error message and modify so that the following rules are met:
  - Volume group must not have a volume that is in an **RC relationship**.
  - Volume group must be of type **generic, thin,** or **thin-clone**.
  - Volume group must not contain a volume that is **size 0 (zero)**.
  - Volume group must not contain a volume that is **not cache Read/Write enabled** (that is **cache=none**, or **cache=readonly**).
  - Volume group must not contain a volume that has **transparent cloud tiering (TCT) enabled**.
  - Volume group must not contain a volume that is an **image mode** volume.
  - Volume group must not contain a volume that is an **access point vdisk**.
  - Volume group must not contain a volume or is a volume group **used by VMware vVols**.
  - Volume group must not have any volumes that have **SCSI Persistent Reservations** or **SCSI Persistent Reservations Registrations**.
  - Volume group must not be restoring or have any volumes that are currently **restoring**.
  - Volume group must not have a replication policy, for example, configured for **policy-based replication**.
- Delete or modify host volume mappings such that the host does not have mappings to volumes in the required volume group(s), or
- Remove volumes from volume groups such that the relationships between the hosts and volumes created by the volume mappings in related volume groups is reduced or are removed.

CMMVC1218E **The command failed because one or more required volume groups has a volume that is the target of a Flash Copy map, this is not supported for volumes in storage partitions, the first volume id is *volume_id*.**

Explanation

The command failed to add a volume group to the storage partition because a volume group that is required in the storage partition has a volume that is the target of a Flash Copy map.

User response

- Delete the flash copy relationship, or delete the volume, or
- Delete or modify host volume mappings such that the host does not have mappings to volumes in the required volume group(s), or
- Remove the mapped volumes for the host from the volume group(s) that are required as a result of volume mappings.

CMMVC1219E **The command cannot be used to manage IP addresses that are configured with VLAN, use rmip or chip command instead.**

Explanation:
chsystemip cannot be used to change or remove the IP with vlan.

User response:
Use **rmip** or **chip** to remove or change IP.

CMMVC1220E **The command could not be initiated because a partition cannot be merged into itself.**

Explanation

A partition cannot be merged into itself as the resulting partition would be unchanged.

User response

None required, partition already in resulting state.

CMMVC1221E **The command failed because IPv4 cannot be changed to IPv6 and other way around.**

Explanation:
Existing IPv4 can be changed to another IPv4 address and IPv6 to another IPv6 address with or without vlan.

User response:
Use IPv4 address if previous IP was IPv4 or IPv6 address.

CMMVC1223E **The command failed because Management IP should be configured on port id which is available on all nodes in the cluster node.**

Explanation:
Cluster IP or Management IP should be configured on port which is available on all nodes in the cluster.

User response:
Use port which is available in all nodes in the cluster.

CMMVC1224E **The command could not be initiated because the command is**

not supported on both systems in the partnership.

Explanation

The partnership is running on mixed software versions and one of them does not support this command.

User response

Update the software on the downlevel system to a software level that supports the command.

CMMVC1225E **The command could not be initiated because one or both of the partitions has a High Availability replication policy and HA is not established.**

Explanation

The partitions have a replication error or are still syncing their configurations, and cannot be merged in this state.

User response

Wait until both the partitions reach established HA status, or remove replication policies from both partitions.

CMMVC1226E **The command could not be initiated because a previous merge of two highly available partitions has not completed on both systems.**

Explanation

A previously executed merge was not completed due a loss of connectivity between the systems. The systems cannot merge another highly available partition until this has been resolved. Please fix all connectivity issues and replication errors in the event log before retrying.

User response

Resolve connectivity problems between systems. Resolve replication error codes in event log then retry.

CMMVC1227E **The command failed because one or more required volume groups has a volume that is configured with more than 1 access I/O group, or the volume is not in the correct I/O group. This is not supported for volumes in storage partitions, the first volume id is *volume_id*.**

Explanation

The command failed to add a volume group to the storage partition because a volume group that is required in the storage partition has a volume that has incorrect I/O group configuration.

User response

- Delete the access I/O groups for the volume, or move the volume to the correct I/O group, or
- Delete or modify host volume mappings such that the host does not have mappings to volumes in the required volume group(s), or
- Remove the mapped volumes for the host from the volume group(s) that are required as a result of volume mappings.

CMMVC1231E **The command failed because the volume group specified contains no thinclone volumes or the specified list of volumes contains a volume that is not thinclone.**

Explanation

The user has specified a list of volumes, or a volume group, that does not contain any thinclone volumes.

User response

Specify a list of thinclone volumes or a volume group that contains thinclone volumes and rerun the command.

CMMVC1232E **The command failed because the volume is zero size.**

Explanation

The volume specified is of zero size so cannot be added to a volume group that is associated with a draft storage partition.

User response

Delete the volume, or change its size.

CMMVC1233E **The command failed because the maximum number of draft storage partitions already exist.**

Explanation

A limited number of draft storage partitions are allowed.

User response

Publish or remove an existing draft storage partition.

CMMVC1234E The command failed because the volume has caching disabled.

Explanation

The volume specified is not cache enabled read/write so cannot be added to a volume group that is associated with a draft storage partition.

User response

Change the caching mode of the volume to read/write.

CMMVC1235E The command failed because the volume has image mode copies.

Explanation

The volume that is specified is an image mode volume. Therefore, it cannot be added to a volume group that is associated with a draft storage partition.

User response

The volume cannot be added to a volume group that is associated with a storage partition.

CMMVC1236E The command failed because the volume is part of a Remote Copy relationship.

Explanation

The volume specified is part of a remote copy relationship, so cannot be added to a volume group that is associated with a draft storage partition.

User response

Remove the volume from the remote copy relationship.

CMMVC1237E The command failed because the volume has Transparent Cloud Tiering enabled.

Explanation

The volume specified has Transparent Cloud Tiering enabled, so cannot be added to a volume group that is associated with a draft storage partition.

User response

Disable Transparent Cloud Tiering for this volume.

CMMVC1238E The command failed because the volume is part of a Flash Copy relationship.

Explanation

The volume specified is part of a flash copy relationship, so cannot be added to a volume group that is associated with a draft storage partition.

User response

Delete the flash copy relationship associated with this volume.

CMMVC1239E The command failed because the volume has the incorrect I/O group configuration.

Explanation

The volume specified has Access I/O groups configured or the volume is in the wrong I/O group, so cannot be added to a volume group that is associated with a draft storage partition.

User response

Remove Access I/O groups from the volume and ensure the volume is in the correct I/O group.

CMMVC1241E The command failed because the volume has SCSI Persistent Reservation Registrations.

Explanation

The volume specified has SCSI Persistent Reservation Registrations active, so cannot be added to a volume group that is associated with a draft storage partition.

User response

Remove SCSI Persistent Reservation Registrations from the volume.

CMMVC1243E The command failed because the volume group is part of a storage partition that has a replication policy.

Explanation:
Snapshots of volume groups in a partition are by default cloned into the same partition, but snapshots cannot be cloned into a partition that is configured for HA.

User response:
To clone the snapshot outside of the partition, use the **-nopartition** parameter in the command-line interface.

CMMVC1244E **The action failed because the maximum number of schedules has been reached.**

Explanation

A snapshot policy cannot have more than two schedules.

User response

User may create another policy with the required schedules.

CMMVC1245E **The command failed because the storage partition already has a disaster recovery link configured.**

Explanation:
It is only possible to make a disaster recovery link to a storage partition that does not have a disaster recovery link configured.

User response:
Only use this command to make a disaster recovery link to a storage partition which has no disaster recovery link configured, or remove the configured disaster recovery link using **-removedrlink**.

CMMVC1247E **The command failed because the volume is currently restoring.**

Explanation

The volume specified is currently restoring, so cannot be added to a volume group that is associated with a draft storage partition.

User response

Wait for the restore to complete and retry the operation.

CMMVC1250E **The command failed because data reduction pools and replication policies cannot both be configured on this system.**

Explanation

This hardware type does not support data reduction pools and replication policies to be configured at the same time.

CMMVC1254E **The command failed because attributes for IPV4 and IPV6 cannot be configured for the same port.**

Explanation:
Cannot mix IPv4 and IPv6 parameter together and passed to this command.

User response:
Either pass all IPv4 parameters or IPv6 parameters.

CMMVC1257E **The command failed because the specified patch is not auto-updateable.**

Explanation:
The specified patch cannot use the auto-update process.

User response:
Retry the command with patch that can be auto-updated.

CMMVC1260E **The command failed because the specified volumes are in different volume groups.**

Explanation

The volumes that are specified on the command-line interface with the **-volumes** argument are not all in the same volume group.

User response

Specify volumes that are in the same volume group.

CMMVC1261E **The command failed because the volume group is not in a storage partition and the replication policy has an async-dr topology.**

Explanation:
The user has attempted to assign an async-dr replication policy to a volume group that is not in a storage partition. The volume group must be in a storage partition with an associated async-dr replication policy, and the replication policy being assigned must be the same as the replication policy associated with the storage partition the volume group is in.

User response:
An async-dr replication policy can only be assigned to a volume group that is associated with a storage partition associated with the same async-dr replication policy. Retry the command with a valid volume group and async-dr replication policy.

CMMVC1262E **The command failed because the volume group is not associated with the same storage partition as the async-dr replication policy.**

Explanation:
The user has attempted to assign an async-dr replication policy to a volume group that is not associated with the same storage partition as the async-dr replication policy. The volume group must be in a storage partition associated with the same async-dr replication policy as the replication policy being assigned to the volume group.

User response:
An async-dr replication policy can only be assigned to a volume group that is associated with a storage partition associated with the same async-dr replication policy. Retry the command with a valid volume group and async-dr replication policy.

CMMVC1263E The command failed because the replication policy is not associated with the correct storage partition.

Explanation:
The user has attempted to create a volume group with an async-dr replication policy that is not associated with the storage partition specified. An async-dr replication policy is associated with a storage partition and therefore the volume group must be created in the same partition.

User response:
Retry the command with the storage partition associated with the specified async-dr replication policy, or retry the command with a replication policy in the associated partition.

CMMVC1264E The command failed because the partner system is running a software version prior to 8.7.0.0.

Explanation

Storage partitions cannot be configured for high availability while the partner is running a code version older than 8.7.0.0 due to missing functionality in the older code version. Configuring high availability for a storage partition is only permitted at the time it is created while the systems are running these code versions.

User response

- Upgrade the partner system to 8.7.0.0 or later and retry the command.
- Create a new storage partition and configure it to be highly available at creation.

CMMVC1265E The command failed as this system is already a member of a Flash Grid.

Explanation

If a system is already a member of a FlashSystem grid:

- It cannot be added to another FlashSystem grid.
- A new FlashSystem grid cannot be created.

User response:
No action.

CMMVC1266E The command failed as this system is not the owner of the Flash Grid.

Explanation

Some operations on the FlashSystem grid, such as accepting new systems, are only permitted from the system that is the owner of the FlashSystem grid.

User response:
Log in to the system that is the owner of the FlashSystem grid and retry the operation.

CMMVC1267E The command failed as the specified system is already a member of the Flash Grid.

Explanation

The system cannot be accepted into the FlashSystem grid as a system with the same IP address is already a member of the FlashSystem grid.

User response:
No action.

CMMVC1268E The command failed as the system is not requesting to join this Flash Grid.

Explanation

The system cannot be accepted into the FlashSystem grid because it is not expected to join. This error can occur if the IP address given on the request to join does not match the IP address of the FlashSystem grid owner, or if the system does not make a request to join the FlashSystem grid.

User response:
Make sure the FlashSystem grid owner IP address that is given in the join command is correct, and that a join request is made to the correct FlashSystem grid.

CMMVC1269E The command failed as the host is not associated with a storage partition configured for high availability.

Explanation

Creating a host with a location, or changing the location of a host, is only permitted if the host is associated with a storage partition configured for high availability.

User response

Do not use location. Associate the specified storage partition with an High Availability replication policy, or add the host to a partition add the host to partition that has a High Availability replication policy.

CMMVC1270E **The command failed as the system does not meet the requirements to be a member of a Flash Grid.**

Explanation

For a system to be a member of a FlashSystem grid, it must meet the following prerequisites:

1. The system must be in the replication layer to ensure that partnerships can be created automatically between systems in the FlashSystem grid for high availability, replication, and migration.
2. The **multiple_io_groups** compatibility mode of the system must be disabled to ensure that the system configures a single I/O group only as required for Storage Partitions.

User response:
Ensure that the system is in the replication layer and the multiple_io_groups compatibility mode is disabled. To convert to a single I/O group system, ensure that all volumes are moved and migrated before you remove all additional I/O groups.

CMMVC1271E **The command failed as the multiple I/O groups feature cannot be enabled as this system is a member of a Flash Grid.**

Explanation

The systems in a FlashSystem grid support a maximum of one I/O group in the system as the FlashSystem grid is used for scaling instead of I/O groups.

User response:
Use the FlashSystem grid instead of I/O groups.

CMMVC1272E **The command failed as the layer must be set to replication as this system is a member of a Flash Grid.**

Explanation

The system must be in the replication layer to ensure that partnerships can be created automatically between systems in the FlashSystem grid to support high availability, replication, and migration.

User response:
No action.

CMMVC1274E **The command failed as the trust store entry is being used for another member of the Flash Grid.**

Explanation

The specified trust store entry is used for another system in the FlashSystem grid, which prevents removal of the entry or its usage for another member of the FlashSystem grid. It also prevents disabling the REST API usage.

User response:
If required, replace the certificate. Remove the member from the FlashSystem grid first to remove the entry.

CMMVC1275E **The command failed as the trust store entry cannot be used for a member of a Flash Grid.**

Explanation

The specified trust store entry cannot be used as it does not meet the following requirements:

- The entry exists.
- REST API usage for the entry is enabled.
- The entry is not the common certificate store.

User response:
Ensure that the entry is valid and REST API usage for the trust store entry is enabled.

CMMVC1279E **The command failed because the remote storage partition already has a disaster recovery link configured.**

Explanation:
It is only possible to make a disaster recovery link to a remote storage partition that does not have a disaster recovery link configured.

User response:
Only use this command to make a disaster recovery link to a remote storage partition which has no disaster recovery link configured, or remove the configured disaster recovery link using **-removedrlink**.

CMMVC1280E **The command failed because the specified remote system is the local system.**

Explanation:
It is only possible to make a disaster recovery link to a valid remote system, not the local system.

User response:
Only use this command to make a disaster recovery link by specifying a valid remote system and not the local system.

CMMVC1281E **The command failed because the remote partition UUID is not valid.**

Explanation:
An invalid remote partition UUID has been provided. It is only possible to make a disaster recovery link using a valid remote partition UUID.

User response:
Only use this command to make a disaster recovery link by specifying a valid remote partition UUID.

CMMVC1282E **The command failed because it is not permitted to make a disaster recovery link from a local non-high-availability storage partition to a high availability storage partition.**

Explanation:
The user has attempted to make a disaster recovery link from a local non-high availability storage partition to a high availability storage partition. If one of the storage partitions to be linked is high availability then the command must be run on the high availability partition.

User response:
Run the **chpartition -makedrlink** command on the high availability storage partition.
Make sure to use the correct arguments for **-remotedrlinkedpartitionuuid** and **-remotesystem** parameters.

CMMVC1284E **The command failed because the source storage partition has a disaster recovery link.**

Explanation:
The user has attempted to merge from a storage partition with a disaster recovery link, the command will fail if the source storage partition has a disaster recovery link.

User response:
Only merge from a storage partition which does not have a disaster recovery link. Use lspartition to check if a storage partition has a disaster recovery link configured.

CMMVC1286E **The command failed because there are no drives of the specified drive class.**

Explanation:
Unable to create the array as the specified drive class does not exist.

User response:
Specify a drive class that exists.

CMMVC1287E **The command failed because parameter clusterip/clusterip_6 is mandatory to configure a new system IP.**

Explanation:
New system IP address cannot be configured with IP address.

User response:
Either pass **clusterip** parameter to configure IPv4 system IP address or **clusterip_6** parameter to configure IPv6 system IP address.

CMMVC1290E **The command failed because the replication policy can not be applied to volume groups in a storage partition.**

Explanation:
The user has attempted to create a volume group in a storage partition with an 2-site-async-dr replication policy.

User response:
Retry the command using a replication policy with the async-dr topology instead, or create the volumegroup outside the storage partition.

CMMVC1291E **The command failed because the disaster recovery link can not be removed as the storage partition contains disaster recovery replicated volume groups.**

Explanation:
The user has attempted to remove a disaster recovery link from a storage partition that contains volume groups configured for disaster recovery replication. It is only possible to remove a disaster recovery link from a storage partition when there are no volume groups configured for disaster recovery.

User response:
Use the **lsvolumegroupreplication** command to identify all volume groups in the storage partition configured for disaster recovery, unassign the replication policies and retry the command.

CMMVC1297E **The command failed because the replication policy can not be**

applied to volume groups not in a storage partition.

Explanation:
The user has attempted to create a volume group with an async-dr replication policy, but the volume group is not in a storage partition.

User response:
Retry the command using a replication policy with the 2-site-async-dr topology instead, or create the volumegroup inside a storage partition.

CMMVC1298E **The command failed because one or more volume groups associated with the storage partition have a 2-site-async-dr replication policy assigned.**

Explanation:
A replication policy cannot be assigned to a storage partition if there are any associated volume groups with a 2-site-async-dr replication policy assigned.

User response:
Identify the associated volume groups that have a 2-site-async-dr replication policy assigned, then complete the migration to an async-dr replication policy for each volume group.

CMMVC1299E **The command failed because the specified remote system is not valid.**

Explanation:
A volume group associated with the storage partition has a 2-site-async-dr replication policy assigned that is configured for replication to a different remote system than the one specified.

User response:
Identify the remote system that associated volume groups are configured for replication to, then retry the command.

CMMVC1300E **The command failed because the set of 2-site-async-dr volume groups associated with the local partition does not match the set of 2-site-async-dr volume groups associated with the specified remote partition.**

Explanation:
The user has attempted to make a disaster recovery link and the set of 2-site-async-dr volume groups associated with the local partition does not match the set of 2-site-async-dr volume groups associated with the remote partition. Any 2-site-async-dr volume groups that are being migrated to async-dr replication policies must have been added to both the local and remote partitions.

User response:
Use the **lsvolumegroup** command to identify all of the volume groups that are in the local and remote partitions. For each volume group that has a 2-site-async-dr replication policy and is in only one of the partitions, use the **mergepartition** command to add it to the required partitions.

CMMVC1304E **The command failed because the source storage partition contains 2-site-async-dr volume groups and the target storage partition has a disaster recovery link.**

Explanation:
The user has attempted to merge from a storage partition with associated 2-site-async-dr volume groups into a storage partition that has a disaster recovery link configured.

User response:
A storage partition containing 2-site-async-dr volume groups cannot be merged into a storage partition that has a disaster recovery link.

CMMVC1305E **The command failed because the source and target storage partitions have 2-site-async-dr volume groups configured for replication to different remote systems.**

Explanation:
The user has attempted to merge storage partitions that both contain volume groups with 2-site-async-dr replication policies. This is only permitted when all such volume groups are configured for replication to the same remote system.

User response:
If both the source and target storage partition have associated 2-site-async-dr volume groups they can only be merged if they are all configured for replication to the same remote system. Identify a target partition that satisfies that requirement and retry the command.

CMMVC1311E **The command failed because the volume group does not have a replication policy assigned.**

Explanation:
This error occurs when replication is restarted on an independent volume group that does not have a replication policy assigned. You can change the replication mode of a volume group to production only if a replication policy is assigned.

User response:
Assign a replication policy to the volume group and retry the command.

CMMVC1321E **The command failed because a storage partition with the**

specified UUID does not exist on the remote system.

Explanation:
The provided remote storage partition UUID does not exist. You can establish a disaster recovery link by using an existing UUID of the remote storage partition only.

User response:
Use the **chpartition** command to establish a disaster recovery link to an existing remote storage partition UUID. Use **lspartition** on the remote system to get the UUID of the remote storage partition.

CMMVC1322E **The command failed because the storage partition does not have a disaster recovery link established.**

Explanation:
The user has attempted to create an async-dr replication policy in a storage partition that does not have a disaster recovery link established. An async-dr replication policy can only be created in a storage partition that has been linked for disaster recovery.

User response:
Establish a disaster recovery link for the storage partition using **chpartition -makedrlink** and then retry the command.

CMMVC1323E **The command failed because the storage partition contains volume groups configured for disaster recovery.**

Explanation:
The user has attempted to remove a storage partition which contains a disaster recovery replicated volume group. It is only possible to remove a storage partition that does not contain a disaster recovery replicated volume group.

User response:
Use the **lsvolumegroup** command to identify the disaster recovery replicated volume groups in the storage partition and remove them, or use **lsreplicationpolicy** to identify the disaster recovery replication polices in the storage partition and remove them, then retry the command.

CMMVC1331E **The command failed because it is not permitted to configure replication for disaster recovery and high availability on the same partnership.**

Explanation

This error occurs when an attempt is made to:

- Establish a disaster recovery link between storage partitions on systems in a partnership with high availability configured, or
- Configure high availability on systems in a partnership with a disaster recovery link established.

Replication for disaster recovery and high availability cannot be configured on the same partnership.

User response:
Configure replication for disaster recovery and high availability by using separate systems.

CMMVC1357E **The command failed because it is not permitted to make a disaster recovery link between two high availability storage partitions.**

Explanation:
This error occurs when an attempt is made to make a disaster recovery link from a local high availability storage partition to a remote high availability storage partition. This attempt results in a four-site configuration, which is not supported.

User response:
If you want to establish a disaster recovery link, then remove high availability from one of the storage partitions by using the **chpartition -noreplicationpolicy** command.

CMMVC1397E **The current user is not allowed to create, change or remove a locally defined remote user when two person integrity is enabled.**

Explanation

This error occurs when a restricted security administrator attempts any operation on a target user that is remote.

User response

To resolve this error, complete the following steps:

- Create a two person integrity request, note down the user name and ID assigned to the request.
- Locate a user with Security Administrator or Restricted Security Administrator access, provide the user name and ID assigned to the two person integrity request and get that user to approve it.
- After it is approved, retry the desired operation on the target remote user before the temporary access expires.

CMMVC4041E **You must use an 0x parameter for the -lba and -vdisklba parameters.**

Explanation:

The parameter format for lsmdisklba needs explaining if using the incorrect format.

User response:
Retry the command using the described parameter format.

CMMVC4042E You must use an 0x parameter for the -lba and -mdisklba parameters.

Explanation:
The parameter format for lsmdisklba needs explaining if using the incorrect format.

User response:
Retry the command using the described parameter format.

CMMVC4043E You must use an 0x parameter for the -mdisklba parameter.

Explanation:
The parameter format for lsmdisklba needs explaining if using the incorrect format.

User response:
Retry the command using the described parameter format.

CMMVC4044E You must use an 0x parameter for the -drivelba parameter.

Explanation:
The parameter format for lsmdisklba needs explaining if using the incorrect format.

User response:
Retry the command using the described parameter format.

CMMVC5000I No message was found for major rc *MAJOR_RC* , minor rc *MINOR_RC* , for action/view id *ACTION_VIEW_ID* .

Explanation

A message is missing.

User response

Contact the support center.

CMMVC5700E The parameter list is not valid.

Explanation

You have entered a list of parameters that is not supported for the command.

User response

Specify a parameter list that is supported for the command, and resubmit the command.

CMMVC5701E No object ID was specified.

Explanation

The command that you have submitted requires that you specify an object identifier name or ID number, and you did not specify an object identifier.

User response

Specify an object ID, and resubmit the command.

CMMVC5702E *VALUE* is below the minimum level.

Explanation

You entered the specified string as a value for a parameter. The parameter requires a minimum value, and the specified string is less than the required minimum value.

User response

Specify a value that is supported by the parameter, and resubmit the command.

CMMVC5703E The value or list starting with *VALUE* is above the maximum permitted for that value or has exceeded the number of items allowed in a list.

Explanation

The specified string was entered as a value for a parameter. The string is either a stand-alone value or the first value in a list of values. If the string is a stand-alone 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.

This error can occur after the **svctask chsystemethernet** command is run, if Ethernet traffic priority values are specified that are not integers in the range 0 - 7.

User response

Specify a value or list of values that is supported by the parameter, and resubmit the command.

CMMVC5704E *VALUE* is not divisible by the permitted step value.

Explanation

You have entered the specified string as a value for a parameter. The string is not a supported value for the parameter. One requirement is that the value is an even multiple of 16, and the specified string does not meet that requirement.

User response

Specify a value that is supported by the parameter, and resubmit the command.

CMMVC5705E A required parameter is missing.

Explanation

The command that you have submitted has at least one required parameter that you have not entered.

User response

Specify all of the required parameters, and resubmit the command.

CMMVC5706E An invalid argument has been entered for the *PARAMETER* parameter.

Explanation

You have entered a value for the specified parameter and the value is not supported for the parameter. The parameter supports a specific set of values.

User response

Specify a value that is supported by the parameter, and resubmit the command.

CMMVC5707E Required parameters are missing.

Explanation

The user has not entered all of the required parameters.

User response

Correct the CLI syntax and rerun the command.

CMMVC5708E A command-line interface parameter is missing its associated arguments.

Explanation

The user has entered one of the command parameters but has not entered any associated arguments.

User response

Correct the CLI syntax and rerun the command.

CMMVC5709E *VALUE* is not a supported parameter.

Explanation

The specified string is not a supported parameter for the command that you have entered.

User response

Specify the correct parameter, and resubmit the command.

CMMVC5711E *VALUE* is not valid data.

Explanation

You have entered the specified string as a value for a parameter. The string is not a supported value for the parameter (for example, the specified string is an incorrect path).

User response

Specify a value that is supported by the parameter, and resubmit the command.

CMMVC5712E Required data is missing.

Explanation

You have entered an incomplete command.

User response

Specify command completely, and resubmit the command.

CMMVC5713E Some parameters are mutually exclusive.

Explanation

Certain commands have two or more parameters that are mutually exclusive. You have submitted a command using at least two mutually exclusive parameters.

User response

Specify a supported combination of parameters, and resubmit the command.

CMMVC5714E The parameter list is empty.

Explanation

Certain parameters require one or more values in a colon separated parameter list. You have specified at least one parameter without the required parameter list.

User response

Specify at least one value for all parameters that require a value, and resubmit the command.

CMMVC5715E The parameter list does not exist.

Explanation

Certain parameters require one or more values in a colon separated parameter list. You have specified at least one parameter without the required parameter list.

User response

Specify at least one value for all parameters that require a value, and resubmit the command.

CMMVC5716E Non-numeric data was entered for the numeric field *FIELD*. Enter a numeric value.

Explanation

The specified string was entered as a value for a parameter that supports only numeric values.

This error can occur after the **svctask chsystemethernet** command is run, if Ethernet traffic priority values are specified that are not integers in the range 0 - 7.

User response

Specify a numeric value in the numeric field, and resubmit the command.

CMMVC5717E No match was found for the specified unit.

Explanation

Certain parameters allow a user to specify a data unit such as mb or kb. You have entered a data unit for a parameter that supports data units, but the data unit that you have entered is not a supported data unit for the parameter.

User response

Specify the correct data unit, and resubmit the command.

CMMVC5718E An unexpected return code was received.

Explanation

The command has completed, but the acknowledgement of the command completion contains a return code that is not defined.

User response

Determine whether or not the command has succeeded. If the command has not succeeded, resubmit the command. If the problem persists, contact IBM technical support for assistance.

CMMVC5719E A value of *VALUE* requires the parameter *PARAMETER* to be specified.

Explanation

Certain commands have required combinations of parameters based on either the entry of a parameter or the value for a parameter. When you enter the specified value, you must enter the specified parameter.

User response

Specify the required parameter, and resubmit the command.

CMMVC5721E *VALUE* is not a valid time stamp format. The valid time stamp format is YYMMDDHHMMSS.

Explanation

The specified value is not a valid time stamp format. The valid format is YYMMDDHHMMSS.

User response

Use the correct time stamp format, and resubmit the command.

CMMVC5722E *VALUE* contains a month value that is not valid. The valid time stamp format is YYMMDDHHMMSS.

Explanation

The month value (MM) that you have specified is not valid.

User response

Specify a valid month value, and resubmit the command.

CMMVC5723E ***VALUE*** **contains a day value that is not valid. The valid time stamp format is YYMMDDHHMMSS.**

Explanation

The day value (DD) that you have specified is not valid.

User response

Specify a valid day value, and resubmit the command.

CMMVC5724E ***VALUE*** **contains an hour value that is not valid. The valid time stamp format is YYMMDDHHMMSS.**

Explanation

The hour value (HH) that you have specified is not valid.

User response

Specify a valid hour value, and resubmit the command.

CMMVC5725E ***VALUE*** **contains a minutes value that is not valid. The valid time stamp format is YYMMDDHHMMSS.**

Explanation

The minutes value (MM) that you have specified is not valid.

User response

Specify a valid minutes value, and resubmit the command.

CMMVC5726E ***VALUE*** **contains a seconds value that is not valid. The valid time stamp format is YYMMDDHHMMSS.**

Explanation

The seconds value (SS) that you have specified is not valid.

User response

Specify a valid seconds value, and resubmit the command.

CMMVC5727E ***VALUE*** **is not a valid filter.**

Explanation

You can filter the output of some views by using the -filtervalue parameter. The specified string that you have entered is not a supported value for the -filtervalue parameter in this view.

User response

Ensure that you use a supported value for the -filtervalue parameter, and resubmit the command.

CMMVC5728E ***%1*** **is not a valid time format. The valid time format is MMDDHHmmYYYY with YYYY<2070.**

Explanation

The specified value should be in the format MMDDHHmmYYYY with YYYY less than 2070.

User response

Follow the correct format, and resubmit the command.

CMMVC5729E **One or more components in the list is not valid.**

Explanation

Certain parameters support one or more items of data in a colon separated list. At least one of the items in the list that you have entered is not correct.

User response

Ensure that you enter supported values in the list, and resubmit the command.

CMMVC5730E ***VALUE*** **is only valid when** ***VALUE*** **has a value of** ***VALUE*** **.**

Explanation:
The specified command and parameter combination that you have entered requires the specified parameter value.

User response:
Ensure that you specify the correct parameter value for the command and parameter combination that you enter, and resubmit the command.

CMMVC5731E ***VALUE*** **can only be entered when** ***VALUE*** **has been entered.**

Explanation

Certain commands have required combinations of parameters based either on the inclusion of a specified parameter, or on the value entered for a specified parameter. When you include the first specified string in the command, you must enter the second specified string as a parameter.

User response

Ensure that you enter a supported combination or parameters and values, and resubmit the command.

CMMVC5733E Enter at least one parameter.

Explanation

You must specify at least one parameter for the command that you have submitted.

User response

Specify at least one parameter, and resubmit the command.

CMMVC5734E A combination of values was entered that is not valid.

Explanation

You have specified a combination of values that is not correct.

User response

Specify a supported combination of values, and resubmit the command.

CMMVC5735E The name entered is not valid. Enter an alphanumeric string that does not start with a number.

Explanation

The first character of an object name cannot be numeric.

User response

Specify an alphanumeric string that does not start with a numeric, and resubmit the command.

CMMVC5737E The parameter *PARAMETER* has been entered multiple times. Enter the parameter only one time.

Explanation

The specified parameter was entered more than once.

User response

Delete all duplicate parameters, and resubmit the command.

CMMVC5738E The argument *ARGUMENT* contains too many characters.

Explanation

The field length of the specified argument is longer than the maximum supported field length for the argument.

User response

Specify the correct argument, and resubmit the command.

CMMVC5739E The argument *ARGUMENT* does not contain enough characters.

Explanation

The field length of the specified argument is less than the minimum supported field length for the argument.

User response

Specify the correct argument, and resubmit the command.

CMMVC5740E The filter flag *VALUE* is not valid.

Explanation

You can filter the output of some views by using the -filtervalue parameter. The specified string that you have entered is not a supported value for the -filtervalue parameter in this view.

User response

Ensure that you use a supported value for the -filtervalue parameter, and resubmit the command.

CMMVC5741E The filter value *VALUE* is not valid.

Explanation

You can filter the output of some views by using the -filtervalue parameter. Each filter has an associated value. The syntax is -filtervalue filter=value. The specified string that you have entered is not a supported value for the -filtervalue filter that you specified in this view.

User response

Ensure that you use a supported value for the -filtervalue filter that you specify, and resubmit the command.

CMMVC5742E **A specified parameter is out of its valid range.**

Explanation

You have entered data that is not in the range of values that is supported for the parameter that you have entered.

User response

Ensure that you enter data values that are supported for the parameter that you enter, and resubmit the command.

CMMVC5743E **A specified parameter does not comply with the step value.**

Explanation

A parameter was specified that does not comply with the step value.

User response

Specify the correct parameter, and resubmit the command.

CMMVC5744E **Too many objects were specified in the command.**

Explanation

There were too many objects specified in the command.

User response

Specify the correct object, and resubmit the command.

CMMVC5745E **Too few objects were specified in the request.**

Explanation

There were not enough objects specified in the command.

User response

Specify the correct object, and resubmit the command.

CMMVC5746E **The requested operation cannot be applied to the object specified.**

Explanation

The requested operation is not valid for this object.

User response

Specify a valid operation, and resubmit the command.

CMMVC5747E **The action requested is invalid - internal error.**

Explanation

The operation that was requested is not valid.

User response

Specify the correct operation, and resubmit the command.

CMMVC5748E **The action requested is invalid - internal error.**

Explanation

The operation that was requested is not valid.

User response

Specify the correct operation, and resubmit the command.

CMMVC5749E **The memory dump file name that is specified exists.**

Explanation

The dump file name that was specified exists.

User response

Specify a different dump file name, and resubmit the command.

CMMVC5750E **The dump file could not be created - the file system is probably full.**

Explanation

The dump file was not created. The file system might be full.

User response

Reduce the size of the file system by deleting obsolete logs or other unnecessary files, then retry the command.

CMMVC5751E **The dump file could not be written to.**

Explanation

The dump file could not be written to disk.

User response

Not applicable.

CMMVC5752E Request failed. The object contains child objects, these must be deleted first.

Explanation

The operation failed because the specified object contains child objects.

User response

Delete the child objects, and resubmit the command.

CMMVC5753E The specified object does not exist or is not a suitable candidate.

Explanation

The specified object does not exist or is not a suitable candidate.

User response

Specify the correct object, and resubmit the command.

CMMVC5754E The specified object does not exist, or the name supplied does not meet the naming rules.

Explanation

The user has specified a either a volume name or volume group name that does not exist.

User response

Specify a volume or volume group that exists and rerun the command.

CMMVC5755E Cannot create as the sizes of the specified objects do not match.

Explanation

The sizes of the specified objects do not match.

User response

Not applicable.

CMMVC5756E Cannot perform the request as the object id is already mapped to another object or is the subject of an FC or RC relationship.

Explanation

The operation failed because the specified object is already mapped.

User response

Specify a different object, and resubmit the command.

CMMVC5757E Self Defining Structure (SDS) defaults not found - internal error.

Explanation

The defaults for the self describing structure were not found.

User response

Not applicable.

CMMVC5758E Object name already exists.

Explanation

The object name already exists.

User response

Specify a unique object name, and resubmit the command.

CMMVC5759E An internal error has occurred - memory could not be allocated.

Explanation

The memory cannot be allocated.

User response

Not applicable.

CMMVC5762E The request did not complete before the timeout period expired.

Explanation

The operation failed because the timeout period expired.

User response

Resubmit the command.

CMMVC5763E The node failed to go online.

Explanation

The node failed to go online.

User response

Not applicable.

CMMVC5764E The mode change request is invalid - internal error

Explanation

The specified mode change is not valid.

User response

Specify a different mode, and resubmit the command.

CMMVC5765E The object specified is no longer a candidate - a change occurred during the request.

Explanation

The specified object is no longer a candidate. A change occurred during the request.

User response

Specify a different object, and resubmit the command.

CMMVC5767E One or more of the parameters specified are invalid or a parameter is missing.

Explanation

One or more of the specified parameters is not valid.

User response

Specify the correct parameter, and resubmit the command.

CMMVC5769E The requested operation requires all nodes to be online - one or more nodes are not online.

Explanation

The operation requires that all nodes be online. One or more nodes are not online.

User response

Check that each node is online, and resubmit the command.

CMMVC5770E The SSH key file supplied is invalid.

Explanation

The file for the SSH key is not valid.

User response

Specify a different file, and resubmit the command.

CMMVC5771E The operation requested could not complete, usually due to child objects existing. To force the operation, specify the force flag.

Explanation

The operation failed, probably, because the object contains child objects.

User response

Specify the -force flag to complete the operation, and resubmit the command.

CMMVC5772E The operation requested could not be performed because an update is in progress.

Explanation

The operation failed because an update is in progress.

User response

Wait for the update to complete, and resubmit the command.

CMMVC5773E The object selected is in the wrong mode to perform the requested operation.

Explanation

The operation failed because the selected object is in the wrong mode.

User response

Specify the correct mode, and resubmit the command.

CMMVC5774E The userid supplied is not valid.

Explanation

The userid is not valid.

User response

Specify a different userid, and resubmit the command.

CMMVC5775E The directory attribute specified is not valid.

Explanation

The directory attribute is not valid.

User response

Specify a different directory, and resubmit the command.

CMMVC5776E The directory listing could not be retrieved.

Explanation

The directory listing could not be retrieved.

User response

Specify a different directory listing, and resubmit the command.

CMMVC5777E The node could not be added to the IO Group, because the other node in the IO Group is in the same power domain.

Explanation

The node was not added to the I/O group because the other node in the I/O Group is in the same power domain.

User response

Specify a different node from another I/O group, and resubmit the command.

CMMVC5778E Cannot create another cluster, a cluster already exists.

Explanation

The cluster was not created because one already exists.

User response

Not applicable.

CMMVC5780E The action could not be completed using the Remote Cluster name. Use the Remote Cluster Unique ID instead.

Explanation

The unique ID of the remote cluster is required for this command.

User response

Specify the unique ID of the remote cluster, and resubmit the command.

CMMVC5781E The cluster ID specified is invalid.

Explanation

The cluster ID is not valid.

User response

Specify a different cluster ID, and resubmit the command.

CMMVC5782E The object specified is offline.

Explanation

The object is offline.

User response

Specify an object that is online, and resubmit the command.

CMMVC5783E The information is not available to complete this command.

Explanation

This error is only returned when the node is in the service state.

User response

None.

CMMVC5784E The cluster name specified is not unique, specify the cluster using the cluster ID.

Explanation

The cluster name is not unique.

User response

Specify the cluster using the cluster ID, and resubmit the command.

CMMVC5785E The file name that is specified contains an invalid character.

Explanation

The file name contains an invalid character.

User response

Specify a valid file name, and resubmit the command.

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

Explanation

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

User response

Not applicable.

CMMVC5787E **The cluster was not created because a cluster already exists.**

Explanation

The cluster was not created because a cluster already exists.

User response

Not applicable.

CMMVC5788E **The service IP address is not valid.**

Explanation

The service IP address is not valid.

User response

Specify the correct service IP address, and resubmit the command.

CMMVC5789E **The cluster was not modified because the IP address, subnet mask, service address, SNMP address, or gateway address is not valid.**

Explanation

The cluster was not modified because the IP address, subnet mask, service address, SNMP address, or gateway address is not valid.

User response

Specify all correct attributes, and resubmit the command.

CMMVC5790E **The node was not added to the cluster because the maximum number of nodes has been reached.**

Explanation

The node was not added to the cluster because the maximum number of nodes has been reached.

User response

Not applicable.

CMMVC5791E **The action failed because an object that was specified in the command does not exist.**

Explanation

An entity that was specified in the command does not exist, therefore the action failed.

User response

Specify the correct entity, and resubmit the command.

CMMVC5792E **The action failed because the I/O group is used for recovery.**

Explanation

The recovery group does not support the command line interface.

User response

Use the management GUI or the **lsiogrp** command to see the list of active I/O groups. Rerun the command with a valid I/O group ID or name.

CMMVC5793E **The node was not added to the cluster because the I/O group already contains a pair of nodes.**

Explanation

The node was not added to the cluster because the I/O group already contains a pair of nodes.

User response

Not applicable.

CMMVC5794E **The action failed because the node is not a member of the cluster.**

Explanation

The node is not a member of the cluster, therefore the action failed.

User response

Specify a node that is contained in the cluster, and resubmit the command.

CMMVC5795E The node was not deleted because an update is in progress.

Explanation

The node was not deleted because an update is in progress.

User response

Wait for the update to complete, and resubmit the command.

CMMVC5796E The action failed because the I/O group that the node belongs to is unstable.

Explanation

A previous configuration command might not yet have completed.

User response

Wait for the previous command to complete, and resubmit the command.

CMMVC5797E The node was not deleted because this is the last node in the I/O group and there are volumes associated with the I/O group.

Explanation

The specified node is the last node in the I/O group and there are volumes associated with the I/O group, therefore the node could not be deleted.

User response

Not applicable.

CMMVC5798E The action failed because the node is offline.

Explanation

The action failed because the node is offline.

User response

Specify a node that is online, and resubmit the command.

CMMVC5799E The shut down was not successful because there is only one online node in the I/O group.

Explanation

There is only one online node is the I/O group, therefore the shut down operation was not successful.

User response

Not applicable.

CMMVC5800E The action failed because an entity that was specified in the command does not exist.

Explanation

The entity that was specified in the command does not exist, therefore the action failed.

User response

Specify a different entity, and resubmit the command.

CMMVC5801E The update of the cluster could not proceed because every node in the cluster must be online. Either delete the node that is offline or bring the node online and resubmit the command

Explanation

The update of the cluster could not proceed because every node in the cluster must be online.

User response

Either delete the node that is offline or bring the node online, and resubmit the command.

CMMVC5802E The update of the cluster could not proceed because there is an I/O group in the cluster that contains only one node. The update requires that each node in an I/O group be shut down and restarted. If there is only one node in an I/O group, I/O operations could be lost if I/O operations are not stopped before beginning the update.

Explanation

The update of the cluster could not proceed because there is an I/O group in the cluster that contains only one node. The update requires that each node in an I/O group be shut down and restarted. If there is only one node in an I/O group, I/O operations could be lost if I/O operations are not stopped before beginning the update.

User response

Either update the cluster using the -force option or specify a different node, and resubmit the command.

CMMVC5803E The entry in the event log was not marked because the error is already fixed or unfixed, or the sequence number could not be found.

Explanation

The entry in the event log was not marked because the sequence number was not found.

User response

Not applicable.

CMMVC5804E The action failed because an object that was specified in the command does not exist.

Explanation

The entity that was specified in the command does not exist, therefore the action failed.

User response

Specify a different entity, and resubmit the command.

CMMVC5805E The progress information was not returned because the FlashCopy statistics are not ready yet.

Explanation

The progress information was not returned because the FlashCopy statistics are not ready yet.

User response

Not applicable.

CMMVC5806E The action failed because an object that was specified in the command does not exist.

Explanation

The entity that was specified in the command does not exist, therefore the action failed.

User response

Specify a different entity, and resubmit the command.

CMMVC5807E The action failed because the managed disk (MDisk) cannot be changed to the specified mode.

Explanation

The action failed because the managed disk (MDisk) cannot be changed to the specified mode.

User response

Not applicable.

CMMVC5808E The action failed because the managed disk (MDisk) does not exist.

Explanation

The action failed because the managed disk (MDisk) does not exist.

User response

Specify a different MDisk, and resubmit the command.

CMMVC5809E The tracing of I/O operations was not started because it is already in progress.

Explanation

The tracing of I/O operations was not started because it is already in progress.

User response

Not applicable.

CMMVC5810E The action failed because the specified resource was unavailable.

Explanation

The resource specified in the action was not available for use.

User response

Fix any errors associated with the specified resource, or reissue the command using an alternative resource.

CMMVC5811E **The quorum index number for the object was not set because the quorum disk does not exist.**

Explanation

An existing quorum disk must be specified before the quorum index number of the object can be set.

User response

Specify an existing quorum disk, and resubmit the command.

CMMVC5812E **The action failed because the managed disk (MDisk) is not in managed mode.**

Explanation

The action is only permitted on MDisks that are currently in the managed mode.

User response

Either add the MDisk to a storage pool, or specify a different MDisk.

CMMVC5813E **The quorum index number for the object was not set because the object has a sector size that is not valid.**

Explanation

The sector size of the specified object will not allow the quorum index number for the object to be set.

User response

Change the sector size of the specified object, or specify a different object, and resubmit the command.

CMMVC5814E **The quorum index number for the managed disk (MDisk) was not set because quorum is not allowed on one or more associated controllers.**

Explanation

The quorum index number for the managed disk (MDisk) was not set because quorum is not allowed on one or more associated controllers.

User response

Specify an MDisk that has quorum enabled on all of its associated controllers, and resubmit the command.

CMMVC5815E **The storage pool was not created because an entity that was specified in the command does not exist.**

Explanation

The storage pool was not created because an entity that was specified in the command does not exist.

User response

Specify a different entity, and resubmit the command.

CMMVC5816E **The action failed because an entity that was specified in the command does not exist.**

Explanation

The action failed because an entity that was specified in the command does not exist.

User response

Specify a different entity, and resubmit the command.

CMMVC5817E **The specified storage pool was invalid.**

Explanation

The storage pool was not renamed because the name was not valid.

User response

Specify a different storage pool name, and resubmit the command.

CMMVC5818E **The storage pool was not deleted because there is at least one MDisk in the group.**

Explanation

The storage pool was not deleted because there is at least one MDisk in the group.

User response

Not applicable.

CMMVC5819E **The managed disk (MDisk) was not added to the storage pool**

because the MDisk is part of another storage pool.

Explanation

The managed disk (MDisk) was not added to the storage pool because the MDisk is part of another storage pool.

User response

Not applicable.

CMMVC5820E **The managed disk (MDisk) was not added to the storage pool because an entity that was specified in the command does not exist.**

Explanation

The managed disk (MDisk) was not added to the storage pool because an entity that was specified in the command does not exist.

User response

Specify a different entity, and resubmit the command.

CMMVC5821E **The managed disk (MDisk) was not added to the storage pool because not enough MDisks were included in the list.**

Explanation

The managed disk (MDisk) was not added to the storage pool because not enough MDisks were included in the list.

User response

Include more MDisks in the list, and resubmit the command.

CMMVC5822E **The managed disk (MDisk) was not added to the storage pool because too many MDisks were included in the list.**

Explanation

The managed disk (MDisk) was not added to the storage pool because too many MDisks were included in the list.

User response

Delete the extra MDisks in the list, and resubmit the command.

CMMVC5823E **The managed disk (MDisk) was not deleted from the storage pool because the MDisk is part of another storage pool.**

Explanation

The managed disk (MDisk) was not deleted from the storage pool because the MDisk is part of another storage pool.

User response

Not applicable.

CMMVC5824E **The managed disk (MDisk) was not deleted from the storage pool because it does not belong to the storage pool.**

Explanation

The managed disk (MDisk) was not deleted from the storage pool because it does not belong to the storage pool.

User response

Not applicable.

CMMVC5825E **The managed disk (MDisk) was not deleted from the storage pool because a volume is allocated from one or more of the specified MDisks. A forced deletion is required.**

Explanation

The managed disk (MDisk) was not deleted from the storage pool because a volume is allocated from one or more of the specified MDisks.

User response

Specify the -force option, and resubmit the command.

CMMVC5826E **The volume was not created because an entity that was specified in the command does not exist.**

Explanation

The volume was not created because an entity that was specified in the command does not exist.

User response

Specify a different entity, and resubmit the command.

CMMVC5827E **The command failed as a result of either an inconsistency between two or more of the entered parameters, or an inconsistency between a parameter and the requested action.**

Explanation

The command failed as a result of an inconsistency between two or more of the entered parameters.

User response

Specify one parameter, and resubmit the command.

CMMVC5828E **The volume was not created because the I/O group contains no nodes.**

Explanation

The volume was not created because the I/O group contains no nodes.

User response

Not applicable.

CMMVC5829E **The image-mode or sequential-mode volume was not created because more than one managed disk (MDisk) is specified.**

Explanation

The image-mode or sequential-mode volume was not created because more than one MDisk is specified.

User response

Specify a different MDisk, and resubmit the command.

CMMVC5830E **The image-mode volume was not created because no managed disk (MDisk) was specified in the command.**

Explanation

The image-mode volume was not created because no managed disk (MDisk) was specified in the command.

User response

Specify a MDisk, and resubmit the command.

CMMVC5831E **The volume was not created because the preferred node for I/O operations is not part of the I/O group.**

Explanation

The volume was not created because the preferred node for I/O operations is not part of the I/O group.

User response

Specify a different node, and resubmit the command.

CMMVC5832E **The property of the volume was not modified because an entity that was specified in the command does not exist.**

Explanation

The property of the volume was not modified because an entity that was specified in the command does not exist.

User response

Specify a different entity, and resubmit the command.

CMMVC5833E **The property of the volume was not modified because there are no nodes in the I/O group.**

Explanation

The property of the volume was not modified because there are no nodes in the I/O group.

User response

Not applicable.

CMMVC5834E **The I/O group for the volume was not modified because the group is a recovery I/O group. To modify the I/O group, use the force option.**

Explanation

The I/O group for the volume was not modified because the group is a recovery I/O group.

User response

Specify the -force option, and resubmit the command.

CMMVC5835E **The volume was not expanded because an entity that was**

specified in the command does not exist.

Explanation

The volume was not expanded because an entity that was specified in the command does not exist.

User response

Specify a different entity, and resubmit the command.

CMMVC5836E **The volume was not shrunk because it is locked.**

Explanation

Commands might still be running in the background.

User response

Wait for all commands to complete. Use the lsmigrate command to view any migrates running in the background.

CMMVC5837E **The action failed because the volume is part of a FlashCopy mapping.**

Explanation

The action failed because the volume is part of a FlashCopy mapping.

User response

Specify a different volume that is not part of a FlashCopy mapping, and resubmit the command.

CMMVC5838E **The action failed because the volume is part of a Remote Copy mapping.**

Explanation

The action failed because the volume is part of a Remote Copy mapping.

User response

Specify a different volume that is not part of a Remote Copy mapping, and resubmit the command.

CMMVC5839E **The volume was not shrunk because an object that was specified in the command does not exist.**

Explanation

The volume was not shrunk because an object that was specified in the command does not exist.

User response

Specify a different object, and resubmit the command.

CMMVC5840E **The volume was not deleted because it is mapped to a host, or it is part of a FlashCopy mapping or Remote Copy relationship, or a cloud snapshot or restore operation is in progress, or it is involved in an image mode migration.**

Explanation

The volume was not deleted for one of the following reasons:

- It is mapped to a host
- It is part of a FlashCopy mapping
- It is in a remote copy relationship
- A cloud snapshot operation is in progress
- A restore operation is in progress

User response

Make sure that you specified the correct volume. If so, take one or more of the following actions:

- If the volume is mapped to a host, remove the host mapping.
- If the volume is part of a FlashCopy mapping, remove the mapping.
- If the volume is in a remote copy relationship, remove it from the relationship.
- If a cloud snapshot or restore operation is in progress, either wait for the operation to finish or cancel the operation.

When the volume is available for deletion, retry the delete command.

CMMVC5841E **The volume was not deleted because it does not exist.**

Explanation

The volume was not deleted because it does not exist.

User response

Specify a different volume, and resubmit the command.

CMMVC5842E **The action failed because an object that was specified in the command does not exist.**

Explanation

The action failed because an entity that was specified in the command does not exist.

User response

Specify a different entity, and resubmit the command.

CMMVC5843E **The host mapping was not created because the volume does not have a capacity greater than zero bytes.**

Explanation

The host map was not created because the volume does not have a capacity greater than zero bytes.

User response

Specify a volume in which its capacity is greater than zero bytes, and resubmit the command.

CMMVC5844E **The host mapping was not created because the SCSI logical unit number (LUN) ID is not valid.**

Explanation

The host map was not created because the SCSI logical unit number (LUN) ID is not valid.

User response

Specify the correct SCSI logical unit number (LUN) ID, and resubmit the command.

CMMVC5845E **The extent was not migrated because an object that was specified in the command does not exist.**

Explanation

The extent was not migrated because an object that was specified in the command does not exist.

User response

Specify a different object, and resubmit the command.

CMMVC5846E **The volume was not migrated because an object that was specified in the command does not exist.**

Explanation

The volume was not migrated because an object that was specified in the command does not exist.

User response

Specify a different object, and resubmit the command.

CMMVC5847E **The volume was not migrated because its associated managed disk (MDisk) is already in the storage pool.**

Explanation

The volume was not migrated because its associated managed disk (MDisk) is already in the storage pool.

User response

Not applicable.

CMMVC5848E **The action failed because the volume does not exist or it is being deleted.**

Explanation

The action failed because the volume does not exist or it is being deleted.

User response

Specify a different volume, and resubmit the command.

CMMVC5849E **The migration failed because some or all of the extents are already being migrated.**

Explanation

The migration failed because some or all of the extents are already being migrated.

User response

Not applicable.

CMMVC5850E **The extent was not migrated because there is a problem with the source extents.**

Explanation

The extent was not migrated because there is a problem with the source extents.

User response

Not applicable.

CMMVC5851E **The extent was not migrated because there is a problem with the target extents.**

Explanation

The extent was not migrated because there is a problem with the target extents.

User response

Not applicable.

CMMVC5852E **The migration failed because there are too many migrations in progress.**

Explanation

The migration failed because there are too many migrations in progress.

User response

Wait for the migration process to complete, and resubmit the command.

CMMVC5853E **The action failed because there was a problem with the group.**

Explanation

An attempt was made to work on a volume which is using a storage pool with one of the following problems:

- The target and source storage pools have different extent sizes (group migrate).
- The target and source storage pools are the same (group migrate).
- The target and source storage pools are different (extents migrate).
- The target group (group migrate) is not valid.
- The source group (group migrate) is not valid.

User response

Ensure that none of the above conditions exist before reissuing the command.

CMMVC5854E **The extent information was not returned because the extent is not used or does not exist.**

Explanation

The extent information was not returned because the extent is not used or does not exist.

User response

Specify the correct extent, and resubmit the command.

CMMVC5855E **The extent information was not returned because the managed disk (MDisk) is not used by any volume.**

Explanation

The extent information was not returned because the managed disk (MDisk) is not used by any volume.

User response

Specify the correct MDisk, and resubmit the command.

CMMVC5856E **The action failed because the volume does not belong to the specified storage pool.**

Explanation

The action failed because the volume does not belong to the specified storage pool.

User response

Specify a different volume, and resubmit the command.

CMMVC5857E **The action failed because the managed disk (MDisk) does not exist or it is not a member of the storage pool.**

Explanation

The action failed because the managed disk (MDisk) does not exist or it is not a member of the storage pool.

User response

Specify a different MDisk, and resubmit the command.

CMMVC5858E **The action failed because the volume is in the wrong mode, the managed disk (MDisk) is in the wrong mode, or both are in the wrong mode.**

Explanation

The action failed because the volume is in the wrong mode, the managed disk (MDisk) is in the wrong mode, or both are in the wrong mode.

User response

Check that the volume and MDisk are in the correct mode, and resubmit the command.

CMMVC5859E **The migration did not complete because an error occurred during the migration of the last extent on an image-mode volume.**

Explanation

The migration did not complete because an error occurred during the migration of the last extent on an image-mode volume.

User response

Not applicable.

CMMVC5860E **The action failed because there were not enough extents in the storage pool.**

Explanation

This error is also returned if a stripe set of MDisks has been specified and one or more of these MDisks does not contain enough free extents to complete the creation of the volume.

User response

In this case, the storage pool reports that it has enough free capacity to create the volume. You can check the free capacity on each MDisk by submitting the command lsfreeextents <mdiskname/ID> . Alternatively, do not specify a stripe set and let the system choose the free extents automatically.

CMMVC5861E **The action failed because there were not enough extents on the managed disk (MDisk).**

Explanation

The action failed because there were not enough extents on the managed disk (MDisk).

User response

Specify another extent, and resubmit the command.

CMMVC5862E **The action failed because the volume is being formatted.**

Explanation

The action failed because the volume is being formatted.

User response

Wait for the volume to be successfully formatted, and resubmit the command.

CMMVC5863E **The migration failed because there are not enough free extents on the target managed disk (MDisk).**

Explanation

The migration failed because there are not enough free extents on the target managed disk (MDisk).

User response

Specify another free extent, and resubmit the command.

CMMVC5864E **The extent information was not returned because the source extent is not used.**

Explanation

The extent information was not returned because the source extent is not used.

User response

Specify a different source extent, and resubmit the command.

CMMVC5865E **The action failed because the extent is out of range for the managed disk (MDisk) or volume specified.**

Explanation

The extent information was not returned because the extent is out of range for the managed disk (MDisk) or volume.

User response

Specify a different extent which is in range for the MDisk or volume and resubmit the command.

CMMVC5866E **The action failed because the extent contains internal data.**

Explanation

The extent was not migrated because the extent contains internal data.

User response

Not applicable.

CMMVC5867E The action failed because the worldwide port name is already assigned or is not valid.

Explanation

The action failed because the worldwide port name is already assigned or is not valid.

User response

Specify a different worldwide port name, and resubmit the command.

CMMVC5868E The action failed because an entity that was specified in the command does not exist.

Explanation

The action failed because an entity that was specified in the command does not exist.

User response

Specify a different entity, and resubmit the command.

CMMVC5869E The host object was not renamed because the host ID or name is not valid.

Explanation

The host object was not renamed because the host ID or name is not valid.

User response

Specify a different host ID or name, and resubmit the command.

CMMVC5870E The host object was not deleted because an entity that was specified in the command does not exist.

Explanation

The host object was not deleted because an entity that was specified in the command does not exist.

User response

Specify the correct entity, and resubmit the command.

CMMVC5871E The host object was not deleted because it is mapped to one or more volumes.

Explanation

There is at least one volume mapped to this host. Therefore the host object cannot be deleted without the loss of data.

User response

Verify that you are attempting to delete the correct host. If so, use the **lshostvdiskmap** command to locate the volumes that are mapped to this host.

If you are sure that these mappings are no longer required, remove them using the **rmvdiskhostmap** command, and retry the original **rmhost** command.

CMMVC5872E The port (WWPN) was not added to the host object because an object that was specified in the command does not exist.

Explanation

The port (WWPN) was not added to the host object because an object that was specified in the command does not exist.

User response

Specify the correct object, and resubmit the command.

CMMVC5873E No matching WWPN.

Explanation

The action failed because there is no matching worldwide port name.

User response

Not applicable.

CMMVC5874E The action failed because the host does not exist.

Explanation

The action failed because the host does not exist.

User response

Specify a different host, and resubmit the command.

CMMVC5875E **The action failed because the volume does not exist.**

Explanation

The action failed because the volume does not exist.

User response

Specify a different volume, and resubmit the command.

CMMVC5876E **The host mapping was not created because the maximum number of mappings has been reached.**

Explanation

The host map was not created because the maximum number of mappings has been reached.

User response

Not applicable.

CMMVC5877E **The host mapping was not created because the maximum number of SCSI LUNs has been allocated.**

Explanation

The host map was not created because the maximum number of SCSI LUNs has been allocated.

User response

Not applicable.

CMMVC5878E **The host mapping was not created because this volume is already mapped to this host.**

Explanation

The host map was not created because this volume is already mapped to this host.

User response

Specify a different volume, and resubmit the command.

CMMVC5879E **The host mapping was not created because a volume is already mapped to this host with this SCSI LUN.**

Explanation

The host map was not created because this SCSI LUN is already assigned to another mapping.

User response

Specify a different SCSI LUN, and resubmit the command.

CMMVC5880E **The volume was not created because a capacity of zero bytes is not allowed for image mode disks.**

Explanation

The host map was not created because the volume has a capacity of zero bytes.

User response

Specify a different volume, and resubmit the command.

CMMVC5881E **The FlashCopy mapping was not created because an entity that was specified in the command does not exist.**

Explanation

The FlashCopy mapping was not created because an entity that was specified in the command does not exist.

User response

Specify a different entity, and resubmit the command.

CMMVC5882E **The FlashCopy mapping was not created because a mapping for the source or target volume already exists.**

Explanation

The FlashCopy mapping was not created because a mapping for the source or target volume already exists.

User response

Specify a different source or target volume, and resubmit the command.

CMMVC5883E **The FlashCopy mapping was not created because the recovery I/O group is associated with the source or target volume.**

Explanation

The FlashCopy mapping was not created because the recovery I/O group is associated with the source or target volume.

User response

Specify a different recovery I/O group, and resubmit the command.

CMMVC5884E The FlashCopy mapping was not created because the source or target volume cannot be a member of a Remote Copy mapping.

Explanation

The FlashCopy mapping was not created because the source or target volume cannot be a member of a Remote Copy mapping.

User response

Specify a different source or target volume, and resubmit the command.

CMMVC5885E The FlashCopy mapping was not created because this source or target volume cannot be a member of a FlashCopy mapping.

Explanation

The FlashCopy mapping was not created because this source or target volume cannot be a member of a FlashCopy mapping.

User response

Specify a different source or target volume, and resubmit the command.

CMMVC5886E The FlashCopy mapping was not created because the source or target volume is associated with the recovery I/O group.

Explanation

The FlashCopy mapping was not created because the source or target volume is associated with the recovery I/O group.

User response

Specify a different source or target volume, and resubmit the command.

CMMVC5887E The FlashCopy mapping was not created because the source or target volume must not be in router mode.

Explanation

The FlashCopy mapping was not created because the source or target volume must not be in router mode.

User response

Specify a different source or target volume, and resubmit the command.

CMMVC5888E The action failed because an entity that was specified in the command does not exist.

Explanation

The action failed because an entity that was specified in the command does not exist.

User response

Specify the correct entity, and resubmit the command.

CMMVC5889E The FlashCopy mapping was not deleted because an entity that was specified in the command does not exist.

Explanation

The FlashCopy mapping was not deleted because an entity that was specified in the command does not exist.

User response

Specify a different entity, and resubmit the command.

CMMVC5890E The FlashCopy mapping or consistency group was not started because starting consistency group 0 is not a valid operation.

Explanation

The FlashCopy mapping or consistency group was not started because starting consistency group 0 is not a valid operation.

User response

Not applicable.

CMMVC5891E **The FlashCopy consistency group was not created because the name is not valid.**

Explanation

The FlashCopy consistency group was not created because the name is not valid.

User response

Specify a different name, and resubmit the command.

CMMVC5892E **The FlashCopy consistency group was not created because it already exists.**

Explanation

The FlashCopy consistency group was not created because it already exists.

User response

Not applicable.

CMMVC5893E **The action failed because an entity that was specified in the command does not exist.**

Explanation

The action failed because an entity that was specified in the command does not exist.

User response

Specify the correct entity, and resubmit the command.

CMMVC5894E **The FlashCopy consistency group was not deleted because you are trying to delete consistency group 0 or the name of the consistency group is not valid.**

Explanation

The FlashCopy consistency group was not deleted because the name of the consistency group is not valid or you are trying to delete consistency group 0.

User response

Specify the correct consistency group, and resubmit the command.

CMMVC5895E **The FlashCopy consistency group was not deleted because it contains mappings. To delete this consistency group, a forced deletion is required.**

Explanation

The FlashCopy consistency group was not deleted because it contains mappings.

User response

Specify that -force option to delete the consistency group.

CMMVC5896E **The FlashCopy mapping was not deleted because the mapping or consistency group is in the preparing state. The mapping or consistency group must be stopped first.**

Explanation

The FlashCopy mapping was not deleted because the mapping or consistency group is in the preparing state. The mapping or consistency group must be stopped first.

User response

Stop the consistency group, and resubmit the command.

CMMVC5897E **The FlashCopy mapping was not deleted because the mapping or consistency group is in the prepared state. The mapping or consistency group must be stopped first.**

Explanation

The FlashCopy mapping was not deleted because the mapping or consistency group is in the prepared state. The mapping or consistency group must be stopped first.

User response

Stop the consistency group, and resubmit the command.

CMMVC5898E **The FlashCopy mapping was not deleted because the mapping or consistency group is in the copying state. The mapping or consistency group must be stopped first.**

Explanation

The FlashCopy mapping was not deleted because the mapping or consistency group is in the copying state. The mapping or consistency group must be stopped first.

User response

Stop the consistency group, and resubmit the command.

CMMVC5899E **The FlashCopy mapping was not deleted because the mapping or consistency group is in the stopped state. To delete the mapping, a forced deletion is required.**

Explanation

The FlashCopy mapping was not deleted because the mapping or consistency group is in the stopped state.

User response

Specify the -force option to delete the mapping.

CMMVC5900E **The FlashCopy mapping was not deleted because the mapping or consistency group is in the suspended state. The mapping or consistency group must be stopped first.**

Explanation

The FlashCopy mapping was not deleted because the mapping or consistency group is in the suspended state. The mapping or consistency group must be stopped first.

User response

Stop the consistency group, and resubmit the command.

CMMVC5901E **The FlashCopy mapping was not prepared because the mapping or consistency group is already in the preparing state.**

Explanation

The FlashCopy mapping was not prepared because the mapping or consistency group is already in the preparing state.

User response

Not applicable.

CMMVC5902E **The FlashCopy mapping was not prepared because the mapping or consistency group is already in the prepared state.**

Explanation

The FlashCopy mapping was not prepared because the mapping or consistency group is already in the prepared state.

User response

Not applicable.

CMMVC5903E **The FlashCopy mapping was not prepared because the mapping or consistency group is already in the copying state.**

Explanation

The FlashCopy mapping was not prepared because the mapping or consistency group is already in the copying state.

User response

Not applicable.

CMMVC5904E **The FlashCopy mapping was not prepared because the mapping or consistency group is already in the suspended state.**

Explanation

The FlashCopy mapping was not prepared because the mapping or consistency group is already in the suspended state.

User response

Not applicable.

CMMVC5905E **The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the idle state. The mapping or consistency group must be prepared first.**

Explanation

The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the idle state.

User response

Prepare the mapping or consistency group, and resubmit the command.

CMMVC5906E **The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the preparing state.**

Explanation

The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the preparing state.

User response

Not applicable.

CMMVC5907E **The FlashCopy mapping or consistency group was not started because the mapping or consistency group is already in the copying state.**

Explanation

The FlashCopy mapping or consistency group was not started because the mapping or consistency group is already in the copying state.

User response

Not applicable.

CMMVC5908E **The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the stopped state. The mapping or consistency group must be prepared first.**

Explanation

The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the stopped state.

User response

Prepare the mapping or consistency group, and resubmit the command.

CMMVC5909E **The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the suspended state.**

Explanation

The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the suspended state.

User response

Not applicable.

CMMVC5910E **The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is in the idle state.**

Explanation

The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is in the idle state.

User response

Not applicable.

CMMVC5911E **The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is in the preparing state.**

Explanation

The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is in the preparing state.

User response

Not applicable.

CMMVC5912E **The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is already in the stopped state.**

Explanation

The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is already in the stopped state.

User response

Not applicable.

CMMVC5913E The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the preparing state.

Explanation

The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the preparing state.

User response

Not applicable.

CMMVC5914E The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the prepared state.

Explanation

The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the prepared state.

User response

Not applicable.

CMMVC5915E The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the copying state.

Explanation

The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the copying state.

User response

Not applicable.

CMMVC5916E The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the suspended state.

Explanation

The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the suspended state.

User response

Not applicable.

CMMVC5917E The FlashCopy mapping was not created because there is no memory in which to create the bitmap.

Explanation

The FlashCopy mapping was not created because there is no memory to create the bitmap.

User response

Not applicable.

CMMVC5918E The FlashCopy mapping was not prepared, either because there are no online nodes in the I/O group or because there are unrecovered FlashCopy mappings or unrecovered Global Mirror or Metro Mirror relationships in the I/O group.

Explanation

This error might be caused by a temporary loss of all of the nodes in the I/O group, which causes all of the FlashCopy mappings and Global and Metro Mirror relationships of the I/O group to be unusable.

User response

Perform the following steps:

1. Ensure that at least one of the nodes in the I/O group of the mapping is online.
2. Fix all of the unfixed events in the event log.
3. Follow the fix procedures.

You might be required to delete and re-add ALL of the FlashCopy maps and Global and Metro Mirror relationships in the I/O group.

Resubmit the command.

CMMVC5919E **The FlashCopy mapping or consistency group was not started, either because there are no online nodes in the I/O group or because there are unrecovered FlashCopy mappings or unrecovered Global Mirror or Metro Mirror relationships in the I/O group.**

Explanation

This error might be caused by a temporary loss of all of the nodes in the I/O group, which causes all of the FlashCopy mappings and Global and Metro Mirror relationships of the I/O group to be unusable.

User response

Perform the following steps:

1. Ensure that at least one of the nodes in the I/O group of the mapping is online.
2. Fix all of the unfixed events in the event log.
3. Follow the fix procedures.

You might be required to delete and re-add ALL of the FlashCopy maps and Global and Metro Mirror relationships in the I/O group.

Resubmit the command.

CMMVC5920E **The FlashCopy mapping was not created because the consistency group is not idle.**

Explanation

The FlashCopy mapping was not created because the consistency group is not idle.

User response

Not applicable.

CMMVC5921E **The properties of the FlashCopy mapping were not modified because the consistency group is not idle.**

Explanation

The properties of the FlashCopy mapping were not modified because the consistency group is not idle.

User response

Not applicable.

CMMVC5922E **The FlashCopy mapping was not created because the destination volume is too small.**

Explanation

The FlashCopy mapping was not created because the destination volume is too small.

User response

Specify a different volume, and resubmit the command.

CMMVC5923E **The FlashCopy mapping cannot be created, either because there are no online nodes in the I/O group or because there are unrecovered FlashCopy mappings in the I/O group.**

Explanation

This error might be caused by a temporary loss of all of the nodes in the I/O group, which causes all of the FlashCopy mappings to be unusable.

User response

Perform the following steps:

1. Ensure that at least one of the nodes in the I/O group of the mapping is online.
2. Fix all of the unfixed events in the event log.
3. Follow the fix procedures.

You might be required to delete and re-add ALL of the FlashCopy maps and Global and Metro Mirror relationships in the I/O group.

Resubmit the command.

CMMVC5924E **The FlashCopy mapping was not created because the source and target volumes are different sizes.**

Explanation

The FlashCopy mapping was not created because the source and target volumes are different sizes.

User response

Specify a different source and target volume that are the same size, and resubmit the command.

CMMVC5925E **The remote cluster partnership was not created because it already exists.**

Explanation

The remote cluster partnership was not created because it already exists.

User response

Specify a different remote cluster partnership, and resubmit the command.

CMMVC5926E The remote cluster partnership was not created because there are too many partnerships.

Explanation

The remote cluster partnership was not created because there are too many partnerships.

User response

Not applicable.

CMMVC5927E The action failed because the cluster ID is not valid.

Explanation

The action failed because the cluster ID is not valid.

User response

Specify the correct cluster ID, and resubmit the command.

CMMVC5928E The action failed because the cluster name is a duplicate of another cluster.

Explanation

The action failed because the cluster name is a duplicate of another cluster.

User response

Specify a different cluster name, and resubmit the command.

CMMVC5929E The Remote Copy partnership was not deleted because it has already been deleted.

Explanation

The Remote Copy partnership was not deleted because it has already been deleted.

User response

Not applicable.

CMMVC5930E The Remote Copy relationship was not created because an object that was specified in the command does not exist.

Explanation

The Remote Copy relationship was not created because an object that was specified in the command does not exist.

User response

Specify the correct object, and resubmit the command.

CMMVC5931E The Remote Copy relationship was not created because the master or auxiliary volume is locked.

Explanation

The Remote Copy relationship was not created because the master or auxiliary volume is locked.

User response

Unlock the master or auxiliary volume, and resubmit the command.

CMMVC5932E The Remote Copy relationship was not created because the master or auxiliary volume is a member of a FlashCopy mapping.

Explanation

The Remote Copy relationship was not created because the master or auxiliary volume is a member of a FlashCopy mapping, and the partner cluster is running a downlevel software version.

User response

Not applicable.

CMMVC5933E The Remote Copy relationship was not created because the master or auxiliary volume is in the recovery I/O group.

Explanation

The Remote Copy relationship was not created because the master or auxiliary volume is in the recovery I/O group.

User response

Not applicable.

CMMVC5934E **The Remote Copy relationship was not created because the master or auxiliary volume is in the router mode.**

Explanation

The Remote Copy relationship was not created because the master or auxiliary volume is in the router mode.

User response

Not applicable.

CMMVC5935E **The action failed because an object that was specified in the command does not exist.**

Explanation

The action failed because an object that was specified in the command does not exist.

User response

Specify the correct object, and resubmit the command.

CMMVC5936E **The action failed because an object that was specified in the command does not exist.**

Explanation

The action failed because an object that was specified in the command does not exist.

User response

Specify the correct object, and resubmit the command.

CMMVC5937E **The action failed because an object that was specified in the command does not exist.**

Explanation

The action failed because an object that was specified in the command does not exist.

User response

Specify the correct object, and resubmit the command.

CMMVC5938E **The Remote Copy consistency group was not deleted because the consistency group contains relationships. To delete the consistency group, the force option is required.**

Explanation

Remote Copy consistency group was not deleted because the consistency group contains relationships.

User response

Specify the -force option to delete the consistency group.

CMMVC5939E **The action failed because the cluster is not in a stable state.**

Explanation

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

User response

Not applicable.

CMMVC5940E **The cluster that contains the auxiliary volume is unknown.**

Explanation

The cluster that contains the auxiliary volume is unknown.

User response

Not applicable.

CMMVC5941E **The cluster that contains the master volume has too many consistency groups.**

Explanation

The cluster that contains the master volume has too many consistency groups.

User response

Not applicable.

CMMVC5942E **The cluster that contains the auxiliary volume has too many consistency groups.**

Explanation

The cluster that contains the auxiliary volume has too many consistency groups.

User response

Not applicable.

CMMVC5943E The specified relationship is not valid.

Explanation

The specified relationship is not valid.

User response

Specify the correct relationship, and resubmit the command.

CMMVC5944E The specified consistency group is not valid.

Explanation

The specified consistency group is not valid.

User response

Specify the correct consistency group, and resubmit the command.

CMMVC5945E The specified master cluster is not valid.

Explanation

The specified master cluster is not valid.

User response

Specify the correct master cluster, and resubmit the command.

CMMVC5946E The specified auxiliary cluster is not valid.

Explanation

The specified auxiliary cluster is not valid.

User response

Specify the correct auxiliary cluster, and resubmit the command.

CMMVC5947E The specified master volume is not valid.

Explanation

The specified master volume is not valid.

User response

Specify the correct master volume, and resubmit the command.

CMMVC5948E The specified auxiliary volume is not valid.

Explanation

The specified auxiliary volume is not valid.

User response

Specify the auxiliary volume, and resubmit the command.

CMMVC5949E The specified relationship is unknown.

Explanation

The specified relationship is unknown.

User response

Specify a different relationship, and resubmit the command.

CMMVC5950E The specified consistency group is unknown.

Explanation

The specified consistency group is unknown.

User response

Specify a different consistency group, and resubmit the command.

CMMVC5951E The operation cannot be performed because the relationship is not a stand-alone relationship.

Explanation

The operation cannot be performed because the relationship is not a stand-alone one.

User response

Not applicable.

CMMVC5952E The relationship and consistency group have different master clusters.

Explanation

The relationship and consistency group have different master clusters.

User response

Not applicable.

CMMVC5953E The relationship and group have different auxiliary clusters.

Explanation

The relationship and group have different auxiliary clusters.

User response

Not applicable.

CMMVC5954E The master and auxiliary volumes are different sizes.

Explanation

The master and auxiliary volumes are different sizes

User response

Not applicable.

CMMVC5955E The maximum number of relationships has been reached.

Explanation

The maximum number of relationships has been reached.

User response

Not applicable.

CMMVC5956E The maximum number of consistency groups has been reached.

Explanation

The maximum number of consistency groups has been reached.

User response

Not applicable.

CMMVC5957E The master volume is already in a relationship.

Explanation

The master volume is already in a relationship.

User response

Specify a different master volume, and resubmit the command.

CMMVC5958E The auxiliary volume is already in a relationship.

Explanation

The auxiliary volume is already in a relationship.

User response

Specify a different auxiliary volume, and resubmit the command.

CMMVC5959E There is a relationship that already has this name on the master cluster.

Explanation

There is a relationship that already has this name on the master cluster.

User response

Specify a different name, and resubmit the command.

CMMVC5960E There is a relationship that already has this name on the auxiliary cluster.

Explanation

There is a relationship that already has this name on the auxiliary cluster.

User response

Specify a different name, and resubmit the command.

CMMVC5961E There is a consistency group that already has this name on the master cluster.

Explanation

There is a consistency group that already has this name on the master cluster.

User response

Specify a different name, and resubmit the command.

CMMVC5962E There is a consistency group that already has this name on the auxiliary cluster.

Explanation

There is a consistency group that already has this name on the auxiliary cluster.

User response

Specify a different name, and resubmit the command.

CMMVC5963E No direction has been defined.

Explanation

No direction has been defined.

User response

Not applicable.

CMMVC5964E The copy priority is not valid.

Explanation

The copy priority is not valid.

User response

Not applicable.

CMMVC5965E The volumes are in different I/O groups on the local cluster.

Explanation

The volumes are in different I/O groups on the local cluster.

User response

Not applicable.

CMMVC5966E The master volume is unknown.

Explanation

The master volume is unknown.

User response

Specify a different master volume, and resubmit the command.

CMMVC5967E The auxiliary volume is unknown.

Explanation

The auxiliary volume is unknown.

User response

Specify a different auxiliary volume, and resubmit the command.

CMMVC5968E The relationship cannot be added because the states of the relationship and the consistency group do not match.

Explanation

The relationship cannot be added because the states of the relationship and the consistency group do not match.

User response

Not applicable.

CMMVC5969E The Remote Copy relationship was not created, either because there are no online nodes in the I/O group or because there are unrecovered FlashCopy mappings or unrecovered Global Mirror or Metro Mirror relationships in the I/O group.

Explanation

This error might be caused by a temporary loss of all of the nodes in the I/O group, which causes all of the FlashCopy mappings and Global and Metro Mirror relationships of the I/O group to be unusable.

User response

Perform the following steps:

1. Ensure that at least one of the nodes in the I/O group is online.
2. Fix all of the unfixed events in the event log.
3. Follow the fix procedures.

You might be required to delete and re-add ALL of the FlashCopy maps and Global and Metro Mirror relationships in the I/O group.

Resubmit the command.

CMMVC5970E The Remote Copy relationship was not created because there is not enough memory.

Explanation

The Remote Copy relationship was not created because there is not enough memory.

User response

Increase the memory with the chiogrp command.

CMMVC5971E The operation was not performed because the consistency group contains no relationships.

Explanation

The operation was not performed because the consistency group contains no relationships.

User response

Not applicable.

CMMVC5972E The operation was not performed because the consistency group contains relationships.

Explanation

The operation was not performed because the consistency group contains relationships.

User response

Not applicable.

CMMVC5973E The operation was not performed because the consistency group is not synchronized.

Explanation

The operation was not performed because the consistency group is not synchronized.

User response

Specify the Force option when starting the consistency group.

CMMVC5974E The operation did not complete because one or more volumes in the consistency group are offline or inaccessible, or the Remote Copy feature is not licensed.

Explanation

This error can occur for one of the following reasons:

- One or more of the source or target volumes for relationships within the consistency group is offline.
- One or more of the source or target volumes are inaccessible because it is participating in a FlashCopy mapping that is prepared or incomplete.
- One of the systems does not have a license for the remote copy function.

User response

Complete the following steps:

1. Ensure that at least one of the nodes in the I/O group of each of the source and target volumes is online.
2. Ensure that the source and target volumes for each relationship in the consistency group are online.
3. Correct any problems that might be holding the source or target volumes offline. For example, a thin-provisioned volume might be held offline if all available space is allocated.
4. Ensure that any FlashCopy mappings that involve the source or target volumes for each relationship are complete.
5. Ensure that the remote copy function is licensed on all systems.
6. Fix all of the unfixed events that are in the event log.

CMMVC5975E The operation was not performed because the cluster partnership is not connected.

Explanation

The operation was not performed because the cluster partnership is not connected.

User response

Not applicable.

CMMVC5976E The operation was not performed because the consistency group is in the freezing state.

Explanation

The operation was not performed because the consistency group is in the freezing state.

User response

Not applicable.

CMMVC5977E **The operation was not performed because it is not valid given the current consistency group state.**

Explanation

The operation was not performed because it is not valid given the current consistency group state.

User response

Not applicable.

CMMVC5978E **The operation was not performed because the relationship is consistent but is not synchronized. Restarting the relationship by using the -force parameter will make the relationship inconsistent until the background copy has completed.**

Explanation

Input transactions have occurred on either the primary or secondary volumes since the ConsistentStopped or Idling state has occurred. Because the relationship is no longer synchronized, the state of the relationship is now Stopped.

The -force parameter of the startrcrelationship command is required when the relationship is not synchronized because consistency would be lost by starting the copy operation. Submitting the startrcrelationship command on an unsynchronized relationship without using the -force parameter is not supported.

If a relationship is in the InconsistentStopped, InconsistentCopying or ConsistentSynchronized state, the -force parameter is not required, but is supported.

User response

Consider using the -force parameter of the startrcrelationship command, if appropriate.

CMMVC5980E **The operation was not performed because the master and auxiliary clusters are not connected.**

Explanation

The operation was not performed because the master and auxiliary clusters are not connected.

User response

Not applicable.

CMMVC5981E **The operation was not performed because the relationship is in the freezing state.**

Explanation

The operation was not performed because the relationship is in the freezing state.

User response

Not applicable.

CMMVC5982E **The operation was not performed because it is not valid given the current relationship state.**

Explanation

The operation was not performed because it is not valid given the current relationship state.

User response

Not applicable.

CMMVC5983E **dump file was not created. This might be due to the file system being full.**

Explanation

dump file was not created. This might be due to the file system being full.

User response

Not applicable.

CMMVC5984E **The dump file was not written to disk. This might be due to the file system being full.**

Explanation

The dump file was not written to disk. This might be due to the file system being full.

User response

Not applicable.

CMMVC5985E **The action failed because the specified directory is not permitted for this command.**

Explanation

You have attempted to copy, delete, or list dumps from a directory that is not valid. A list of valid directories for these commands is provided in the documentation.

User response

Ensure that the directory you specify is valid, and resubmit the command.

CMMVC5986E **The tracing of I/O operations was not started because the volume or managed disk (MDisk) failed to return any statistics.**

Explanation

The tracing of I/O operations was not started because the volume or managed disk (MDisk) failed to return statistics.

User response

Not applicable.

CMMVC5987E ***VALUE*** **is not a valid command line option.**

Explanation

The specified string that you have entered is not a supported command line option.

User response

Specify a supported option, and resubmit the command.

CMMVC5988E **command should not be run by the root userid. Use the admin userid.**

Explanation

This command should not be issued if you are logged in with a root user ID. Use the admin userid.

User response

Log off of the root user ID and log in as admin.

CMMVC5989E **The operation was not performed because the relationship is offline or the Remote Copy feature is not licensed.**

Explanation

This error can occur when a relationship is offline or when the remote copy function is not licensed on one of the systems.

User response

Complete the following steps:

1. Ensure that the volumes in the relationship are online.
2. Ensure that the remote copy function is licensed on all systems.

CMMVC5990E **The FlashCopy consistency group was not stopped as there are no FlashCopy mappings within the group.**

Explanation

The FlashCopy consistency group was not stopped as there are no FlashCopy mappings within the group.

User response

Not applicable.

CMMVC5991E **The Remote Copy consistency group was not stopped as there are no Remote Copy relationships within the group.**

Explanation

The Remote Copy consistency group was not stopped as there are no Remote Copy relationships within the group.

User response

Not applicable.

CMMVC5992E **The Remote Copy consistency group was not stopped as there are no Remote Copy relationships within the group.**

Explanation

The Remote Copy consistency group was not stopped as there are no Remote Copy relationships within the group.

User response

Not applicable.

CMMVC5993E The specified update package does not exist.

Explanation

The specified update package does not exist.

User response

Not applicable.

CMMVC5994E Error in verifying the signature of the update package.

Explanation

The system could not verify the signature of the update package due to the following reasons:

- There is not enough space on the system to copy the file.
- The package is incomplete or contains errors.

User response

If the copy failed with an error indicating that there was insufficient space on the system, free up additional space on your system. Otherwise, ensure that the cluster time and date stamp on the signature is correct. (For example, the time and date cannot be in the future.)

CMMVC5995E An error prevented the unpacking of the update package.

Explanation

The system disk is too full to allow the update package to be unpacked.

User response

Use the cleardumps command with the parameter -prefix /home/admin/upgrade/ to clear unused files, then reboot the node before attempting to unpack the update package again.

CMMVC5996E The specific update package cannot be installed over the current version.

Explanation

The update package is not compatible with the current version or the system.

User response

Check the available update packages and find the correct update package for your current version and for your system. If the update package is correct for your system, check the version requirements for the package. You might have to update the current version to an intermediate version before you update to the latest version. (For example, if your current version is 1 and you are trying to update to version 3, you might need to update to version 2 before applying the version 3 update.)

CMMVC5999W Featurization for this facility has not been enabled.

Explanation

Featurization for this facility has not been enabled.

User response

Not applicable.

CMMVC6000W Featurization for this facility has not been enabled.

Explanation

Featurization for this facility has not been enabled.

User response

Not applicable.

CMMVC6001E The FlashCopy consistency group was not started as there are no FlashCopy mappings within the group.

Explanation

The FlashCopy consistency group was not started as there are no FlashCopy mappings within the group.

User response

Create a FlashCopy within the appropriate group.

CMMVC6002E This command can only be run on a node that is in the service state.

Explanation

This command can only be run on a node that is in the service state.

User response

Not applicable.

CMMVC6003E **This command cannot be run on a node that is in the service state.**

Explanation

This command cannot be run on a node that is in the service state.

User response

Not applicable.

CMMVC6004E **The delimiter value *VALUE* is invalid.**

Explanation

The specified value is not a valid delimiter value.

User response

Specify a different delimiter.

CMMVC6005E **The view request failed as the specified object is not a member of an appropriate group.**

Explanation

A view was request on an object that has been incorrectly initialized.

User response

Ensure that the object is correctly initialized before resubmitting the view request.

CMMVC6006E **The managed disk (MDisk) was not deleted because the resource was busy.**

Explanation

An attempt was made to delete an MDisk from a storage pool that is being used as a source and destination for migration operations.

User response

Ensure that the storage pool is not being used for migration operations before reissuing the command.

CMMVC6007E **The two passwords that were entered do not match.**

Explanation

The two passwords entered for verification of your password change were not the same.

User response

Re-enter the passwords.

CMMVC6008E **The key already exists.**

Explanation

An attempt was made to load a duplicate SSH key.

User response

Not applicable.

CMMVC6009E **Unable to allocate a block of memory in which to copy the returned data.**

Explanation

The command line was unable to allocate a block of memory in which to copy the results of the query.

User response

Resubmit the command. If the problem persists, contact IBM technical support for assistance.

CMMVC6010E **Unable to complete the command as there are insufficient free extents, or the command requested an expansion of 0 size.**

Explanation

There are not enough free extents to meet the request.

User response

Not applicable.

CMMVC6011E **This cluster is part of a remote cluster partnership. Because this update package will make changes to the cluster state, it cannot be applied to the current code level until all remote cluster partnerships are deleted.**

Explanation

You have attempted to apply software when a Remote Copy relationship to a remote cluster exists.

User response

Delete the Remote Copy relationship to the remote clusters, and resubmit the command.

CMMVC6012W The virtualized storage capacity is approaching the amount that you are licensed to use.

Explanation

The requested action has completed. However, the limits permitted by the license you purchased are approaching.

User response

Subsequent actions might require that you increase your licensed limits.

CMMVC6013E The command failed because there is a consistency group mismatch on the aux cluster.

Explanation

The action has failed as there was a difference in attributes between the Metro Mirror consistency groups involved.

User response

Ensure that the attributes of the two Metro Mirror consistency groups match before resubmitting the command.

CMMVC6014E The command failed because the requested object is either unavailable or does not exist.

Explanation

The command failed because the requested object is either unavailable or does not exist.

User response

Ensure that all parameters have been correctly entered. If this is the case the determine why the object is unavailable, then resubmit the command.

CMMVC6015E A delete request is already in progress for this object.

Explanation

A delete request is already in progress for this object.

User response

Not applicable.

CMMVC6016E The action failed as there would be, or are, no more disks in the storage pool.

Explanation

The action failed as there would be, or are, no more disks in the I/O group.

User response

Ensure that all parameters have been correctly entered.

CMMVC6017E A parameter or argument contains invalid characters. Ensure that all characters are ASCII.

Explanation

The command-line interface (CLI) will only accept ASCII input.

User response

Ensure that all input to the CLI is ASCII, then resubmit the command.

CMMVC6018E The update pre-install process failed.

Explanation

The update failed as there was an error during the preprocessing. The package is either not valid or corrupted.

User response

Ensure the package is a valid update package. Download the package from the source location again as it might have been corrupted during a network transfer.

CMMVC6019E The update failed because a node pended.

Explanation

The update failed because a node pended as the update was in progress.

User response

Ensure that all nodes are online and available before restarting the update process.

CMMVC6020E **The update failed because the system was unable to distribute the package to all of the nodes.**

Explanation

The system could not complete the process of updating files. A full disk is a possible cause.

User response

Ensure that all nodes are online, and use the cleardumps command to clean the updates directory.

CMMVC6021E **The system is currently busy performing another request. Try again later.**

Explanation

The requested action failed as the system is processing another request.

User response

Wait before resubmitting the request.

CMMVC6022E **The system is currently busy performing another request. Try again later.**

Explanation

The requested action failed as the system is processing another request.

User response

Wait before resubmitting the request.

CMMVC6023E **The system is currently busy performing another request. Try again later.**

Explanation

The requested action failed as the system is processing another request.

User response

Wait before resubmitting the request.

CMMVC6024E **The auxiliary volume entered is invalid.**

Explanation

The auxiliary volume is entered as a parameter in the command-line interface is not a valid auxiliary volume.

User response

Select a valid auxiliary volume, and resubmit the command.

CMMVC6025E **The RC consistency group Master cluster is not the local cluster.**

Explanation

The auxiliary volume is entered as a parameter in the command-line interface is not a valid auxiliary volume.

User response

Resubmit the command with a consistency group that belongs to the local cluster.

CMMVC6026E **The RC consistency group is not in the stopped state.**

Explanation

The action failed as the Metro Mirror consistency group is not in the stopped state.

User response

Ensure that the Metro Mirror consistency group is in the stopped state before resubmitting the command.

CMMVC6027E **The RC consistency group is not the primary master.**

Explanation

The RC consistency group requested in the command is not the Metro Mirror primary master.

User response

Ensure that the parameters have been entered correctly on the command line.

CMMVC6028E **This package cannot be applied to the current code level because it contains changes to the cluster state and there are remote cluster partnerships defined.**

Explanation

The action failed because there is a connected remote cluster. The update cannot be applied because it would put the remote cluster at a different code level than the local cluster.

User response

Ensure that the cluster partnership is unconfigured before resubmitting the command. Ensure that you unconfigure the remote cluster and update the code on it before reconfiguring the cluster partnership.

CMMVC6029E **All nodes must have identical code level before a concurrent code update can be performed.**

Explanation

The concurrent update failed as two or more nodes were at differing code levels. All nodes must be at the same code level before a software update can be performed.

User response

Use the service assistant to bring all nodes to the same level before resubmitting the concurrent update.

CMMVC6030E **The operation was not performed because the FlashCopy mapping is part of a consistency group. The action must be performed at the consistency group level.**

Explanation

An attempt was made to stop a FlashCopy mapping. This failed as the FlashCopy mapping is part of a consistency group.

User response

Issue the stop command to the FlashCopy consistency group. This will stop all FlashCopies within that group that are in progress.

CMMVC6031E **The operation was not performed because the FlashCopy consistency group is empty.**

Explanation

An attempt was made to prestart an empty FlashCopy consistency group.

User response

Not applicable.

CMMVC6032E **The operation was not performed because one or more of the entered parameters is invalid for this operation.**

Explanation

An parameter that is not valid was entered for the command.

User response

If attempting to change the I/O group to which the volume belongs, ensure that the volume is not already a part of the group.

CMMVC6033E **The action failed due to an internal error.**

Explanation

An internal error caused the action to fail.

User response

Not applicable.

CMMVC6034E **The action failed because the maximum number of objects has been reached.**

Explanation

The action failed because the maximum number of objects has been reached.

User response

Not applicable.

CMMVC6035E **The action failed as the object already exists.**

Explanation

An operation was requested to create an object that already exists.

User response

Ensure that the name you are attempting to apply to a new object does not exist, or change the name before re-issuing the command.

CMMVC6036E **An invalid action was requested.**

Explanation

The action failed because it is not a valid action with the command that was issued.

User response

Issue an action that is valid with the command.

CMMVC6037E The action failed as the object is not empty.

Explanation

The action failed because an object was specified.

User response

Resubmit the command without specifying an object.

CMMVC6038E The action failed as the object is empty.

Explanation

The action failed because an object was not specified.

User response

Specify an object, and resubmit the command.

CMMVC6039E The action failed as the object is not a member of a group.

Explanation

The action failed because the object is not a member of a group.

User response

Specify an object that is part of a group, and resubmit the command.

CMMVC6040E The action failed as the object is not a parent.

Explanation

The action failed because the object is not a parent object.

User response

Specify an object that is a parent, and resubmit the command.

CMMVC6041E The action failed as the cluster is full.

Explanation

The action failed because the cluster is full.

User response

Remove data from the cluster, and resubmit the command.

CMMVC6042E The action failed as the object is not a cluster member.

Explanation

The action failed because the object is not a member of the cluster.

User response

Specify an object that is a member of the cluster, and resubmit the command.

CMMVC6043E The action failed as the object is a member of a group.

Explanation

The action failed because the object is a member of a group.

User response

Specify an object that is not a member of a group, and resubmit the command.

CMMVC6044E The action failed as the object is a parent.

Explanation

The action failed because the object is a parent object.

User response

Specify an object that is not a parent object, and resubmit the command.

CMMVC6045E The action failed, as the -force flag was not entered.

Explanation

The action failed because the -force option was not entered.

User response

Specify the -force option in the command.

CMMVC6046E The action failed as too many candidates were selected.

Explanation

The action failed because too many candidates were specified.

User response

Specify fewer candidates in the command.

CMMVC6047E The action failed as too few candidates were selected.

Explanation

An action was requested with too few candidate objects.

User response

Determine the correct number of candidates required for the specific command and reissue the command.

CMMVC6048E The action failed as the object is busy.

Explanation

The action failed because the object is busy.

User response

Not applicable.

CMMVC6049E The action failed as the object is not ready.

Explanation

The action failed because the object is not ready.

User response

Not applicable.

CMMVC6050E The action failed as the command was busy.

Explanation

The action failed because the command is busy.

User response

Not applicable.

CMMVC6051E An unsupported action was selected.

Explanation:
The hardware platform does not support snapshot objects.

User response:
None.

CMMVC6052E The action failed as the object is a member of a FlashCopy mapping.

Explanation

The object is a member of a FlashCopy mapping, thus it cannot be deleted.

User response

Specify an object that is not a member of a FlashCopy mapping, or remove the object from the FlashCopy mapping.

CMMVC6053E An invalid WWPN was entered.

Explanation

A worldwide port name (WWPN) that is not valid was specified.

User response

Specify a valid WWPN.

CMMVC6054E The action failed as not all nodes are online.

Explanation

One or more nodes are not online.

User response

Check that each node is online, and resubmit the command.

CMMVC6055E The action failed as an update is in progress.

Explanation

The action failed because a software update is in progress.

User response

Wait for the software update to complete, and resubmit the command.

CMMVC6056E The action failed as the object is too small.

Explanation

The action failed because the object is too small.

User response

Specify a different object, and resubmit the command.

CMMVC6058E The action failed as the object is in the recovery HWS.

Explanation

An attempt was made to perform an operation on a node that is in the recovery I/O group.

User response

Get the node into one of the other I/O groups and reissue the command.

CMMVC6059E The action failed as the object is in an invalid mode.

Explanation

The action failed because the object is in the wrong mode.

User response

Check that the object is in the correct mode, and resubmit the command.

CMMVC6060E The action failed as the object is being deleted.

Explanation

The action failed because the object is being deleted.

User response

Not applicable.

CMMVC6061E The action failed as the object is being resized.

Explanation

The action failed because the object is being resized.

User response

Check that the object is in the correct mode, and resubmit the command.

CMMVC6062E The action failed as the object is being moved between HWS.

Explanation

An attempt was made to perform an action against an object that is currently being moved between I/O groups.

User response

Re-issue the command when the move operation has completed.

CMMVC6063E The action failed as there are no more disks in the group.

Explanation

An attempt was made to perform an action against a group that contained no disks.

User response

Either add disks to the group and reissue the command, or select another group against which to execute the action.

CMMVC6064E The action failed as the object has an invalid name.

Explanation

An attempt was made to create or rename an object using a name that is not valid.

User response

Use a name that meets the naming standards and reissue the command.

CMMVC6065E The action failed as the object is not in a group.

Explanation

An attempt was made to perform an action on an object that was not in an appropriate group.

User response

Ensure that the object is a member of an appropriate group and reissue the command.

CMMVC6066E The action failed as the system is running low on memory.

Explanation

The system is running low on memory.

User response

Not applicable.

CMMVC6067E The action failed as the SSH key was not found.

Explanation

An attempt was made to perform an action using an SSH key that does not exist.

User response

Reissue the command using a key that does exist.

CMMVC6068E **The action failed as there are no free SSH keys.**

Explanation

An attempt was made to use an SSH key when there are no free SSH keys.

User response

Upload additional keys and reissue the command.

CMMVC6069E **The action failed as the SSH key is already registered.**

Explanation

An attempt was made to register an SSH key that was already registered.

User response

Not applicable.

CMMVC6070E **An invalid or duplicated parameter, unaccompanied argument, or incorrect argument sequence has been detected. Ensure that the input is as per the help.**

Explanation

The parameters entered for a command were not valid.

User response

Correct the parameters and reissue the command.

CMMVC6071E **The host mapping was not created because the volume is already mapped to a host.**

Explanation

The volume is already mapped to a host.

User response

Not applicable.

CMMVC6073E **The maximum number of files has been exceeded.**

Explanation

The maximum number of files has been exceeded.

User response

Not applicable.

CMMVC6074E **The command failed as the extent has already been assigned.**

Explanation

The command failed as the extent has already been assigned.

User response

Assign a different extent, and resubmit the command.

CMMVC6075E **The expand failed as the last extent is not a complete extent.**

Explanation

The expand failed as the last extent is not a complete extent.

User response

Assign a different extent, and resubmit the command.

CMMVC6076E **The command failed because the volume cache is not empty. Either wait for the cache to flush or use the force flag to discard the contents of the cache.**

Explanation

The command failed due to an error during the flushing of the volume.

User response

Not applicable.

CMMVC6077E **WARNING - Unfixed errors should be fixed before applying an update. Depending on the nature of the errors, they might cause the update process to fail. It is highly recommended to fix these errors before proceeding. If a particular error cannot be fixed, contact the support center.**

Explanation

Unfixed errors should be fixed before applying an update. Depending on the nature of the errors, they might cause the update process to fail. It is highly recommended to fix these errors before proceeding.

User response

If the error cannot be fixed, contact the support center.

CMMVC6078E **The action failed because the object is in an invalid mode.**

Explanation

An attempt was made to perform an action against an object in a mode that did not allow for that action to be performed.

User response

Get the object into a suitable mode and reissue the command.

CMMVC6079E **Metadata recovery could not complete the operation because a parameter is invalid.**

Explanation

Metadata recovery could not complete the operation because a parameter is not valid.

User response:

CMMVC6081E **Metadata Recovery is busy processing the previous operation.**

Explanation

Metadata Recovery is busy processing the previous operation.

User response:

CMMVC6082E **The attempt to abort metadata recovery failed because the previous operation has completed.**

Explanation

The attempt to cancel metadata recovery failed because the previous operation has completed.

User response

None.

CMMVC6083E **Metadata recovery could not find a valid dumpfile required for the rebuild operation.**

Explanation

Metadata recovery could not find a valid dumpfile required for the rebuild operation.

CMMVC6084E **Metadata recovery could not create/open/write the scan file, the disk might be full.**

Explanation

Metadata recovery could not create/open/write the scan file, the disk might be full.

User response:

CMMVC6085E **Metadata recovery could not create/open/write the dump file, the disk might be full.**

Explanation

Metadata recovery could not create/open/write the dump file, the disk might be full.

User response:

CMMVC6086E **Metadata recovery could not create/open/write the progress file, the disk might be full.**

Explanation

Metadata recovery could not create/open/write the progress file, the disk might be full.

User response:

CMMVC6087E **Metadata recovery could not map the buffers necessary to complete the operation.**

Explanation

Metadata recovery could not map the buffers necessary to complete the operation.

User response:

CMMVC6088E **The lba at which metadata recovery was requested does not contain metadata.**

Explanation

The lba at which metadata recovery was requested does not contain metadata.

User response:

CMMVC6089E **The metadata at the requested lba is flagged as invalid.**

Explanation

The metadata at the requested lba is flagged as not valid.

User response:

CMMVC6090E **The metadata header checksum verification failed.**

Explanation

The metadata header checksum verification failed.

User response:

CMMVC6091E **The metadata region checksum verification failed.**

Explanation

The metadata region checksum verification failed.

User response

Contact your administrator. Resubmit the command after the administrator confirms that the problem is resolved.

CMMVC6092E **The metadata recovery operation was aborted.**

Explanation

The metadata recovery operation was cancelled.

User response

Check your command and ensure that your input is correct. Resubmit the command. If you get the same error, contact yoour administrator.

CMMVC6093E **Metadata recovery internal error - (read only)**

Explanation

Metadata recovery internal error - (read only)

User response

Contact your administrator. When the administrator lets you know that the problem has been resolved, resubmit the command.

CMMVC6095E **Metadata recovery encountered the end of the disk.**

Explanation

Metadata recovery encountered the end of the disk.

User response

Contact the administrator and let them know about this error. The administrator must take care of this problem before you can continue.

CMMVC6096E **The metadata recovery task could not be initiated because the required back-end resource could not be found.**

Explanation

The back-end resource that is required for the task is unavailable.

User response:
Ensure that the required back-end resource is available, and reinitiate the task.

CMMVC6097E **The metadata recovery task could not be initiated because the system was unable to send the required I/O to the back-end resource.**

Explanation

The back-end resource is possibly not configured properly.

User response:
Ensure that the required back-end resource is accessible, and reinitiate the task.

CMMVC6098E **The copy failed as the specified node is the configuration node.**

Explanation

The copy failed because the specified node is the configuration node.

User response

Check your command. Correct the specified node and resubmit..

CMMVC6100E ***OPTION*** **not consistent with** ***ACTION***

Explanation

The specified option is not supported for the specified action.

User response

Remove the option, and resubmit the command.

CMMVC6101E ***OPTION*** **not consistent with** ***OPTION***

Explanation

The two specified options cannot be used together.

User response

Remove one of the options, and resubmit the command.

CMMVC6102E ***OPTION*** **and** ***OPTION*** **are alternatives**

Explanation

The two specified options are alternatives, and cannot be used together.

User response

Remove one of the options, and resubmit the command.

CMMVC6103E **Problem with** ***FILENAME*** **:** ***DETAILS***

Explanation

A problem occurred when opening the specified file. Determine the cause of the problem and correct it before trying again.

User response

Correct the problem, and resubmit the command.

CMMVC6104E **Action** ***ACTION*** **not run**

Explanation

An unexpected error occurred.

User response

Contact IBM technical support for assistance.

CMMVC6105E **Different names for source** ***SOURCE_CLUSTER_NAME*** **and target** ***TARGET_CLUSTER_NAME*** **clusters**

Explanation

The backup configuration cannot be restored to the target cluster because the source and target cluster have different names.

User response

Perform one of the following actions: (1) Use a different backup configuration. (2) Delete the cluster and recreate it with the same name as that stored in the backup configuration file.

CMMVC6106W **Target cluster has non-default id_alias** ***ALIAS*** **.**

Explanation

The specified id_alias of the target cluster is a non-default value. Clusters should have the default value. The non-default value suggests that the cluster is customized and is not suitable for restoration. Restoration changes the id_alias.

User response

Change the id_alias to a default value, and resubmit the command.

CMMVC6107E ***NUMBER_OF_OBJECTS*** **io_grp objects in target cluster;** ***NUMBER_OF_REQUIRED_OBJECTS*** **are required**

Explanation

The number of I/O groups in the target cluster is not sufficient to accommodate the I/O groups defined in the backup configuration file. Determine why there are not enough I/O groups.

User response

Correct the problem, and resubmit the command.

CMMVC6108I **Disk controller system with a WWNN of** ***WWNN_VALUE*** **found.**

Explanation

A disk controller system with the required WWNN has been found.

User response

Not applicable.

CMMVC6109E **Backup file version** ***version_id*** **is not compatible with current version** ***version_id***

Explanation

The backup file was generated on a cluster with a different version number than that of your current system.

User response

Make sure that you are copying the most up-to-date backup file onto the cluster. If you believe your backup file is correct, contact your service support representative.

CMMVC6110E Bad code level: *VALUE* .

Explanation

An unexpected error occurred.

User response

Contact IBM technical support for assistance.

CMMVC6111E The cluster code level could not be determined from *VALUE* .

Explanation

The code level of the cluster could not be determined. The code level should be of the format x.y.z, where x, y, and z are integers.

User response

If the cause of the problem cannot be determined, contact IBM technical support for assistance.

CMMVC6112W *OBJECT_TYPE OBJECT_NAME* has a default name.

Explanation

An object in the cluster has a default name. This can cause problems when restoring a cluster because default names are changed during restoration. Object IDs are also changed during restoration.

User response

Choose an appropriate name for each object in the cluster, and resubmit the command.

CMMVC6113E The command *COMMAND* has failed with return code *RETURN_CODE* .

Explanation

An attempt to run a command remotely failed using secure communications.

User response

Determine the cause of the problem, and resubmit the command. Specific steps are dependent on what command was running and what return code was given.

CMMVC6114E No help for action *ACTION* .

Explanation

There is no help for the specified action topic.

User response

Not applicable.

CMMVC6115W Feature *FEATURE_PROPERTY* mismatch: *VALUE* expected; *VALUE* found.

Explanation

The features in the backup configuration file and the target cluster do not match. There should be an exact match between the two. Nevertheless, the restore of the configuration can continue.

User response

Not applicable.

CMMVC6116I Feature match for *FEATURE* .

Explanation

The features in the backup configuration file and the target cluster are an exact match.

User response

Not applicable.

CMMVC6117E *FIX_OR_FEATURE* is not available.

Explanation

An unexpected error has occurred.

User response

Contact IBM technical support for assistance.

CMMVC6118I *TYPE* with *PROPERTY* *PROPERTY_VALUE* and *PROPERTY* *PROPERTY_VALUE* found.

Explanation

An object in the cluster has been found with the correct properties.

User response

Not applicable.

CMMVC6119E ***TYPE* with *PROPERTY* *PROPERTY_VALUE* not found.**

Explanation

An object in the cluster with the correct properties has not been found. Restoration cannot proceed without the object.

User response

Determine why the object cannot be found. Ensure that the object is available, and resubmit the command.

CMMVC6120E **Target is not the configuration node**

Explanation

The target is not the configuration node.

User response

Redirect the action against the configuration node, and resubmit the command.

CMMVC6121E **No cluster id or id_alias in backup configuration.**

Explanation

Neither the cluster id_alias nor the ID can be extracted from the backup configuration file.

User response

If the cause of the problem cannot be determined, contact IBM technical support for assistance.

CMMVC6122E **No *TYPE* with *PROPERTY VALUE* is present in the table.**

Explanation

An unexpected error has occurred.

User response

Contact IBM technical support for assistance.

CMMVC6123E **No *PROPERTY* for *TYPE NAME* .**

Explanation

An unexpected error has occurred.

User response

Contact IBM technical support for assistance.

CMMVC6124E **No *TYPE* with *PROPERTY VALUE***

Explanation

An unexpected error has occurred.

User response

Contact IBM technical support for assistance.

CMMVC6125E **No unique ID for *TYPE NAME***

Explanation

An unexpected error has occurred.

User response

Contact IBM technical support for assistance.

CMMVC6126E **No *TYPE* with unique ID *VALUE***

Explanation

An unexpected error has occurred.

User response

Contact IBM technical support for assistance.

CMMVC6127I **The SSH key *IDENTIFIER* for *USER* is already defined; the SSH key will not be restored**

Explanation

An identical SSH key for this user is already defined on the cluster. Therefore, the key in the backup file will not be restored.

User response

Specify a different SSH key, and resubmit the command.

CMMVC6128W ***DIRECTORY***

Explanation

The files in the specified directory cannot be listed.

User response

Determine why the files cannot be listed, correct the problem, and resubmit the command.

CMMVC6129E **Host mapping objects have VDisk_UID values that are not consistent.**

Explanation

All of the host mapping objects do not have the same number for the volume LUN instance. Therefore, there is a possibility the backup configuration file is corrupt. The LUN instance number should be the same for all host mapping objects that are associated with a specific volume. The LUN instance number is incorporated into the volume ID property.

User response

Determine why the LUN instance number is not the same, correct the problem, and resubmit the command.

CMMVC6130W **Inter-cluster *PROPERTY VALUE* will not be restored.**

Explanation

The restoration of inter-cluster objects is not supported.

User response

Not applicable.

CMMVC6131E **No location cluster information**

Explanation

An unexpected error has occurred.

User response

Contact IBM technical support for assistance.

CMMVC6132E **The object *OBJECT* of type *TYPE* has a property *PROPERTY* with an incorrect value *INCORRECT_VALUE* . The operation cannot proceed until the property has the correct value *CORRECT_VALUE* . Take administrative action to change the value and try again.**

Explanation

The specified object has the specified property of the specified type with the specified incorrect value. The property most likely reflects the state of the object.

User response

Change the state to the required value, and resubmit the command.

CMMVC6133E **Required *TYPE* property *PROPERTY* not found**

Explanation

An unexpected error has occurred.

User response

Contact IBM technical support for assistance.

CMMVC6134E **No argument for *OPTION***

Explanation

No argument has been supplied for the specified option, which requires an argument.

User response

Supply an argument, and resubmit the command.

CMMVC6135E **Argument *VALUE* for *OPTION* is not valid.**

Explanation

The specified argument that you have supplied is not valid for the specified option.

User response

Supply an valid argument, and resubmit the command.

CMMVC6136W **No SSH key file *FILENAME***

Explanation

The specified file, which should contain the SSH key, is not present and will not be restored. The backup operation will continue.

User response

No action is required. You might have to manually restore the key.

CMMVC6137W **No SSH key file *FILENAME*; key not restored**

Explanation

An SSH key cannot be restored because the specified file, which is expected to contain the SSH key, is not present. The restore operation will continue.

User response

After the restore is complete, locate the file containing the key, and perform one of the following actions: (1) Rename the file so that it has the correct name, and resubmit the command. (2) Restore the key manually using the addsshkey command.

CMMVC6138E ***OPTION*** **is required**

Explanation

An option is missing. The option might be listed as optional, but circumstances make the option mandatory.

User response

Supply the option, and resubmit the command.

CMMVC6139E **Incorrect XML tag nesting in** ***FILENAME***

Explanation

There is a problem with the content of a configuration file. There is a problem parsing the XML in the file, because the XML records are not consistent. The file might be corrupt, or the file has been truncated.

User response

Replace this copy with a good copy, and resubmit the command. If the problem persists, contact IBM technical support for assistance.

CMMVC6140E **No default name for type** ***TYPE***

Explanation

An unexpected error has occurred.

User response

Contact IBM technical support for assistance.

CMMVC6141E **The option** ***OPTION*** **does not support an argument.**

Explanation

An argument has been supplied for an option that does not support one.

User response

Remove the argument, and resubmit the command.

CMMVC6142E **Existing** ***OBJECT_TYPE OBJECT_NAME*** **has a non-default name.**

Explanation

The specified object in the target default cluster has a non-default name. This suggests that the cluster was customized. The cluster is therefore not suitable for restoration.

User response

Reset the cluster as per the instructions for restoring the cluster configuration, and resubmit the command.

CMMVC6143E **The required configuration file** ***FILENAME*** **does not exist.**

Explanation

A file that is critical for successful operation is missing.

User response

Check your command. Specify the correct configuration file and resubmit the command.

CMMVC6144W **The object with default name** ***NAME*** **has been restored as** ***SUBSTITUTE_NAME*** **.**

Explanation

An object with a default name bas been restored with a different name. Ensure that you account for this name change when using the restored cluster in the future. To avoid this problem in the future, choose an appropriate name for each object in the cluster.

User response

Choose an appropriate name for each object in the cluster.

CMMVC6145I **First use the** ***COMMAND*** **-prepare command.**

Explanation

This advisory is given prior to CMMVC6103E when an intermediate file is missing.

User response

.The command you submitted cannot be processed at the moment. Follow the message and submit a different command first.

CMMVC6146E **Problem parsing *OBJECT_TYPE* data: *LINE***

Explanation

An unexpected error has occurred.

User response

Contact the support center.

CMMVC6147W ***TYPE NAME* has a name beginning with *PREFIX* .**

Explanation

An object has been encountered that has a name beginning with the specified reserved prefix. The only valid reason for an object with this kind of name is that a restoration command did not complete successfully.

User response

Ensure that no object uses the reserved prefix in its name, and resubmit the command.

CMMVC6148E **Target cluster has *NUMBER_OF_EXISTING_OBJECTS* objects of type *TYPE* instead of *NUMBER_OF_REQUIRED_OBJECTS* .**

Explanation

The target cluster does not have the specified required number of objects of the specified type.

User response

Correct the problem, and resubmit the command.

CMMVC6149E **An action is required.**

Explanation

An action is required to run the command.

User response

Supply an action, and resubmit the command.

CMMVC6150E **The action *ACTION* is not valid.**

Explanation

The specified action that you have entered is not valid.

User response

Specify a valid action, and resubmit the command.

CMMVC6151E **The option *OPTION* is not valid.**

Explanation

The specified option that you have entered is not valid.

User response

Specify a valid option, and resubmit the command.

CMMVC6152E **Volume *VDISK_NAME* instance number *INSTANCE_NUMBER* is not valid.**

Explanation

The volume cannot be restored because the instance number, which must be a hexadecimal number, is not valid.

User response

Contact IBM technical support for assistance.

CMMVC6153E ***OBJECT* not consistent with *ACTION***

Explanation

The specified object is not supported for the specified action.

User response

Remove the object, and resubmit the command.

CMMVC6154E **Required *OBJECT_TYPE* property *PROPERTY_NAME* has a null value.**

Explanation

An unexpected error has occurred.

User response

Contact IBM technical support for assistance.

CMMVC6155I **The command *COMMAND* processing has completed successfully.**

Explanation

Only information and warning messages are issued.

User response

Not applicable.

CMMVC6156W ***COMMAND*** **processing completed with errors.**

Explanation

Processing was not successful.

User response

Not applicable.

CMMVC6157E **Object is required.**

Explanation

An object or target was not specified for the command.

User response

Refer to the documentation for the command and specify all required parameters.

CMMVC6164E ***variable_error_message***

Explanation

This message is generated by the system and varies depending on the circumstances in which it was created.

User response

For help with a specific message, contact your service support representative.

CMMVC6165E **The target is not the original configuration node with a WWNN of** ***WWNN_VALUE*** **.**

Explanation

A backup configuration can only be restored to the original configuration node.

User response

Re-create the default cluster with the correct configuration node, and resubmit the command.

CMMVC6166E **The property** ***PROPERTY*** **of the object** ***OBJECT*** **has changed during svcconfig restore -execute.**

Explanation

The integrity of the restoration cannot be guaranteed.

User response

Resubmit the command from svcconfig restore -prepare.

CMMVC6169E ***variable_error_message***

Explanation

This message is generated by the system and varies depending on the circumstances in which it was created.

User response

For help with a specific message, contact your service support representative.

CMMVC6171I ***percentage*****% metadata compression**

Explanation

The message displays the extent of metadata compression.

User response

This message is informational only. No action is required.

CMMVC6174I **Pausing for** ***num_minutes*** **minute(s) after node addition(s)**

Explanation

The system must process the addition of one or more nodes for the amount of time that is displayed.

User response

This message is informational only. No action is required.

CMMVC6175I **Resuming after pause**

Explanation

The system completed processing the addition of one or more nodes.

User response

This message is informational only. No action is required.

CMMVC6180E ***object_type* [*object_id*] not restored because [*object_type* | *object_property*] is missing**

Explanation

The error message has two variants. One specifies a missing object type, such as a drive, and the other specifies a missing object property, such as a name. In either case, the system cannot continue with object creation because of the missing object or property.

User response

Add the specified object or property to the system. If you believe that the object or property already exists, contact your service support representative.

CMMVC6180W ***object_type object_identifier* not restored because *object* with *property property_value* is missing**

Explanation

The system cannot continue with object creation because a dependent object is missing.

User response

Add the specified dependent object to the system. If you believe that the object already exists, contact your service support representative.

CMMVC6181E ***object_type object_id* has *property property_value*; should be *property_value_in_backup*[; use -force to override | overridden]**

Explanation

The displayed object has a property value in its configuration that does not match the value for that property in the backup.

User response

Check the properties in the current configuration. The error message might include one of the following additions:

use -force to override
If this version of the message is displayed, you can retry the command with the addition of the **-force** parameter to retain the property value that is in the current configuration.

overridden
This version of the message is displayed for informational purposes only. The property value in the current configuration is retained. No user action is required.

Otherwise, you must correct the property value in the current configuration before you retry the command.

CMMVC6182E ***object_type object_name* with *property property_value* cannot be [restored | backed up].**

Explanation

Either the displayed object was not restored, or was not backed up, as shown.

User response

The user response is different for each object type. Contact your service support representative.

CMMVC6186E **io_grp *io_group_name* restored with id *new_id* instead of *old_id* .**

Explanation

This situation can occur when the configuration node is different from the node that was used to create the original cluster. This change affects the SCSI Inquiry value for the I/O group.

User response

This message is a warning only. No user response is required.

CMMVC6187W **Extraneous *object_type object_name* discovered**

Explanation

Recovery identified an extraneous object of the displayed type and named it *object_name*.

User response

This message is only a warning. No action is required.

CMMVC6188W **VARIANT #1: Current Config inconsistent with backup *object_type object_id* does not exist in backup**
VARIANT #2: Current Config inconsistent with backup *object_type object_id* now has *wrong_property wrong_value* was *correct_value*

Explanation

VARIANT #1: The object of the displayed type and ID is in the current configuration but not in the backup file. The object will not be recovered.

VARIANT #2: The recovery process found an inconsistency. The object of the displayed type and ID now has a property with an incorrect value.

User response

VARIANT #1: Re-create the missing object in the backup file after the recovery process is complete.

VARIANT #2: Manually restore the correct value after the recovery process is complete.

CMMVC6189W **E-mail server settings could not be restored due to missing configuration information**

Explanation

An attempt was made to restore email server settings that were not found on the backup device.

User response

Make sure that you specified the correct email server and that you made no typographical errors. If everything looks correct, contact your service support representative.

CMMVC6200E **Discovery did not complete within time limit - check that this is expected**

Explanation

User response

Contact IBM Support.

CMMVC6201E **System layer must be changed to *new_layer* using chsystem before configuration can be restored**

Explanation

User response

Contact IBM Support.

CMMVC6202E **This command can only be run by the superuser**

Explanation

User response

Contact IBM Support.

CMMVC6203E **Enclosure serial number *serial_no* was found for enclosure ID *enclosure_1* but it is already in use for enclosure ID *enclosure_2***

Explanation

User response

Contact IBM Support.

CMMVC6204E **Node *node_name* in backup is part of io_grp *io_group_1*, but the node is currently set to io_grp *io_group_2***

Explanation

User response

Contact IBM Support.

CMMVC6205E **Cannot retrieve VPD for node *node_name* because it is currently in service mode.**

Explanation

User response

Contact IBM Support.

CMMVC6206E **Timeout exceeded waiting for *object_type object_id* to have expected *property_name expected_value*, actual *actual_value*.**

Explanation

User response

Contact IBM Support.

CMMVC6207E **Cannot restore encryption as more USB devices are required. Required *req_number_devices* device(s) but found *actual_number_devices* device(s).**

Explanation

User response

Add the required number of USB devices and retry the command.

CMMVC6208E **The *feature_id* feature cannot be partially configured prior to a restore operation. The current state is *current_state*.**

Explanation

User response

Contact IBM Support.

CMMVC6209E **Unable to determine system code version.**

Explanation

User response

Contact IBM Support.

CMMVC6210E **Current Config inconsistent with backup because metadatavdisk with vdisk_id *volume_id* already exists.**

Explanation

User response

Contact IBM Support.

CMMVC6211E **The system contains fewer USB ports than required to recover encryption automatically. Manually set up encryption and then run the -prepare command again.**

Explanation

User response

Manually set up encryption and then run the **-prepare** command again.

CMMVC6212E **The host port mode did not change to transitional within the time limit.**

Explanation

The value of the `fctargetportmode` variable should change to transitional within 1 minute. If it does not, the T4 recovery fails.

User response

Contact IBM Support.

CMMVC6213E **Active-active remote copy relationship *relationship_id* not added to consistency group *group_id*.**

Explanation

User response

Contact IBM Support.

CMMVC6215W **Ensure that the system certificate is trusted by all key servers before running execute.**

Explanation

A new cluster was created as part of the T4 store process. The new cluster has a new cluster (client) certificate. This certificate must be exported and installed on all key servers so that subsequent recovery commands in the execute phase, such as **mkkeyserver**, succeed.

User response

The system SSL certificate must be trusted by all key servers before you run the **svcconfig restore -execute** command. To export the system certificate, use the **chsystemcert -export** command and then install the certificate on the key servers. Alternately, to generate a signing request to be signed by a certificate authority, use the **chsystemcert -mkrequest** command and provide the Certificate Signing Request (CSR) to the certificate authority. To install the signed certificate, use the **chsystemcert -install -file** command.

CMMVC6216E **Failed to generate key server certificates from the backup file.**

Explanation

As part of the T4 restore script, if a key server configuration existed in the backup, key server SSL certificates are regenerated and written to new files, which are then used by the restore process. This error

means that the restore process failed to regenerate the key server certificates.

User response

Contact your service support representative to manually restore the key server configuration, including server certificates.

CMMVC6218E **Before you continue with the T4 recovery, the key server username and password that are configured on the system must match username and password settings in the backup file of the configuration. Current settings: username_set:%1, password_set:%2. Backup settings: username_set:%3, password_set:%4**

Explanation

The T4 restore script checks that the current key server username and password settings match the settings in the backup file. This error means that the settings do not match and must be configured before the restore can proceed.

User response

Before you continue with the T4 recovery, complete these steps:

1. Verify that the correct configuration backup file is being used.
2. Change the key server username and password on the system to match the settings in the backup file.

CMMVC6219E **The key server SSL certificates cannot be restored automatically from the backup. Manually set up key server encryption and then run the -prepare command again.**

Explanation

If key server configuration exists in the backup that is being restored during Tier 4 (T4) recovery, any key server SSL certificates are regenerated and written to new files that are used by the restoration process. This error occurs when the restoration process fails to regenerate the key server certificates because they are too large. This error typically occurs when multiple certificates are supplied, for example when a certificate chain is used.

User response

Manually configure key server encryption by using either the management GUI or the CLI, and then run the **svcconfig restore -prepare** command again.

CMMVC6220E **The LDAP administrator password cannot be restored automatically from the backup. Manually set the password and then run the -prepare command again.**

Explanation

This error occurs when the LDAP administrator password cannot be restored automatically because the required cluster AES key is no longer available.

User response

Use the **chldap -username** *LDAP_admin_username* **-password** *password* command to set the LDAP administrator user name and password manually. Then, retry the command.

CMMVC6221W **The superuser cannot be locked automatically and must be manually locked.**

Explanation

This error occurs when the configuration backup file indicates that the superuser account is locked on the system. The restoration process cannot automatically lock the superuser because the restoration process runs as the superuser, and users are not allowed to lock themselves.

User response

Use the **chuser** command to manually lock the superuser account.

CMMVC6222W **Multi-factor authentication for the superuser cannot be enabled automatically and must be manually enabled.**

Explanation:
A recovery procedure has been run on the system. Before the disaster occurred, multi-factor authentication was enabled for the superuser. To ensure access to the system is not lost during recovery, multi-factor authentication for the superuser is automatically disabled during recovery. Once the system recovery procedure has completed successfully, multi-factor authentication must be manually re-enabled for the superuser.

User response

Confirm that the system recovery is complete, then run the command **chsecurity -superusermultifactor** *yes* to enable multi-factor authentication for the superuser.

CMMVC6223E **Before you can continue with the restore, the multi-factor authentication credentials for IBM Security Verify must be entered manually to match the credentials of the previous configuration. The credentials can be obtained for your application from IBM Security Verify.**

Explanation:
As part of the **svcconfig** restore procedure, the system tries to restore the credentials used for multifactor authentication. The credentials in the backup configuration cannot be decrypted. Therefore, the credentials must be entered manually. The credentials can be obtained for the application from the IBM Security Verify.

User response

Run the **chauthmultifactorverify** command to set the OpenID Client ID, OpenID Client Secret, CLI Client ID, and CLI Client Secret, then try the restore procedure again.

CMMVC6224E **Before you can continue with the restore, the single sign-on credentials must be entered manually to match the credentials of the previous configuration.**

Explanation:
As part of the **svcconfig** restore procedure, the system will try to restore the credentials used for single sign-on. The credentials in the backup configuration cannot be decrypted and therefore the credentials must be entered manually.

User response

Run the **chauthsinglesignon** command to set the Client ID and Client Secret, then try the restore procedure again.

CMMVC6225W **The system certificate must be manually exported using the chsystemcert-export command and installed as a signer certificate in IBM Security Verify, using the id_alias from the lssystem command as the friendly name.**

Explanation:
When Multi-factor Authentication was previously configured, the system certificate was exported and installed in IBM Security Verify as a signer certificate. As part of the **svcconfig** restore procedure, a new system has been created and therefore a new system certificate has been generated. This means the old system certificate in use by IBM Security Verify is no longer valid. Until this is resolved, users will not be able to authenticate using secondary authentication factors when logging into the GUI.

User response

Once the restore procedure completes successfully, manually export the new system certificate using the **chsystemcert -export** command. Then install this as a signer certificate in IBM Security Verify. Run the **lssystem** command to find the system id_alias, then use this as the value for the friendly name in IBM Security Verify.

CMMVC6227W **Two person integrity cannot be enabled automatically from the backup. Manually enable two person integrity.**

Explanation:
This error occurs when an attempt is made to enable two person integrity automatically from the backup. Two person integrity must be enabled manually from a backup.

User response:
To manually enable two person integrity from the backup, see Configuring two person integrity (TPI) .

CMMVC6230E **Encryption was configured using only a recovery key. Manually set up encryption and then run the -prepare command again.**

Explanation

The user only had a recovery key enabled.

User response

In order to continue with restoration and the user must set up some form of encryption.

CMMVC6231E **Encryption recovery keys cannot be restored automatically. Manually configure a new recovery key after the restore process has completed**

Explanation

The user had a recovery key enabled with at least one other type of encryption.

User response

To re-enable the recovery the user must manually configure it once the restoration process has been fully complete.

CMMVC6245E The action failed because the maximum number of objects has been reached.

Explanation:
This error occurred when you tried to create a new two person integrity request, while already maximum number of requests exist. Maximum number of such requests allowed are four.

User response:
Wait until the request expires naturally or manually cancel one of the four existing request to create a new request. After that retry the command.

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.

User response

- If the MDisk identity has changed since it was last discovered, submit the 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

CMMVC6300E The create failed because the source and target volumes are members of FlashCopy mappings that belong to different I/O groups.

Explanation

All FlashCopy mappings in a tree of connected mappings must be in the same I/O group. The new FlashCopy mapping that you attempted to create would have linked two existing trees that are in different I/O groups.

User response

You have three options. The first option is to resubmit the command and specify a different source or target volume. The second option is to delete all of the existing mappings that contain the source volume and resubmit the command. The third option is to delete all of the existing mappings that contain the target volume and resubmit the command.

CMMVC6301E The create failed because the specified consistency group does not exist.

Explanation

The FlashCopy mapping was not created because the consistency group that you specified does not exist. You must create a consistency group before you can place a mapping in that group.

User response

Either create the FlashCopy consistency group that you specified and resubmit the command, or resubmit the command and specify an existing consistency group.

CMMVC6302E The create failed because the resulting tree of FlashCopy mappings would exceed the upper limit.

Explanation

Either the source volume or the target volume, or both, are already members of other FlashCopy mappings. The FlashCopy mapping was not created because the new FlashCopy mapping that you attempted to create would have linked two existing mapping trees into a single tree that exceeds the maximum number of mappings that are supported for a single tree.

User response

You have two options. The first option is to resubmit the command and specify a different source or target volume. The second option is to delete a sufficient number of the existing FlashCopy mappings in which either the source or the target volume is a member so that the combined mapping tree does not exceed the maximum number of mappings that are supported for a single tree, and resubmit the command.

CMMVC6303E The create failed because the source and target volumes are the same.

Explanation

A particular volume cannot be both the source and the target in a FlashCopy mapping. The FlashCopy mapping was not created because you have specified the same volume as both the source and the target.

User response

Resubmit the command and specify source and target volumes that are not identical.

CMMVC6304E The create failed because the source volume does not exist.

Explanation

You must specify an existing volume as the source of a FlashCopy mapping. The FlashCopy mapping was not created because the source volume that you specified does not exist.

User response

Either create the source volume that you specified and resubmit the command, or resubmit the command and specify an existing volume as the source.

CMMVC6305E The create failed because the target volume does not exist.

Explanation

You must specify an existing volume as the target of a FlashCopy mapping. The FlashCopy mapping was not created because the target volume that you specified does not exist.

User response

Either create the target volume that you specified and resubmit the command, or resubmit the command and specify an existing volume as the target.

CMMVC6306E The create operation failed because the source volume is the member of a FlashCopy mapping whose grain size is different to that specified.

Explanation

All FlashCopy mappings that are in a tree of connected mappings must have the same grain size. The FlashCopy mapping was not created because the source volume that you specified is either the source or the target volume of another FlashCopy mapping, and the grain size of the other mapping is different from the grain size that you specified for the mapping that you attempted to create.

User response

You have two options. The first option is to delete all of the FlashCopy mappings that contain the source volume that you specified where the grain size of the FlashCopy mapping is different from the grain size that you specified, and resubmit the command. The second option is to resubmit the command and do not specify the grain size attribute.

CMMVC6307E The create operation failed because the target volume is the member of a FlashCopy mapping whose grain size is different to that specified.

Explanation

All FlashCopy mappings that are in a tree of connected mappings must have the same grain size. The FlashCopy mapping was not created because the target volume that you specified is either the source or the target volume of another FlashCopy mapping, and the grain size of the other mapping is different from the grain size that you specified for the mapping that you attempted to create.

User response

You have two options. The first option is to delete all of the FlashCopy mappings that contain the target volume that you specified where the grain size of the FlashCopy mapping is different from the grain size that you specified, and resubmit the command. The second option is to resubmit the command and do not specify the grain size attribute.

CMMVC6308E The create operation failed because the source volume is the member of a FlashCopy mapping

whose IO group is different to that specified.

Explanation

All FlashCopy mappings in a tree of connected mappings must be in the same I/O group. The FlashCopy mapping was not created because the source volume that you specified is the source or target volume in another FlashCopy mapping and the I/O group of the other FlashCopy mapping is different from the I/O group that you specified.

User response

You have two options. The first option is to delete all of the FlashCopy mappings that contain the source volume that you specified where the FlashCopy mapping is in a different I/O group from the I/O group that you specified, and resubmit the command. The second option is to resubmit the command and do not specify the I/O group attribute. If you perform the second option, the default value of the I/O group attribute is used.

CMMVC6309E **The create operation failed because the target volume is the member of a FlashCopy mapping whose IO group is different to that specified.**

Explanation

All FlashCopy mappings in a tree of connected mappings must be in the same I/O group. The FlashCopy mapping was not created because the target volume that you specified is the source or target volume in another FlashCopy mapping and the I/O group of the other FlashCopy mapping is different from the I/O group that you specified.

User response

You have two options. The first option is to delete all of the FlashCopy mappings that contain the target volume that you specified where the FlashCopy mapping is in a different I/O group from the I/O group that you specified, and resubmit the command. The second option is to resubmit the command and do not specify the I/O group attribute. If you perform the second option, the default value of the I/O group attribute is used.

CMMVC6310E **The modify failed because the specified FlashCopy mapping does not exist.**

Explanation

You cannot modify a FlashCopy mapping that does not exist. The modify command failed because the FlashCopy mapping that you specified does not exist.

User response

Resubmit the command and specify an existing FlashCopy mapping.

CMMVC6311E **The command failed because the source volume is the target of a FlashCopy mapping that is in the specified consistency group.**

Explanation

A particular volume cannot be both the source of one FlashCopy mapping and the target of another FlashCopy mapping in the same consistency group. The FlashCopy mapping was not created because the source volume of the FlashCopy mapping that you attempted to create is already the target volume of a FlashCopy mapping in the consistency group that you specified.

User response

Resubmit the command and specify a different consistency group.

CMMVC6312E **The command failed because the target volume is the source of a FlashCopy mapping that is in the specified consistency group.**

Explanation

A particular volume cannot be both the source of one FlashCopy mapping and the target of another FlashCopy mapping in the same consistency group. The FlashCopy mapping was not created because the target volume of the FlashCopy mapping that you attempted to create is already the source volume of a FlashCopy mapping in the consistency group that you specified.

User response

Resubmit the command and specify a different consistency group.

CMMVC6313E **The command failed because the specified background copy rate is invalid.**

Explanation

The command failed because the background copy rate that you specified is not a supported value.

User response

Either resubmit the command and specify a supported value for the background copy rate, or resubmit the command and do not specify the background copy rate attribute. If you do not specify the background copy rate attribute, the default background copy rate value is used.

CMMVC6314E The command failed because the specified cleaning rate is not valid.

Explanation

The command failed because the cleaning rate that you specified is not a supported value.

User response

Either resubmit the command and specify a supported value for the cleaning rate, or resubmit the command and do not specify the cleaning rate attribute. If you do not specify the cleaning rate attribute, the default cleaning rate value is used.

CMMVC6315E The command failed because the specified grain size is not valid.

Explanation

The command failed because the grain size that you specified is not a supported value.

User response

Either resubmit the command and specify a supported value for the grain size, or resubmit the command and do not specify the grain size attribute. If you do not specify the grain size attribute, the default grain size value is used.

CMMVC6319E The command has failed because a combination of IPv4 and IPv6 parameters were entered.

Explanation

The task accepts either IPv4 or IPv6 parameters. You cannot specify a combination of IPv4 and IPv6 parameters for this task.

User response

Specify only IPv4 or only IPv6 parameters, and resubmit the task.

CMMVC6320E The command has failed because the IPv4 address is not valid.

Explanation

The valid IPv4 address format is d.d.d.d, where d is a decimal value from 0-255.

User response

Specify a valid IPv4 address, and resubmit the task.

CMMVC6321E The command has failed because the IPv4 subnet mask is not valid.

Explanation

The valid IPv4 address format is d.d.d.d, where d is a decimal value from 0-255.

User response

Specify a valid IPv4 subnet mask, and resubmit the task.

CMMVC6322E The command has failed because the IPv4 gateway address is not valid.

Explanation

The valid IPv4 address format is d.d.d.d, where d is a decimal value from 0-255.

User response

Specify a valid IPv4 gateway address, and resubmit the task.

CMMVC6323E The command has failed because the IPv6 address is not valid.

Explanation

Valid IPv6 address formats are:

- x:x:x:x:x:x:x:x
- x:x:x:x: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: 0:0:0:0:0: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

User response

Specify a valid IPv6 address, and resubmit the task.

CMMVC6324E **The command has failed because the IPv6 prefix is not valid.**

Explanation

The value that you entered for an IPv6 address prefix is not a valid IPv6 address prefix.

User response

Specify a valid IPv6 address prefix, and resubmit the task.

CMMVC6325E **The command has failed because the IPv6 gateway address is not valid.**

Explanation

Valid IPv6 address formats are:

- x:x:x:x:x:x:x:x
- x:x:x:x: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: 0:0:0:0:0: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

User response

Specify a valid IPv6 gateway address, and resubmit the task.

CMMVC6326E **The command has failed because the IPv4 service state address is not valid.**

Explanation

The valid IPv4 address format is d.d.d.d, where d is a decimal value from 0-255.

User response

Specify a valid IPv4 service state address, and resubmit the task.

CMMVC6327E **The command has failed because the IPv6 service state address is not valid.**

Explanation

Valid IPv6 address formats are:

- x:x:x:x:x:x:x:x
- x:x:x:x: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: 0:0:0:0:0: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

User response

Specify a valid IPv6 service state address, and resubmit the task.

CMMVC6328E **The command has failed because the console address is not valid.**

Explanation

The valid IPv4 address format is d.d.d.d, where d is a decimal value from 0-255.

Valid IPv6 address formats are:

- x:x:x:x:x:x:x:x
- x:x:x:x: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: 0:0:0:0:0: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

User response

Specify a valid console address, and resubmit the task.

CMMVC6329E The command has failed because the IP address or DNS is not valid.

Explanation

The valid IPv4 address format is d.d.d.d, where *d* is a decimal value in the range 0 - 255.

Valid IPv6 address formats are:

- x:x:x:x:x:x:x:x
- x:x:x:x:x:x:d.d.d.d

where *d* is a decimal value in the range 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: 0:0:0:0:0: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

User response

Retry the command, specifying a valid IP address or DNS name.

CMMVC6330E The command has failed because an IPv6 address was specified and the cluster does not have an IPv6 address.

Explanation

The cluster can only communicate with a server through an IPv6 address if an IPv6 cluster management IP address is configured.

User response

Either configure the cluster to have an IPv6 cluster management address or specify an IPv4 address, and resubmit the task.

Note: You do not need to remove the IPv4 address if you configure the cluster to have an IPv6 cluster management address.

CMMVC6331E The command has failed because an IPv4 address was specified and the cluster does not have an IPv4 address.

Explanation

The cluster can only communicate with a server through an IPv4 address if an IPv4 cluster management IP address is configured.

User response

Either configure the cluster to have an IPv4 cluster management address or specify an IPv6 address, and resubmit the task.

Note: You do not need to remove the IPv6 address if you configure the cluster to have an IPv4 cluster management address.

CMMVC6332E The command has failed because an IPv6 email server address was specified and the cluster does not have an IPv6 address.

Explanation

The cluster can only communicate with a server through an IPv6 address if an IPv6 cluster management IP address is configured.

User response

Either configure the cluster to have an IPv6 cluster management address or use an email server that has an IPv4 address, and resubmit the task.

Note: You do not need to remove the IPv4 address if you configure the cluster to have an IPv6 cluster management address.

CMMVC6333E **The command has failed because an IPv4 email server address was specified and the cluster does not have an IPv4 address.**

Explanation

The cluster can only communicate with a server through an IPv4 address if an IPv4 cluster management IP address is configured.

User response

Either configure the cluster to have an IPv4 cluster management address or use an email server that has an IPv6 address, and resubmit the task.

Note: You do not need to remove the IPv6 address if you configure the cluster to have an IPv4 cluster management address.

CMMVC6334E **The command failed as the email port number supplied is invalid.**

Explanation

The value that you entered for an email port number is not a valid email port number.

User response

Specify a valid email port number, and resubmit the task.

CMMVC6335E **The command failed as the combination of parameters provided are either mutually incompatible or would leave the cluster without a functioning protocol stack.**

Explanation

You have submitted a task with a combination of parameters and parameter values that is not supported or that does not provide the minimum amount of required information.

User response

Ensure that you specify a supported combination of parameters and parameter values, and resubmit the task.

CMMVC6336E **The volume copy was not created because the grain size must be 32, 64, 128 or 256.**

Explanation

You have supplied an incorrect value for the -grainsize parameter when you attempted to create a thin-provisioned volume copy.

User response

Specify a supported grain size, and resubmit the command.

CMMVC6337E **The action failed because the warning size must be a multiple of 512 bytes.**

Explanation

You are attempting to create a thin-provisioned volume copy but you have entered an incorrect value for the -warning parameter. The value can either be a percentage of the volume capacity or an absolute value that is a multiple of 512 bytes.

User response

Enter a supported warning value, and resubmit the command.

CMMVC6338E **The action failed because the warning size can not be larger than the virtual size.**

Explanation

You are attempting to create a thin-provisioned volume copy but you have entered an incorrect value for the -warning parameter. The warning value cannot be greater than the volume capacity.

User response

Enter a supported warning value, and resubmit the command.

CMMVC6339E **The volume copy was not created because the virtual size was not provided.**

Explanation

You are attempting to create an image-mode thin-provisioned volume but you did not set the -size parameter.

User response

Resubmit the command using the -size parameter.

CMMVC6340E The action failed because the value supplied for real size is not a multiple of 512 bytes.

Explanation

You are attempting to create or resize a thin-provisioned volume copy but you have entered an incorrect value for the -rsize parameter. All sizes must be integer multiples of 512 bytes.

User response

Resubmit the command using a supported -rsize parameter value.

CMMVC6341E The action failed because the volume copy is not thin-provisioned or compressed.

Explanation

You are attempting to run a command that is valid only for thin-provisioned or compressed volumes.

User response

Specify a thin-provisioned or compressed volume, and resubmit the command.

CMMVC6342E The volume copy was not shrunk because its real size cannot be less than its used size.

Explanation

You are attempting to reduce the real size that is allocated to a thin-provisioned volume copy, but the command cannot be initiated because it would make the real size less than the size that is currently used.

User response

Determine the used size of the volume copy, and resubmit the command using a -rsize parameter value that is greater than or equal to the used size.

CMMVC6343E The volume copy was not shrunk because its real size can not be negative.

Explanation

You are attempting to reduce the real size that is allocated to a thin-provisioned volume copy, but the command cannot be initiated because it would make the real size less than zero.

User response

Determine the real size of the volume copy, and resubmit the command using a supported -rsize parameter value.

CMMVC6344E The repair operation cannot start because the volume copy is already being repaired.

Explanation

You are attempting to repair a thin-provisioned or compressed volume copy, but the copy is already being repaired.

User response

Specify the correct volume and copy parameters, and resubmit the command.

CMMVC6345E The repair operation cannot start because the volume copy was created using -import but the cluster could not recognize its format.

Explanation

You are attempting to repair a thin-provisioned or compressed volume copy that is reporting corrupt metadata. The cluster cannot repair the volume copy because it was not recognized as a valid thin-provisioned or compressed volume when it was imported into this cluster. The most probable cause is that the wrong MDisk was used when the volume copy was imported.

User response

Delete the volume copy, and resubmit the import operation using the same MDisk that was exported from the original cluster.

CMMVC6346E The repair operation cannot start because the thin-provisioned volume copy was created using -import with a real size that is too small.

Explanation

You are attempting to repair a thin-provisioned volume copy that is reporting corrupt metadata. The cluster cannot repair the volume copy because although it was recognized as a valid thin-provisioned volume when it was imported into this cluster, the real size allocated to the volume copy is too small. The most probable cause is that the incorrect value was supplied with -rsize parameter when the volume copy was imported.

User response

Delete the volume copy. Resubmit the import operation either using a larger value for -rsize, or supplying the -rsize parameter without a value to let the system choose a real size.

CMMVC6347E **The specific update package cannot be installed on this hardware level.**

Explanation

The version of software that you are attempting to install does not support the hardware level of the configuration node.

User response

Check the release notes for the version of software that you want to install. Ensure that the version of software that you install supports the hardware level of all of the nodes in the cluster, and resubmit the task.

CMMVC6348E **The command failed as there was not enough information provided to process successfully.**

Explanation

You have submitted a task with a combination of parameters and parameter values that does not provide the minimum amount of required information.

User response

Ensure that you specify a supported combination of parameters and parameter values, and resubmit the task.

CMMVC6349E **The command was not initiated because the volume cache has been lost and you have not specified the -force option.**

Explanation

You must specify the -force option when you move a volume from one I/O group to another and the volume has lost cache data.

User response

Resubmit the command and specify the -force option.

CMMVC6350E **The command failed because there is insufficient mirror bitmap space.**

Explanation

The command failed because there is insufficient free memory to allocate the bitmap that is needed for volume mirroring or formatting in the I/O group. Mirroring bitmaps are temporarily used for the formatting of non-mirrored volumes and for tracking the synchronization of mirrored volumes.

User response

Complete one of the following actions:

- Submit a **chiogrp** command to increase the bitmap space.
- Remove volume mirrors from the I/O group.
- Create the volume with formatting turned off.

Resubmit the command that caused this error.

CMMVC6351E **The command failed because the volume is not mirrored.**

Explanation

Only mirrored volumes are supported for this command.

User response

Perform one of the following actions:

- Submit the appropriate command for a volume that is not mirrored.
- Submit a addvdiskcopy command to add a copy to the volume, and resubmit the command that caused this error.

CMMVC6352E **The command failed because the number of copies of this volume would exceed the limit.**

Explanation

You cannot exceed the limit on the number of copies that are supported for a volume.

User response

Submit a rmvdiskcopy or splitvdiskcopy command to decrease the number of volume copies, and resubmit the command that caused this error.

CMMVC6353E The command failed because the copy specified does not exist.

Explanation

You must specify an existing copy for this command.

User response

Submit an lsvdiskcopy command to show all of the available copies for this volume. Select a copy that exists, and then resubmit the command that caused this error.

CMMVC6354E The command failed because a copy is not synchronized.

Explanation

The copy that you specify for this command must be a synchronized copy.

User response

Use the lsvdisksyncprogress command to view the synchronization status. Wait for the copy to synchronize. If you want the synchronization process to complete more quickly, increase the rate by submitting a chvdisk command. When the copy is synchronized, resubmit the command that caused this error.

CMMVC6355E The command failed because an image mode copy is not synchronized.

Explanation

An attempt was made to remove an image mode copy, and the data on the copy is not synchronized with the host-accessible copy.

User response

Wait for the volume copy to resynchronize, then retry the command. Alternatively, specify the **-discardimage** or the **-force** parameter to force the delete operation.

CMMVC6356E The command failed because a copy is not synchronized and -force was not specified.

Explanation

When you specify a copy for this command, the copy must be synchronized unless you also specify the -force parameter.

User response

Perform one of the following actions:

- Use the lsvdisksyncprogress command to view the synchronization status. Wait for the copy to synchronize. If you want the synchronization process to complete more quickly, increase the rate by submitting a chvdisk command. When the copy is synchronized, resubmit the command that caused this error.
- Resubmit the command and specify the -force parameter.

Note: When you specify the -force parameter with the command that caused this error, the entire volume copy is resynchronized.

CMMVC6357E The command failed because the copy specified is not synchronized and -force was not specified.

Explanation

When you specify a copy for this command, the copy must be synchronized unless you also specify the -force parameter.

User response

Perform one of the following actions:

- Use the lsvdisksyncprogress command to view the synchronization status. Wait for the copy to synchronize. If you want the synchronization process to complete more quickly, increase the rate by submitting a chvdisk command. When the copy is synchronized, resubmit the command that caused this error.
- Resubmit the command and specify the -force parameter.

Note: When you specify the -force parameter with the command that caused this error, the created volume is no longer guaranteed to have identical data to the original volume when the split is performed.

CMMVC6358E The command failed because the copy specified is the only synchronized copy or in fast formatting state.

Explanation

The command failed because the copy specified is the only synchronized copy.

User response

1. Use the **lsvdisksyncprogress** command to view the synchronization status. If the specified copy is the only synchronized copy, wait for another copy to synchronize.
2. Use the **lsvdiskprogress** command to view the fast formatting progress. If the copy is formatting, wait for the formatting process to complete.
3. When the copy synchronizes and the formatting process completes, retry the command that caused the error.

You can use the **chvdisk -syncrate** command to increase the copy synchronization rate.

CMMVC6359E The command failed because there are insufficient online synchronized copies.

Explanation

This error occurs when at least one of the volume copies is offline.

User response

Fix all of the errors that are associated with the volume copies, and resubmit the command.

CMMVC6363E The command failed because the Logical Block Address (LBA) specified is invalid for this volume.

Explanation

You must specify a Logical Block Address (LBA) that is a valid address for this volume.

User response

Use the lsvdisk command to obtain the volume size, and resubmit the command that caused this error using a logical block address that is in range.

CMMVC6364E The command failed because the logical block address (LBA) requested is too large for the disk.

Explanation

You have specified an LBA in conjunction with a volume or MDisk, but the LBA is too large and does not exist on the disk.

User response

Check the size of the disk, and resubmit the command using an LBA that exists on the disk.

CMMVC6365E The command timed out.

Explanation

The command has not completed in a reasonable amount of time. Processing of the command required the software to wait for a set of MDisk reads or writes to complete, and the predefined reasonable wait time has been exceeded.

User response

Resolve any MDisk or fabric event log entries, and resubmit the command.

CMMVC6366E One or more nodes in the cluster has hardware that is not supported by the new code.

Explanation

The version of code that you are attempting to install does not support the hardware in at least one node in the cluster.

User response

Check the release notes for the version of code that you want to install. update hardware so that all of the hardware in the cluster is supported by the new version of code, and resubmit the task.

CMMVC6367E A remote cluster is running software that is incompatible with the new software package.

Explanation

The version of software that you are attempting to install on the local cluster does not support the version of software that is installed on the remote cluster.

User response

Check the release notes for the version of software that you want to install. Perform one of the following actions:

- update the software on the remote cluster to a version that is supported by the version of software that you want to install on the local cluster before you update the software on the local cluster.
- Delete the cluster partnership to stop all remote copy relationships between the clusters, and resubmit the task.

CMMVC6368E The new code might be incompatible with the remote cluster.

Explanation

The version compatibility between clusters cannot be checked because the remote cluster is not accessible.

User response

Perform one of the following actions:

- Ensure that the link to the remote cluster is functioning properly, and resubmit the task.
- Delete the cluster partnership to stop all remote copy relationships between the clusters, and resubmit the task.

CMMVC6369W The FlashCopy storage capacity that the cluster is using is approaching the FlashCopy storage capacity that is licensed.

Explanation

You are being warned that the FlashCopy storage capacity license might be exceeded soon.

User response

update the FlashCopy storage capacity license to prevent recurrence of this warning message.

CMMVC6370W The Remote Copy storage capacity that the cluster is using is approaching the Remote Copy storage capacity that is licensed.

Explanation

You are being warned that the Remote Copy storage capacity license might be exceeded soon.

User response

update the Remote Copy storage capacity license to prevent recurrence of this warning message.

CMMVC6372W The virtualized storage capacity that the cluster is using is approaching the virtualized storage capacity that is licensed.

Explanation

You are being warned that the virtualized storage capacity license might be exceeded soon.

User response

update the virtualized storage capacity license to prevent recurrence of this warning message.

CMMVC6373W The virtualized storage capacity that the cluster is using exceeds the virtualized storage capacity that is licensed.

Explanation

You are being warned that the virtualized storage capacity license has been exceeded.

User response

Update the virtualized storage capacity license to prevent recurrence of this warning message.

CMMVC6374W The FlashCopy storage capacity that the cluster is using exceeds the FlashCopy storage capacity that is licensed.

Explanation

You are being warned that the FlashCopy storage capacity license has been exceeded.

User response

update the FlashCopy storage capacity license to prevent recurrence of this warning message.

CMMVC6375W The Remote Copy storage capacity that the cluster is using exceeds the Remote Copy storage capacity that is licensed.

Explanation

You are being warned that the Remote Copy storage capacity license has been exceeded.

User response

update the Remote Copy storage capacity license to prevent recurrence of this warning message.

CMMVC6394E The command failed because an attempt to make the volume cache empty took too long.

Explanation

The failed command must empty the volume cache before attempting the requested action to ensure that data is preserved. The empty volume cache subtask has taken too long, and therefore the command that

you have submitted was not initiated so that other configuration activity can occur.

The system continues attempting to empty the volume cache.

The storage associated with the volume is probably overloaded.

User response

Wait a few minutes to allow the volume cache to empty. Resubmit the command.

Alternatively, you can use the -force parameter, if the command supports the -force parameter, to bypass the empty volume cache subtask. However, specifying the -force parameter will discard cache data for the volume. Only use the -force flag with this command if you do not intend to use the existing contents of the volume.

In addition to the above actions, investigate the performance of the network storage devices associated with this volume. The performance of host applications using these devices might be degraded.

Remedial action to resolve a performance problem enables host application performance to return to optimal conditions, and prevents this error message from recurring when you resubmit the command that caused this error.

CMMVC6399E The command failed because there is not enough memory available for reservation.

Explanation

At least one node in the cluster cannot reserve the required amount of memory. This might be caused by pinned data in the cache.

User response

Check for events in the event log. Follow the fix procedures to resolve the problem.

CMMVC6400E The command failed because a specified managed disk (MDisk) is already in use.

Explanation

You cannot specify an MDisk for this command if it is already in a storage pool or is being used as an image mode volume.

User response

Specify an MDisk that is not being used as an image mode volume and is not in a storage pool, and resubmit the command.

CMMVC6401E The command failed because one or more of the specified managed disks (MDisks) that you have specified are not in the required storage pool.

Explanation

The command requires that all of the MDisks that you specify must be in the same storage pool.

User response

Ensure that all of the MDisks that you specify are in the same storage pool, and resubmit the command.

CMMVC6402E The command failed because the managed disk (MDisk) is not in the required storage pool.

Explanation

All of the MDisks that you specify must be in the required storage pool. At least one of the source MDisks that you have specified in the command is not in the required storage pool.

User response

Ensure that all of the MDisks that you specify are in the storage pool that you specify, and resubmit the command.

CMMVC6403E The command failed because the target managed disk (MDisk) is not in the required storage pool.

Explanation

All of the MDisks that you specify must be in the required storage pool. At least one of the target MDisks that you have specified in the command is not in the required storage pool.

User response

Ensure that all of the MDisks that you specify are in the storage pool that you specify, and resubmit the command.

CMMVC6404E The command failed because the source and target storage pools must be different.

Explanation

The source and target storage pools that you specify for a cross storage pool migration must be different.

User response

Ensure that the source and target storage pools that you specify for a cross storage pool migration are different, and resubmit the command.

CMMVC6405E The command failed because the target copy was not specified.

Explanation

A target copy must be specified when you use migrations on a volume and more than one volume copy exists.

User response

Specify the target copy, and resubmit the command.

CMMVC6406E The command failed because the specified storage pool does not exist.

Explanation

At least one of the storage pools that you have specified in the parameter list does not exist.

User response

Ensure that each of the storage pools that you specify exists, and resubmit the command.

CMMVC6407E The command failed because the storage pool is invalid.

Explanation

At least one storage pool ID is above the maximum value that is available for the system.

User response

Ensure that each storage pool ID that you specify in the parameter list exists, and resubmit the command.

CMMVC6408E The command failed because too few storage pools were specified.

Explanation

You must specify the number of storage pools that is consistent with the other parameters and parameter values that you specify with the command.

User response

Refer to the command documentation for valid combinations of parameters and parameter values. Use a valid combination of parameters and values, and resubmit the command.

CMMVC6409E The command failed because too many storage pools were specified.

Explanation

You must specify the number of storage pools that is consistent with the other parameters and parameter values that you specify with the command.

User response

Refer to the command documentation for valid combinations of parameters and parameter values. Use a valid combination of parameters and values, and resubmit the command.

CMMVC6410E The command failed because too few managed disks (MDisks) were specified.

Explanation

You must specify the number of MDisks that is consistent with the other parameters and parameter values that you specify with the command.

User response

Refer to the command documentation for valid combinations of parameters and parameter values. Use a valid combination of parameters and values, and resubmit the command.

CMMVC6411E The command failed because too many managed disks (MDisks) were specified.

Explanation

You must specify the number of MDisks that is consistent with the other parameters and parameter values that you specify with the command.

User response

Refer to the command documentation for valid combinations of parameters and parameter values. Use a valid combination of parameters and values, and resubmit the command.

CMMVC6412E **The command failed because the storage pool extent size is above maximum permitted size.**

Explanation

You cannot specify a storage pools extent size that is larger the maximum size.

User response

Specify a storage pool extent size that is less than or equal to the maximum size, and resubmit the command.

CMMVC6413E **The command failed because the managed disk (MDisk) is invalid.**

Explanation

At least one MDisk ID is above the maximum value that is available for the system.

User response

Ensure that each MDisk ID that you specify in the parameter list exists, and resubmit the command.

CMMVC6414E **The command failed because the managed disk (MDisk) is currently being migrated.**

Explanation

When you submit this command, you cannot specify an MDisk that is being migrated.

User response

Either wait until the migration has completed for the MDisk that you specify, or specify a different MDisk, and resubmit the command.

CMMVC6415E **The command failed because the storage pool warning threshold is too low.**

Explanation

You must specify a storage pool warning threshold that is equal to or greater than the minimum size.

User response

Specify a storage pool warning threshold that is equal to or greater than the minimum size, and resubmit the command.

CMMVC6416E **The command failed because the storage pool warning threshold is too high.**

Explanation

You must specify a storage pool warning threshold size that is equal to or less than the size of the storage pool when all of the MDisks have been added, or you must specify a storage pool warning percentage that is equal to or less than the maximum warning threshold percentage.

User response

Specify valid values for the storage pool warning threshold size or percentage, and resubmit the command.

CMMVC6417E **The command failed because the storage pool warning threshold is invalid.**

Explanation

To specify the warning threshold there must be at least one managed MDisk in the storage pool.

User response

Ensure that there is at least one MDisk defined for the storage pool or remove the warning threshold, and resubmit the command.

CMMVC6418E **The command failed because the volume is in the process of being resized.**

Explanation

When you submit this command, you cannot specify a volume that is being resized.

User response

Wait for the resize volume operation to complete. If you still want to submit this command after the operation has completed, resubmit the command.

CMMVC6419E **The command failed because one or more of the specified managed disks (MDisks) are in the process of being deleted.**

Explanation

When you submit this command, you cannot specify an MDisk that is being deleted with the -force option.

User response

Wait for the delete MDisk operation to complete. Do not include any MDisks that have been deleted in the list of MDisks that you specify, and resubmit the command.

CMMVC6420E The command failed because the cluster was not stable when the command was issued.

Explanation:
All nodes in the I/O group must be online when you enter this command. This error indicates that at least one node went offline while the command was active.

User response:
Check the status of the nodes by using the lsnodecanister command for enclosure-based systems or the lsnode command for appliance-based systems. If any node in the specified I/O group is not currently online, follow maintenance procedures for that node. When all nodes are online, repeat the command.

CMMVC6421E The specified size exceeds the maximum permitted for this feature.

Explanation

- A **chiogrp** command was attempted where the maximum bitmap memory for one or more functions was exceeded. The maximum supported amount of bitmap memory for the RAID function, the Volume Mirroring function and the Metro Mirror or Global Mirror Copy Services functions is 512 MB. The maximum supported amount of bitmap memory for the FlashCopy® function is 2048 MB.
- Alternatively, a remote copy was attempted that used a memory size of 512 MB when the value of the **remote_copy_free_memory** variable reached 0. This scenario is unlikely.

User response

Retry the command, specifying a smaller amount of memory where necessary.

CMMVC6422E The specified size is too big. Total size of memory across all features exceeds maximum permitted.

Explanation

- The maximum combined amount of memory across all functions other than FlashCopy® is 552 MB. A **chiogrp** command was used to change the memory allocations for an I/O group where the specified **-size** value (which defaults to units of megabytes) caused the total memory to exceed the maximum.
- Alternatively, an attempt was made to create a large HyperSwap volume, but enough memory was not available to complete the request.

User response

Review the amount of memory that is used by each feature by running an **lsiogrp** command that specifies the name or ID of an I/O group. Run one or more **chiogrp** commands to redistribute the amount of memory that each feature uses without exceeding the maximum. For more information, run the **help chiogrp** command.

CMMVC6423E The Send Inventory email operation failed because email is not started.

Explanation

The send inventory email functionality has been enabled but the email service has not been started.

User response

Disable the send inventory email functionality or start the email service.

CMMVC6424E The Send Inventory email operation failed because there are no inventory email users.

Explanation

The send inventory functionality has been enabled but no email users with the ability to receive inventory emails have been created.

User response

Either turn off the send inventory email functionality or create an email user account that is capable of receiving inventory emails. Refer to the documentation for the mke-mailuser command for help on creating email users.

CMMVC6425E The action failed because the maximum number of objects has been reached.

Explanation

The action failed because the maximum number of objects has been reached.

User response

Check the object specified in the command and determine if you need to specify a different object. Make the correction and resubmit the command.

CMMVC6426E **The command failed because a specified managed disk (MDisk) is already in use.**

Explanation

You cannot specify an MDisk that is already configured as an image mode volume.

User response

Specify an unmanaged disk, and resubmit the task.

CMMVC6427E **The command failed because one or more of the specified managed disks (MDisks) are not in the required storage pool.**

Explanation

The create volume task requires that all of the MDisks that you specify must be in the same storage pool.

User response

Ensure that all of the MDisks that you specify are in the same storage pool, and resubmit the task.

CMMVC6428E **The command failed because the source managed disk (MDisk) is not in the required storage pool.**

Explanation

The task requires that all of the source MDisks that you specify must be in the same storage pool.

User response

Ensure that all of the source MDisks that you specify are in the same storage pool, and resubmit the task.

CMMVC6429E **The command failed because the target managed disk (MDisk) is not in the required storage pool.**

Explanation

The task requires that all of the target MDisks that you specify must be in the same storage pool.

User response

Ensure that all of the target MDisks that you specify are in the same storage pool, and resubmit the task.

CMMVC6430E **The command failed because the target and source storage pools must be different.**

Explanation

The cross storage pool migration task does not support specifying the same storage pool to be both the source and target storage pool.

User response

Specify a source storage pool and a target storage pool that are not identical, and resubmit the task.

CMMVC6431E **The command failed because the target copy was not specified.**

Explanation

When you use migrations on a volume and there is more than one copy, you must specify which copy to use as the target copy.

User response

Specify the target copy, and resubmit the task.

CMMVC6432E **The command failed because the specified storage pool does not exist.**

Explanation

All of the storage pools that you specify must already exist.

User response

Ensure that all of the storage pools that you specify already exist, and resubmit the task.

CMMVC6433E **The command failed because the storage pool is invalid.**

Explanation

All of the storage pool IDs that you specify must have a value that is less than or equal to the maximum supported storage pool ID value.

User response

Ensure that all storage pools have supported ID values. Ensure that all of the storage pools that you specify already exist, and resubmit the task.

CMMVC6434E The command failed because too few storage pools were specified.

Explanation

The combination of parameters and parameter values that you have specified is not supported. The task requires that you specify more storage pools than the number that you have specified.

User response

Specify a supported combination of parameters and parameter values, and resubmit the task.

CMMVC6435E The command failed because too many storage pools were specified.

Explanation

The combination of parameters and parameter values that you have specified is not supported. The task requires that you specify fewer storage pools than the number that you have specified.

User response

Specify a supported combination of parameters and parameter values, and resubmit the task.

CMMVC6436E The command failed because too few managed disks (MDisks) were specified.

Explanation

The combination of parameters and parameter values that you have specified is not supported. The task requires that you specify more MDisks than the number that you have specified.

User response

Specify a supported combination of parameters and parameter values, and resubmit the task.

CMMVC6437E The command failed because too many managed disks (MDisks) were specified.

Explanation

The combination of parameters and parameter values that you have specified is not supported. The task requires that you specify fewer MDisks than the number that you have specified.

User response

Specify a supported combination of parameters and parameter values, and resubmit the task.

CMMVC6438E The command failed because the storage pool extent size is above maximum permitted size.

Explanation

The storage pool extent size that you have specified is greater than the supported maximum value.

User response

Specify a supported storage pool extent size, and resubmit the task.

CMMVC6439E The command failed because the managed disk (MDisk) is invalid.

Explanation

Each MDisk ID must have a value that is less than or equal to the maximum supported MDisk ID value.

User response

Ensure that all of the MDisks have supported ID values. Ensure that all of the MDisks that you specify already exist, and resubmit the task.

CMMVC6440E The command failed because the managed disk (MDisk) is currently being migrated.

Explanation

When you submit this task, you cannot specify an MDisk that is being migrated.

User response

Ensure that the MDisk that you specify is not migrating, and resubmit the task. If you want to specify the same MDisk and resubmit the task, ensure that the migration for that MDisk has completed before you resubmit the task.

CMMVC6441E The command failed because the storage pool warning threshold is too low.

Explanation

The value that you have specified for the storage pool warning threshold is less than the minimum supported value.

User response

Specify a supported value for the storage pool warning threshold, and resubmit the task.

CMMVC6442E The command failed because the storage pool warning threshold is too high.

Explanation

Either the value for the storage pool warning percentage is greater than the maximum supported value, or the storage pool warning disk size is greater than the storage pool capacity.

User response

Specify supported values for storage pool warning percentage and disk size, and resubmit the task.

CMMVC6443E The command failed because the storage pool warning threshold is invalid.

Explanation

If you submit this command and specify a storage pool warning threshold percentage, you must specify a storage pool that contains at least one MDisk and you must specify a supported value for the storage pool warning threshold percentage.

User response

Either do not specify a storage pool warning threshold percentage, or specify a supported value for the storage pool warning threshold percentage and specify a storage pool that contains at least one MDisk, and resubmit the task.

CMMVC6444E The command failed because the volume is in the process of being resized.

Explanation

You cannot specify a volume that is being resized when you submit this task.

User response

Wait for the resize volume task to complete. You can specify the same volume and resubmit this task only after the resize volume task that is in progress has completed.

CMMVC6445E The command failed because one or more of the specified managed disks (MDisks) are in the process of being deleted.

Explanation

You cannot specify an MDisk that is being force deleted.

User response

Wait until all force delete MDisk tasks have completed. Ensure that all of the MDisks that you specify still exist, and resubmit the task.

CMMVC6446E The command failed because the storage pools have different extent sizes.

Explanation

This task requires that the extent size of the source storage pool and the extent size of the target storage pool must be identical.

User response

If you want to resubmit this command, ensure that the source and target storage pools have the same extent size. If you want to move a volume to a storage pool that has a different extent size, you must use the procedure that is documented in the technical notes.

CMMVC6447E The command failed because the volume is currently being migrated.

Explanation

You cannot specify a volume that is being migrated.

User response

Either wait until the volume migration process has completed and resubmit the task, or specify a volume that is not being migrated and resubmit the task.

CMMVC6448E Deleting this node will cause data loss for resources associated with the I/O group of this node.

Explanation

This node contains resources which are vital to the I/O group and unavailable elsewhere. Removing this node will cause a loss of customer data.

It is recommended that this node not be removed unless the customer data supported by it is of no consequence.

User response

The -force option must be used to remove this node.

CMMVC6449E The operation was not performed because the partnership owns Global or Metro Mirror relationships or consistency groups.

Explanation

The cluster partnership cannot be removed while there are Global or Metro Mirror relationships or consistency groups that are configured in the local cluster and that are associated with the remote cluster of the partnership.

User response

Identify all of the Global or Metro Mirror relationships or consistency groups in the local cluster that are configured between this cluster and the remote cluster of the partnership. Remove all of the relationships and groups that you have identified, and resubmit the task.

Note: Do not remove relationships or groups that are associated with a different cluster, and do not remove relationships or groups that are contained entirely within the local cluster.

CMMVC6450W A FlashCopy mapping was created but physical_flash is not enabled.

Explanation

The create FlashCopy mapping task has succeeded. However, physical_flash should be enabled when you create a FlashCopy mapping in the physical disk license scheme.

User response

Ensure that you have the appropriate virtualization license for the cluster configuration that you want to enable. Ensure that the license settings for this cluster match the license.

Delete the FlashCopy mapping or enable physical_flash.

CMMVC6451W A Global Mirror or Metro Mirror relationship was created but physical_remote is not enabled.

Explanation

The create Global Mirror or Metro Mirror relationship task has succeeded. However, physical_remote should be enabled when you create a Global Mirror or Metro Mirror relationship and the cluster uses the physical disk license scheme.

User response

Ensure that you have the appropriate virtualization license for the cluster configuration that you want to enable. Ensure that the license settings for this cluster match the license.

Delete the Global Mirror or Metro Mirror relationship or enable physical_remote.

CMMVC6452W You are using the physical disk license scheme but the values for physical_flash and physical_remote are not set.

Explanation

The task has succeeded. However, you should enable physical_flash before you create a FlashCopy mapping and you should enable physical_remote before you create a Global Mirror or Metro Mirror mapping.

User response

Enable physical_flash before you create a FlashCopy mapping. Enable physical_remote before you create a Global Mirror or Metro Mirror relationship.

CMMVC6453W You have disabled the physical disk license scheme but the capacity license scheme is not set.

Explanation

The task has succeeded. However, you should configure a license scheme before you create a FlashCopy, Global Mirror or Metro Mirror relationship. You can configure a physical disk license scheme or a capacity license scheme, but not both.

User response

If you do not have a virtualization feature license that is valid for this cluster, contact your IBM sales representative and obtain a license. Ensure that the license settings for this cluster match the license that you have for this cluster.

CMMVC6454E The command failed because the physical disk license scheme is not enabled.

Explanation

You can only enable physical_flash or physical_remote when the physical disk license scheme is enabled.

User response

Ensure that you have the appropriate virtualization license for the cluster configuration that you want to enable. Ensure that the license settings for this cluster match the license. Resubmit the task if it supported by the license.

CMMVC6455E The command failed because a capacity license scheme parameter was specified but the physical disk license scheme is enabled.

Explanation

You cannot enable the capacity license scheme or specify a capacity license scheme parameter while the cluster is using the physical disk license scheme.

User response

Ensure that you have the appropriate virtualization license for the cluster configuration that you want to enable. Ensure that the license settings for this cluster match the license. Resubmit the task if it supported by the license.

CMMVC6456E The command failed because a physical disk license scheme parameter was specified but the capacity license scheme is enabled.

Explanation

You cannot enable the physical disk license scheme or specify a physical disk license scheme parameter while the cluster is using the capacity license scheme.

User response

Ensure that you have the appropriate virtualization license for the cluster configuration that you want to enable. Ensure that the license settings for this cluster match the license. Resubmit the task if it supported by the license.

CMMVC6457E One or more quorum disks are on the specified controller.

Explanation

You cannot disable the setting that allows a controller to support a quorum disk while a quorum disk is configured on the controller.

User response

Move all quorum disks from the controller to a different storage system using the setquorum command, and resubmit this task.

CMMVC6458E The specified controller cannot support quorum disks.

Explanation

The controller type of the controller that you specified does not support quorum disks.

User response

Specify a controller that has a controller type that supports quorum disks, and resubmit the task.

CMMVC6459E The mkrcrelationship command failed because the same volume was specified as the master and auxiliary volume.

Explanation

A relationship cannot be created from a volume to itself. The mkrcrelationship command requires that you specify two different volumes for the master and auxiliary positions. These can be two volumes in the local cluster, or a volume in each of two different clusters.

User response

Specify a master volume and an auxiliary volume that are not identical to each other, and resubmit the task.

CMMVC6460E The command failed because the migration source is offline.

Explanation

The source of the migration is offline. The offline source is either an image mode MDisk or the entire storage pool.

User response

- If you submitted the rmmdisk command and specified a regular MDisk, determine the storage pool to which the source MDisk is defined, and follow the procedure for bringing the storage pool

online. There will be an entry in the event log for the corresponding storage pool.

- If you submitted the rmmdisk command and specified an image mode MDisk, determine the source MDisk and follow the procedure for bringing the image mode MDisk online. There will be an entry in the event log for the corresponding MDisks.
- If you submitted a command to migrate a copy of an image mode volume, determine the corresponding source MDisk and follow the procedure for diagnosing problems with MDisks. There will be an entry in the event log for the corresponding MDisks.
- If you submitted any other command to migrate a volume copy, determine the storage pool to which the volume is defined, and follow the procedure for bringing the storage pool online. There will be an entry in the event log for the corresponding storage pool.

CMMVC6461E **The command failed because starting the migration will result in volumes going offline in the source storage pool.**

Explanation

A migration from an image mode volume will use the source storage pool and the source storage pool assumes the combined state of the image mode MDisk and the storage pool. If the online or offline states of the image mode MDisk and the storage pool are different on different nodes, the source volume might go offline or all of the volumes in the source storage pool might go offline.

User response

For each node, note the online or offline states of the source volume and the source storage pool. If one entity is online and the other is offline, bring online whichever is offline. Taking the online entity offline is not recommended because other volumes might go offline.

CMMVC6462E **The command failed because starting the migration will result in volumes going offline because the target storage pool is offline.**

Explanation

The migration process assigns the volume an online or offline state based on the states of the source and target storage pools. In this case, based on the offline state of the target storage pool the volume that is currently online would have been taken offline. The command cannot be initiated because this action is not supported. There will be an entry in the event log for the corresponding storage pool.

User response

For each node, note the online or offline state of the source and target storage pools. For each node, if one of these two storage pools is online and the other is offline, bring online whichever storage pool is offline. Taking the online storage pool offline is not recommended because other volumes might go offline.

CMMVC6463E **The command failed because Starting the migration will result in volumes going offline because a target MDisk is offline.**

Explanation

The volume is currently online. The migration process assigns the volume an online or offline state based on the states of the source and target MDisks. In this case, based on the offline state of the target MDisk, the volume would have been taken offline. The task cannot be initiated because this action is not supported.

User response

Bring the target MDisk online by following the recommended procedure for bringing an MDisk online, and resubmit the command.

CMMVC6464E **The Create FlashCopy mapping task cannot be initiated because the size of the source volume is being changed by a previously submitted task.**

Explanation

You cannot submit this task while the Change volume size task is in progress.

User response

Wait until the Change volume size task completes, and then resubmit the task.

CMMVC6465E **The Create FlashCopy mapping task cannot be initiated because the size of the target volume is being changed by a previously submitted task.**

Explanation

You cannot submit this task while the Change volume size task is in progress.

User response

Wait until the Change volume size task completes, and then resubmit the task.

CMMVC6466E **The Create FlashCopy mapping task cannot be initiated because an identical map already exists.**

Explanation

A map between the source and target volumes that you have specified is defined. You cannot define a map that is exactly the same as a map that is already defined.

User response

Specify a unique map when you submit this task.

CMMVC6467E **The Create FlashCopy mapping task cannot be initiated because a FlashCopy map with the same target volume already exists in the consistency group.**

Explanation

You cannot create more than one FlashCopy map with the same target volume in the same consistency group.

User response

Specify a target volume for the FlashCopy map that is unique to the consistency group when you submit this task.

CMMVC6468E **The Start or Prepare FlashCopy mapping task cannot be initiated because the target volume is the source of a different FlashCopy map that is being restored.**

Explanation

You cannot start or prepare a map while the target of the map is the source volume of another FlashCopy mapping that is being restored.

User response

Ensure that the target volume in the map that you are attempting to start or prepare is not the source volume of another FlashCopy mapping that is being restored when you submit the task. You could stop the associated map that is being restored, or you could wait for the map that is being restored to reach the Idle_or_Copied state.

CMMVC6469E **The Split stop FlashCopy map task cannot be initiated because the mapping is either being restored or is not in the copy complete state.**

Explanation

You cannot split stop a FlashCopy map while it is being restored or is not in the copy complete state.

User response

Ensure that the map is not being restored and is in the copy complete state when you submit this task.

CMMVC6470E **The Start or Prepare FlashCopy mapping task cannot be initiated because the target volume is being used by a different FlashCopy map.**

Explanation

You cannot start or prepare a map while the target of the map is also the target volume of another map that is in one of the following states: copying, stopping, suspended, prepared or preparing.

User response

Ensure that the target volume in the map that you are attempting to start or prepare is not the target volume of another FlashCopy mapping that is in one of the unsupported states when you submit this task.

CMMVC6471E **The Create cluster partnership task cannot be initiated because a cluster in the existing partnership has a downlevel code version that does not support this configuration.**

Explanation

One scenario in which this error occurs is when a cluster at a higher version is partnered to a cluster at a lower version that does not support multiple cluster mirroring, and you attempt to create another partnership to a cluster at the higher version to implement multiple cluster mirroring. Adding a partnership to a third cluster is not supported while at least one cluster in the current partnership is at the lower version.

User response

Either update the downlevel cluster software version to a version that supports this task or remove the

partnership to the cluster that has the downlevel software version, and resubmit the task.

CMMVC6472E **The Create cluster partnership task cannot be initiated because the remote cluster with which you are attempting to create a partnership has a downlevel code version that does not support this configuration.**

Explanation

The code versions of the clusters in the existing partnership do not support a partnership with a cluster that has the code version of the remote cluster with which you are attempting to create a partnership. If a cluster at version 5.1.0 or later is already in a partnership with another cluster at version 5.1.0 or later, you can only add a partnership to a cluster at version 5.1.0 or later, and cannot add a partnership to a cluster at version 4.3.1 or earlier. If a cluster at version 5.1.0 or later is already in a partnership with another cluster at version 4.3.1 or earlier, you cannot add another partnership while the partnership with the cluster at version 4.3.1 exists. If a cluster is not in a partnership, you can create a partnership between it and a cluster at any version. One scenario in which this error occurs is when you attempt to add a partnership with a remote cluster at version 4.3.1 or earlier to a cluster at version 5.1.0 or later that is already in partnership with another cluster at version 5.1.0 or later.

User response

Either update the downlevel cluster code version to a version that supports this task or remove all existing partnerships from the cluster to which you want to partner the cluster that has the downlevel version, and resubmit the task.

CMMVC6473E **The partnership task cannot be initiated because the supported maximum number of accessible remote clusters would be exceeded.**

Explanation

With multiple cluster mirroring, you can build a configuration of a chain of clusters. There is a limit to the number of clusters that you can configure in the chain. The task would have resulted in exceeding the supported maximum number of clusters in a chain.

User response

Ensure that the resulting configuration is supported when you submit this task.

CMMVC6474E **The Create partnership task cannot be initiated because there is a Global Mirror or Metro Mirror relationship or consistency group that has a deleted partnership.**

Explanation

You must resolve the unpartnered objects error that is related to the deleted partnership with a Global Mirror or Metro Mirror relationship or consistency group before you can create a partnership from the local cluster to more than one other cluster.

User response

Resolve the unpartnered objects error, and resubmit the task. To resolve the error, either delete the unpartnered Global Mirror or Metro Mirror relationship or consistency group from the deleted partnership, or create a partnership for the unpartnered objects.

CMMVC6475E **The Add relationship to group task cannot be initiated because the master cluster of the relationship that you are attempting to add to the group is the auxiliary cluster of the group, and the auxiliary cluster of the relationship that you are attempting to add to the group is the master cluster of the group.**

Explanation

All of the relationships within a group must have the same master cluster as the group and must have the same auxiliary cluster as the group. The determination as to which cluster is assigned as the master cluster when you create a relationship or consistency group is based on the cluster from which you submit the task.

User response

Perform one of the following three options:

- Delete the group and create the group so that the master cluster of the group is identical to the master cluster of the relationship and the auxiliary cluster of the group is identical to the auxiliary cluster of the relationship.
- Delete the relationship and create the relationship so that the master cluster of the relationship is identical to the master cluster of the group and the

auxiliary cluster of the relationship is identical to the auxiliary cluster of the group.

- Specify a group and a relationship that have identical master clusters and identical auxiliary clusters.

Resubmit the task.

CMMVC6478E The Enable remote authentication service task cannot be initiated because the server settings are not configured.

Explanation

You cannot enable the remote authentication service until the server has been configured with all of the required settings. You must specify the user name, password, and remote authentication server URL, and if required, the SSL certificate.

User response

Ensure that the server settings are configured correctly, and resubmit the task.

CMMVC6479E The task cannot be initiated because the user group table is full.

Explanation

The maximum supported number of user groups is already configured in the user group table.

User response

Delete a user group that is not required from the table, and resubmit the task.

CMMVC6480E The task cannot be initiated because the user group that you have specified is not defined.

Explanation

You must specify a user group that exists in the user group table.

User response

Either create the user group that you had specified or specify an existing user group, and resubmit the task.

CMMVC6481E The Modify user group task cannot be initiated because you have specified a default user group.

Explanation

Examples of default user groups are SecurityAdmin, Administrator, CopyOperator, Service, and Monitor.

User response

Specify a user group that is not a default user group when you submit this task.

CMMVC6482E The Delete user group task cannot be initiated because you have specified a default user group.

Explanation

Examples of default user groups are SecurityAdmin, Administrator, CopyOperator, Service, and Monitor.

User response

Specify a user group that is not a default user group when you submit this task.

CMMVC6483E The task cannot be initiated because the user group name that you have specified already exists.

Explanation

Each user group must have a unique name.

User response

If you want to define a new user group with the name that you had specified, you must first delete the existing user group that has that same name. Specify a user group name that does not exist when you submit this task.

CMMVC6484E The task cannot be initiated because the role that you have specified is not supported.

Explanation

Examples of valid roles are SecurityAdmin, Administrator, CopyOperator, Service, and Monitor.

User response

Specify a supported role, and resubmit the task.

CMMVC6485E The Delete user group task has failed because there is at least one user that is defined as a member of the group, and you did not specify the -force parameter.

Explanation

You cannot delete a user group that is not empty unless you specify the -force parameter. If you use the -force parameter when you delete a user group, all of the users that were in the deleted user group are added to the Monitor user group.

User response

Ensure that you specify the correct user group. For each member of the specified user group that you want to belong to a user group other than Monitor, move that member to the desired group. If the user group has at least one member, specify the -force parameter when you submit the task.

CMMVC6486E **The task cannot be initiated because the user table is full.**

Explanation

The maximum supported number of users is already configured in the user table.

User response

Delete a user that is not required from the table, and resubmit the task.

CMMVC6487E **The task cannot be initiated because the user name that you have specified already exists.**

Explanation

Each user must have a unique name.

User response

If you want to define a new user with the name that you had specified, you must first delete the existing user that has that same name. Specify a user name that does not exist when you submit this task.

CMMVC6488E **The task cannot be initiated because you have specified a user group ID that is not correct.**

Explanation

You must specify a valid user group ID when you submit this task.

User response

Specify a valid user group ID, and resubmit the task.

CMMVC6489E **The task cannot be initiated because you have specified more than one password.**

Explanation

This task allows you to specify only one password.

User response

Specify only one password, and resubmit the task.

CMMVC6490E **The task cannot be initiated because you have specified both a user group and the use of the remote authentication service.**

Explanation

You cannot specify a user group when you specify the use of the remote authentication service.

User response

Either specify a user group or specify the use of the remote authentication service, but not both, and resubmit the task.

CMMVC6491E **The task cannot be initiated because an SSH key and password were not specified for the remote authentication service.**

Explanation:
An SSH key and password are required for the remote authentication service.

User response

Specify a valid SSH key and password when issuing this task.

CMMVC6492E **The task cannot be initiated because you have specified a local user but you have not specified a user group.**

Explanation

You must specify a user group when you specify a local user for this task.

User response

Specify a valid user group if you specify a local user when you submit this task.

CMMVC6493E **The task cannot be initiated because the user that you have specified is not defined.**

Explanation

You must specify a user that exists in the user table.

User response

Either create the user that you had specified or specify an existing user, and resubmit the task.

CMMVC6494E The task cannot be initiated because you cannot remove a default user.

Explanation

Examples of default users are SecurityAdmin, Administrator, CopyOperator, Service, and Monitor.

User response

Specify a user that is not a default user when you submit this task.

CMMVC6495E The task cannot be initiated because the user superuser must be a local user.

Explanation

You cannot define the user superuser to use the remote authentication service.

User response

Ensure that you have specified the correct user, and resubmit the task.

CMMVC6496E The task cannot be initiated because you cannot remove the superuser password.

Explanation

The user superuser must always have a password defined.

User response

Ensure that you have specified the correct user when you submit the task.

CMMVC6497E The task cannot be initiated because the user that you have specified does not have a password defined.

Explanation

You cannot remove a password that does not exist.

User response

Ensure that you have specified the correct user when you submit the task.

CMMVC6498E The task cannot be initiated because the user that you have specified does not have an SSH key defined.

Explanation

You cannot remove an SSH key that does not exist.

User response

Ensure that you have specified the correct user when you submit the task.

CMMVC6499E The task has failed because the SSH key that you have specified is already defined for another user.

Explanation

A single SSH key cannot be defined for more than one user.

User response

Either specify a unique SSH key for the user that you had specified or delete the user that has the SSH key that you had specified, and resubmit the task.

CMMVC6500E The action failed because the source and destination volumes are the same.

Explanation

The action failed because the source and destination volumes are the same.

User response

Recheck your command and make the correction to the source or destination or both that you have specified. Then resubmit the command after you have made the correction.

CMMVC6501E The action failed because the node hardware is incompatible with the current I/O group member.

Explanation

The action failed because the node hardware is incompatible with the current I/O group member.

User response

Recheck your command for the specified I/O group and verify that it is correct. Make the correction and resubmit the command. If your original command is correct, do further research to correct this pproblem.

CMMVC6502E The FlashCopy mapping was not prepared because preparing consistency group 0 is not a valid operation.

Explanation

The FlashCopy mapping was not prepared because preparing consistency group 0 is not a valid operation.

User response

Recheck your command and ensure that you specified the correct consistency group. Make the correction and resubmit the command. If you specified the correct consistency group, then more research is necessary to correct this problem.

CMMVC6503E The FlashCopy mapping or consistency group was not stopped because stopping consistency group 0 is not a valid operation.

Explanation

The FlashCopy mapping or consistency group was not stopped because stopping consistency group 0 is not a valid operation.

User response

Recheck your command to ensure that you specified the FlashCopy mapping or consistency group that you intended. Make the correction and resubmit the command. If the command was correct, more research is needed before you can resubmit the command.

CMMVC6504E The task cannot be initiated because the SSH key file that you have specified does not contain a valid SSH key.

Explanation

You must specify an SSH key file that contains a valid SSH key.

User response

Specify an SSH key file that contains a valid SSH key, and resubmit the task.

CMMVC6505E The task has failed because an error has occurred while communicating with the authentication service.

Explanation

The cluster is configured to use an authentication service to control which users are authorized to access the cluster. An error has occurred while the cluster was attempting to contact the authentication service. The error is probably the result of an incorrect configuration, either of the cluster or of the authentication service. This error occurs if the SSL certificate, user name or password is incorrect.

User response

Ensure that the authentication service is functioning properly. Ensure that the cluster authentication service configuration is correct. Resubmit the task.

CMMVC6506E The task has failed because a timeout has occurred while communicating with the authentication service.

Explanation

The cluster is configured to use an authentication service to control which users are authorized to access the cluster. A timeout has occurred while the cluster was attempting to contact the authentication service. This timeout is probably the result of a TCP/IP network problem or of an incorrect configuration. Configuring the incorrect IP address or protocol in the authentication service URL causes this error. The protocol can be either http or https.

User response

Ensure that the cluster authentication service configuration is correct. Ensure that the Ethernet network between the cluster and the authentication service is functioning properly. Ensure that the authentication service is functioning properly. Resubmit the task.

CMMVC6507E The task has failed because the authentication service reports an incorrect user name or password.

Explanation

The cluster is configured to use an authentication service to control which users are authorized to access the cluster.

If the password for the user name has recently been changed on the authentication service, you might be required to force the cluster to refresh its authentication cache. You can force the refresh using the cluster console View Cluster Properties, Remote Authentication panel or by submitting the Command-Line Interface command chauthservice -refresh.

User response

Ensure that the user name and password that you use is correct.

If the password for the user name has recently been changed on the authentication service, force the cluster to refresh its authentication cache.

If the user name that you are using also has a password configured on the cluster, ensure that the password that is configured on the cluster is identical to the password that is configured for that user name on the authentication service.

Resubmit the task.

CMMVC6508E **The task has failed because the authentication service reports that the authentication token has expired.**

Explanation

The cluster is configured to use an authentication service to control which users are authorized to access the cluster. The authentication token, which is saved as a browser cookie, has expired. You can modify the token expiration property that is set by the authentication service to reduce the frequency of this error in the future.

User response

Either acquire a new authentication token or log in with a user name and password, and resubmit the task.

CMMVC6510E **The task has failed because the user name or password is not correct.**

Explanation

The password that you are using does not match the password that is configured on the cluster for the user name that you are using.

User response

Enter the correct user name or password, and resubmit the task.

CMMVC6511E **The task has failed because the cluster is not configured correctly to use the authentication service.**

Explanation

The user name that you are using is configured to be authenticated using an authentication service, but either the cluster is not configured to use an authentication service or the function is not enabled.

User response

If you want to use an authentication service, configure the cluster to use the service.

If you do not want to use an authentication service, modify the configuration of the user name on the cluster to remove the designation for the use of the authentication service.

Resubmit the task.

CMMVC6512E **The task has failed because you cannot both create a new quorum disk and set that new disk to active using the same command.**

Explanation

The create new quorum disk task and set disk to active task must be done using two separate tasks.

User response

Submit a create new quorum disk task. When that task has completed, submit a task to activate the new disk.

CMMVC6513E **The task has failed because you cannot activate a quorum disk until all of the quorum disks have been initialized.**

Explanation

The initialization process for at least one disk has not yet completed. You cannot select a disk as the active disk until the initialize process for all of the quorum disks has completed.

User response

Wait until the initialize quorum disk process has completed for all of the quorum disks, and resubmit the task.

CMMVC6514E **The task has failed because the disk that you have selected to activate is not online.**

Explanation

A disk must be online to be eligible for activation.

User response

Either bring the disk that you have selected online or select a different disk that is already online, and resubmit the task.

CMMVC6515E The task has failed because at least one quorum disk is in the Excluded state.

Explanation

You cannot activate a quorum disk when one or more of the quorum disks are in the Excluded state.

User response

Either create additional quorum disks or change the configuration so that none of the quorum disks is in the Excluded state. Ensure that none of the quorum disks are in the Excluded state, and resubmit the task.

CMMVC6516E The command has failed because an IPv4 cluster address cannot be removed while remote IPv4 services are configured.

Explanation

The configured management IP address protocols determine whether IPv4 or IPv6 or both are enabled on the cluster. If a cluster does not have an IPv4 cluster address the IPv4 protocol stack will not be enabled, and therefore remote services such as email servers or SNMP servers cannot be accessed through an IPv4 address.

User response

If you can only access the service through an IPv4 address and you need to continue to use the service, you will also have to continue to specify an IPv4 cluster address even if you do not intend to manage your cluster through this address.

Otherwise, re-configure the cluster so that all remote services use only IPv6 addresses, and resubmit the task to remove the IPv4 cluster address.

CMMVC6517E The command has failed because an IPv6 cluster address cannot be removed while remote IPv6 services are configured.

Explanation

The configured management IP address protocols determine whether IPv4 or IPv6 or both are enabled on the cluster. If a cluster does not have an IPv6 cluster address the IPv6 protocol stack will not be enabled, and therefore remote services such as email servers or SNMP servers cannot be accessed through an IPv6 address.

User response

If you can only access the service through an IPv6 address and you need to continue to use the service, you will also have to continue to specify an IPv6 cluster address even if you do not intend to manage your cluster through this address.

Otherwise, re-configure the cluster so that all remote services use only IPv4 addresses, and resubmit the task to remove the IPv6 cluster address.

CMMVC6518E The task has failed because no roles are defined for the current user on the cluster.

Explanation

The cluster has been configured to use an authentication service to control which users are authorized to access the cluster. The user's credentials were accepted by the authentication service, but none of the groups defined for the user on the authentication service match any of the user groups that are defined on the cluster.

User response

Perform the following steps in sequence:

1. Determine which user groups are defined for the user on the authentication service.
2. Ensure that at least one user group that is defined for the user on the authentication service is also defined on the cluster.
3. Ensure that at least one user group that is defined for the user on both the authentication service and the cluster has its 'remote' parameter set to 'enabled'.
4. Resubmit the task.

CMMVC6519E The task has failed because you cannot change the user group of the 'superuser' account to anything other than 'SecurityAdmin'.

Explanation

The user group that is assigned to the user name 'superuser' must always be 'SecurityAdmin'. This assignment cannot be changed.

User response

Ensure that you specify a user account other than 'superuser' if you submit a task to change the user group of a user account from 'SecurityAdmin' to a different user group.

CMMVC6520E **You cannot use this task to modify the properties of the current user because those properties are only defined by an authentication service.**

Explanation

The current user is not defined on the cluster. The current user is defined on an authentication service, and the cluster is configured to use that authentication service. You must use the authentication service to change the current user's password.

If you want to enable command-line interface (CLI) access to the cluster by using an SSH key, you must define the current user on the cluster and associate the SSH key with that user. If you also want to continue using the authentication service for the current user, you must enable the 'remote' setting for the new current user account that you create on the cluster.

User response

If you want to change your password, use the authentication service for that task.

If you want to enable command-line interface (CLI) access to the cluster by using an SSH key, define your user account on the cluster and associate the ssh key with that definition. If you also want to continue using the authentication service to authorize your user account, enable the 'remote' setting for your newly created user account on the cluster.

CMMVC6521E **The task cannot be initiated because it would have resulted in a user account definition for a local user that specifies neither a password, an SSH key nor a certuid.**

Explanation

The definition of a local user must always specify neither a password, an SSH key nor a certuid.

User response

When you submit this task, ensure that you have specified the correct user account and parameters, and that all local user definitions would still specify neither a password, an SSH key nor a certuid after the task completes.

CMMVC6522E **Authorization has failed.**

Explanation

An SSH login attempt has failed. This message will be followed by a second message that will contain detailed information about the cause of the error.

User response

Follow the instructions in the second error message to solve the problem.

CMMVC6523E **The URL that you have entered is not valid.**

Explanation

The URL must start with either http:// or https:// and must use only the following characters: A-Z, a-z, 0-9, - _ : [ ] . ~ / %.

User response

Ensure that the URL that you enter starts with one of the supported strings and contains only supported characters, and resubmit the task.

CMMVC6524E **The name that you have entered is not valid. The name cannot begin or end with a space character, and the name cannot contain any of the following characters: * : , \' %**

Explanation

A space cannot be the first or last character in the name that you enter. Also, the following characters are not supported anywhere in the name: * : , \"'' %

User response

Ensure that the name that you enter does not begin or end with the space character and that it does not contain any of the unsupported characters listed above, and resubmit the task.

CMMVC6525E **The password that you have entered is not valid. The password cannot begin or end with a space character.**

Explanation

A space cannot be the first or last character in the password that you enter.

User response

Ensure that the password that you enter does not begin or end with the space character, and resubmit the task.

CMMVC6526E **The Create Volume task cannot be initiated because the number of copies that you have requested is not equal to the number of unique storage pools that you have specified.**

Explanation

When you submit this task, you must specify a unique storage pool for each volume copy that you request.

User response

Specify the same number of unique storage pools as the number of volume copies that you request, and resubmit the task.

CMMVC6527E **The name that you have entered is not valid. The name can contain letters, numbers, spaces, periods, dashes, and underscores. The name must begin with a letter or an underscore. The name must not begin or end with a space.**

Explanation

A number or space cannot be the first character and a space cannot be the last character in the name that you enter. Also, the following characters are not supported anywhere in the name: * : , "" ' % #

User response

Ensure that the name you enter does not begin with a number, does not begin or end with a space character, and does not contain any of the unsupported characters listed above, and resubmit the task.

CMMVC6528E **The command cannot be initiated because the MDisk mode is not set to Array.**

Explanation

Any MDisk that you specify for this command must be a local MDisk that is an array of LDisks. The mode of the MDisk that you have specified is not Array.

User response

Either select a different MDisk that is a local MDisk and is an array of LDisks, or configure the system so that the MDisk that you have specified is a local MDisk and is an array of LDisks, and resubmit the command.

CMMVC6529E **The command cannot be initiated because the maximum supported number of MDisks already exists.**

Explanation

This command requires that an MDisk is available for array creation. There are no available MDisks for array creation because the maximum number of MDisks is already configured on the cluster.

User response

Ensure that a local MDisk is available, and resubmit the command. To make a local MDisk available for this task, either delete an array on an existing local MDisk or remove a SAN attached MDisk and configure a local MDisk.

Some other considerations to fix this error:

- Each distributed array occupies 16 slots, starting at an mdisk id divisible by 16, therefore you might also consider the need to delete 16 unwanted mdisks starting on a mdiskid boundary divisible by 16. If you need additional information see lsmdisk.
- Or, you can delete an unwanted distributed array to make space for a distributed array.
- You must also issue the detectmdisk command to ensure that the mdisk inventory is updated after deleting mdisks.

CMMVC6530E **The command cannot be initiated because the maximum supported number of arrays already exists.**

Explanation

The cluster already has the maximum number of arrays that it can support. The command attempted to add a new array.

User response

Remove an array that is no longer needed, and resubmit the command.

CMMVC6532E **The command cannot be initiated because there is insufficient free memory that is available to the I/O group.**

Explanation

This command requires that there is sufficient free memory available for the specified I/O group to allocate the memory that is required for the new array.

User response

Ensure that there is sufficient memory available to the I/O group, and resubmit the command. You can increase the amount of memory that is allocated to the I/O group. You can also reduce the amount of memory that is used by reducing the number volume mirrors or Copy Services relationships in the I/O group.

CMMVC6533E **The command cannot be initiated because the specified array member does not exist in the selected array.**

Explanation

This command requires that the array member that you specify is an LDisk. It is possible that the array member that you specified was an LDisk that was recently deconfigured due to an error. You can use the lsarraymember command to display the available members of an array.

User response

Select an array member that has an associated LDisk, and resubmit the command.

CMMVC6534E **The command cannot be initiated because the drive that you have specified does not exist.**

Explanation

You have specified a drive ID that is not defined.

User response

Use the lsdrive command to display existing drive IDs. Specify only existing drive IDs, and resubmit the command.

CMMVC6535E **The command cannot be initiated because you have specified an incorrect number of drives to configure an array using the RAID geometry that you have specified.**

Explanation

Each RAID geometry requires a minimum number of available drives in order to configure an array using that geometry. For example, a RAID 6 geometry requires that you specify at least four available drives. The number of drives that you have specified is less than the minimum number of drives that are required for the RAID geometry that you have specified.

User response

Ensure that you specify a sufficient number of drives to accommodate the RAID geometry that you specify, and resubmit the command. You might want to specify a different number of drives or a different RAID geometry.

CMMVC6536E **The command cannot be initiated because you have specified more drives than the specified RAID geometry permits.**

Explanation

The number of drives that you specify must be within the supported range of number of drives that is supported for the RAID geometry that you specify. For example, a RAID 1 geometry requires that you specify exactly two available drives.

User response

Specify a number of available drives that is supported for the RAID geometry that you specify, and resubmit the command.

CMMVC6537E **The command cannot be initiated because the drive that you have specified has a Use property that is not supported for the task.**

Explanation

The drive has a Use that is not unused.

User response:
Choose a different drive, or change this drive's use to unused.

CMMVC6538E **The command cannot be initiated because at least one of the drives that you have specified has a Use property that is not Candidate.**

Explanation

Every drive that you specify for this command must have a Use property of Candidate. You can submit

the lsdrive command to display the Use property of existing drives.

User response

Ensure that all of the drives that you specify have a Use property of Candidate, and resubmit the command.

CMMVC6539E **The command cannot be initiated because the array does not have sufficient redundancy.**

Explanation

The array must have sufficient redundancy when you submit this command. The task that you have requested would have taken the array offline.

User response

Fix all errors that are related to the array that you have specified and restore redundancy to the array before you resubmit the command.

CMMVC6540E **The task cannot be initiated because the thin-provisioned grain size is too small to accommodate the provisioned capacity that you have requested for the volume.**

Explanation

The provisioned capacity that you requested requires a larger number of grains than the supported maximum for the specified grain size.

User response

Either increase the grain size, decrease the requested provisioned capacity of the volume, or both, and then resubmit the task.

CMMVC6541E **The task cannot be initiated because the provisioned capacity that you have requested for the volume is larger than the maximum capacity that is supported for the extent size.**

Explanation

The extent size of the storage pool that you selected would require a larger number of extents than the supported maximum to accommodate the provisioned capacity that you requested for the volume.

User response

Either select a different storage pool that has an extent size that is large enough to accommodate the requested provisioned capacity or specify a provisioned capacity that is supported for the extent size of the storage pool that you selected, and resubmit the task.

CMMVC6542E **The remote authentication task has failed.**

Explanation

An error has occurred while attempting to authenticate a user account using a remote authentication service. You can run the svc_snap task to gather cluster information that can be used in problem determination.

User response

Contact IBM technical support for assistance.

CMMVC6543E **The task cannot be initiated because you can only specify a direct-attached managed drive when you submit the task.**

Explanation

The drive that you have specified either is not managed or is not a local drive.

User response

Specify a direct-attached MDisk when you submit this task.

CMMVC6544E **The task cannot be initiated at this time because the direct-attached managed drive that you have specified is too busy. Resubmit the task when the drive is less busy.**

Explanation

The task takes approximately thirty seconds to complete. When the direct-attached managed drive is busy, the time that is required to complete the task increases. When the drive is too busy, the task cannot complete in a reasonable amount of time.

User response

Resubmit the task when the direct-attached managed drive is less busy.

CMMVC6545E **The Apply Drive Software task has failed to access the software download image.**

Explanation

Either the image file cannot be read, the validation signature is incorrect, the drive type or software type is not correct, or the image file is corrupted.

User response

Reinstall the software download image, and resubmit the task. If the problem persists, contact IBM technical support for assistance.

CMMVC6546E **A drive error was detected during the Apply Drive Software task.**

Explanation

At least one of the drives that was requested for update by the Apply Drive Software task has failed.

User response

Either remove the failed drives from the list of drives for the Apply Drive Software task and resubmit the task, or perform problem determination on the failed drives.

CMMVC6547W **The Download FPGA firmware task has been initiated. The MDisk remains Offline while the task is in progress. Do not remove power from the drive or node while the task in is progress.**

Explanation

The task might take approximately fifteen minutes to complete. When the task completes, the drive status changes to Online automatically.

User response

Ensure that electrical power is continuously being supplied to the node and the drive, at least until the task completes and the drive status changes to Online.

CMMVC6548E **The FPGA firmware cannot be applied because the drive has a use other than candidate.**

Explanation

Updating a drive FPGA level is not guaranteed to maintain data integrity, therefore the drive must not be part of an array. To ensure this, the drive must have a use of "candidate" before the package can be applied.

User response

If the drive is currently in the "failed" state, run through all maintenance actions required for the drive before continuing. If the drive is a spare or is unused, the drive use can be changed through the GUI or through the **chdrive** command. If a drive is currently part of an array, a hot-spare drive must be configured and the drive use changed to "failed" before changing the use to "candidate".

CMMVC6549E **The Authentication task has failed because the authentication service URL that you have specified is not a valid URL.**

Explanation

This error might be caused by the authentication service not operating correctly or by an incorrect URL being defined for the authentication service. You can use the chauthservice command to change the URL that is defined in the cluster for the authentication service.

User response

Ensure that the authentication service is operating correctly. Ensure that the authentication service URL that is defined in the cluster is correct, and resubmit the task.

CMMVC6550E **The Authentication task has failed because the network address that is specified in the authentication service URL cannot be resolved.**

Explanation

The authentication service URL that is defined in the cluster has a network address that cannot be resolved. You can use the chauthservice command to change the URL that is defined in the cluster for the authentication service.

User response

Ensure that the authentication service is operating correctly. Ensure that the authentication service URL that is defined in the cluster is correct. Ensure that the network connection between the cluster and the authentication service is operating correctly, and resubmit the task.

CMMVC6551E **The Authentication task has failed because the combination**

of user name and password that is defined in the cluster for authorization by the authentication service is not defined on the authentication service.

Explanation

The authentication service has refused an authentication request from the cluster. You can use the chauthservice command to change the user name or the password that is defined in the cluster for the authentication service.

User response

Ensure that the user name and password combination that is defined in the cluster for the authentication service is also defined on the authentication service, and resubmit the task.

CMMVC6552E **The Authentication task has failed because an SSL connection could not be established with the authentication service.**

Explanation

This error might be caused by an incorrect SSL configuration on the authentication service server or by a rejection by the authentication service server of the SSL certificate that is configured on the cluster. You can use the chauthservice command to set the SSL certificate that is defined in the cluster for the authentication service server.

User response

Ensure that the SSL configuration on the authentication service server is correct and that the SSL certificate that is defined in the cluster for the authentication service server is correct, and resubmit the task.

CMMVC6553E **The task cannot be initiated because at least one quorum disk is not in the correct state.**

Explanation

All of the quorum disks must have a state of Online when you set an MDisk to be the active quorum disk.

User response

Ensure that all of the quorum disks have a state of Online, and resubmit the task.

CMMVC6554E **The Authentication task has failed because the user name that was received from the authentication service is not a valid cluster user name.**

Explanation

The cluster user name cannot exceed 256 characters in length, and cannot contain any of the following characters:

- colon :
- percent sign %
- comma ,
- double quote""
- single quote '

User response

Change the definition of the user name in the remote authentication service so that it conforms to the cluster user name requirements, and resubmit the task.

CMMVC6555E **The Authentication task has failed because the authentication service either sent an incorrect response, or it sent a response that indicates that the authentication request has failed for a reason other than incorrect authentication credentials.**

Explanation

Either the format of the response from the authentication service is not valid or the response indicates a failure to authenticate that is not related to the credentials that were being authenticated.

User response

Ensure that the authentication service is functioning correctly, and resubmit the task. If the problem persists, contact the authentication service technical support for assistance.

CMMVC6556E **The task cannot be initiated because an error has occurred while attempting to read a file.**

Explanation

The task specified the name of a file on the file system of the cluster configuration node. The specified file cannot be opened. This error might be caused by a typographical error in the file name that you specified

or by a failover of the configuration node to a different node than the node into which you are currently logged in.

User response

Ensure that the file has been copied to the current configuration node and that you are logged in to that node, specify the correct file name, and resubmit the task.

CMMVC6557E The task cannot be initiated because the file that you have specified is too large.

Explanation

The task specified the name of a file on the file system of the cluster configuration node. The specified file cannot be used because it exceeds the maximum size supported for the task. If the file has been corrupted, you can copy the correct version of the file onto the configuration node to restore the correct file size. The maximum file size is described in the task help.

User response

Specify the correct file name and ensure that the size of the file does not exceed the supported maximum file size for this task, and resubmit the task.

CMMVC6558E The command cannot be initiated because it might cause volumes to go Offline. Refer to the cluster Command-Line Interface (CLI) command help for this command.

Explanation

You are being warned that this command might result in taking volumes Offline. After you completely understand the possible consequences by reading the command help, you can override the safety precautions and avoid this message by using the -force flag.

User response

1. Submit the lsnode dependantvdisks command to determine which volumes will go Offline if you resubmit this command using the -force flag. If you received this message when you submitted the applysoftware command, you must submit the lsnode dependantvdisks command for every node in the cluster; for all other commands you must submit the lsnode dependantvdisks command for the node that you specified as a parameter in the command that generated this message.
2. This step is required because it is critically important that you understand the implications of using the -force flag for the specific command that you have submitted: Refer to the CLI command help to determine what safety precautions are bypassed when you use the -force flag. The ignored precautions differ, depending on the command.
3. If you want to bypass the safety precautions when you resubmit the command, you must use the -force flag.

CMMVC6559E The Add or Change email user command has failed because you have specified a user type of 'support' and you have specified either the -warning or -info parameter value as 'on'.

Explanation

The user type 'support' is intended to be used to indicate that the user is from a hardware maintenance support service external to your organization. Therefore, only events with the more serious notification type of 'error' can be sent to a 'support' user type.

User response

Ensure that you have specified the correct user type. If you want this user to receive warning or information notifications, do not specify the '-usertype support' parameter and value. If you specify the user type as 'support', you must specify the -warning and -info parameters as 'off'.

CMMVC6560E The command has failed because the specified IP address is already in use by the cluster.

Explanation

You cannot specify an IP address that is already configured to be used by the cluster.

User response

Ensure that the IP address that you specify is not already configured for use by the cluster, and resubmit the task.

CMMVC6561E The set quorum active task has failed because either another set quorum active task is in progress or the selected disk cannot be set as the active quorum disk.

Explanation

This is a multi-step task and can take from a few seconds to several minutes to complete. Only one set quorum active task can be in progress at any specified time. This error has one of two causes. Either another set quorum task is already in progress, or the internal cluster logic did not accept your request to make the selected disk the active quorum disk.

User response

Check the state of the MDisks and complete any outstanding fix procedures. If another set quorum active task might be in progress, wait for several minutes for that task to complete, and resubmit this task. If you have received this error when there is no other set quorum active task in progress, specify a different disk to replace the current active quorum disk and specify the same quorum index number, and resubmit this task.

CMMVC6562E **The requested size exceeds the maximum supported value.**

Explanation

You have submitted a command that has a size parameter and an associated unit option that has a default value of Megabytes (MB, 2e20 bytes) when the -unit option is not specified. The value that you have specified for the size parameter in combination with the specified or default unit value is greater than the maximum supported size of (2e64 - 1) bytes.

User response

Ensure that the size that you specify is correct for the value of the unit option that is defaulted or specified, and that the size is not greater than the maximum supported size, and resubmit the task.

CMMVC6563E **The command has failed because a user that you have specified is not configured to receive email notifications.**

Explanation

All of the users that you specify as a target recipient in the **testemail** command must already have at least one of the following email notification flags set to 'on': -error, -warning, or -info.

User response

Ensure that all of the users that you specify have at least one email notification flag set to 'on', and resubmit the command.

CMMVC6564E **You cannot make this user a remote user because the password type is not valid for a remote user.**

Explanation

The remote authentication server has requirements that do not accept legacy type passwords. This user has a legacy type password.

User response

Either specify a new password and resubmit the command, or first modify the password and then resubmit the command to designate remote authentication for this user.

CMMVC6565E **The command has failed because the specified node is not online.**

Explanation

The command requires that the status of the node that you specify is Online.

User response

Ensure that the node that you specify has a status of Online when you submit this command.

CMMVC6566E **The command cannot be submitted because specifying the -failover parameter requires that you also specify either the -name, -iscsialias or -noiscsialias parameter.**

Explanation

You have not specified the required failover data that is required when you specify the -failover parameter.

User response

Ensure that you want to specify the -failover parameter. When you specify the -failover parameter with this command, ensure that you also specify either the -name, -iscsialias or -noiscsialias parameter.

CMMVC6567E **The Apply Drive Software task cannot be initiated because no download images were found in the package file.**

Explanation

The drive software update package file was unpacked but no download software images were found in the package.

User response

Acquire a valid flash drive software update package file, and resubmit the task using the new package file.

CMMVC6568E **The Apply Drive Software task cannot be initiated because for at least one of the specified drives, the specified file does not contain an image for that drive's technology.**

Explanation

The package file documentation lists the drive types for which there are images.

User response

Acquire a valid drive software update package file that contains an image for this drive type, and resubmit the task using the new package file.

CMMVC6569E **The Apply Drive Software task has failed because no download images were found in the package file of this software type.**

Explanation

The package file documentation lists the drive types and software types for which there are images. The value of the -type parameter that you enter for software type is case-sensitive.

User response

Ensure that the value that you enter for the -type parameter exactly matches the software type that is contained in the flash drive software update package file, and resubmit the command.

CMMVC6570E **The command was not initiated because the cache mode of the volume is already in the state that you had requested.**

Explanation

You have issued a change volume cache mode command but requested the current mode, so there would not have been a change. Therefore, the command was ignored.

User response

List the volume properties to determine the current cache mode. If you want to change the cache mode, ensure that you specify a cache mode that is different from the current cache mode, and resubmit the command.

CMMVC6571E **The command has failed because the I/O group that manages the volume that you specified was offline when you submitted the command. You can use the -force flag to force the operation, which might result in the loss of cache data.**

Explanation

If you submit this command without the -force flag, the I/O group that manages the volume that you specify must have a state of Online.

Note: Use of the -force flag when you change the cache mode might result in loss of the cache data for the volume, depending on the current cache mode and requested cache mode. One example of a risk of potential loss of cache data would be changing the cache mode from readwrite to none.

User response

Either follow service procedures to bring the I/O group online or specify the -force flag to force the change of the cache mode of the volume, and resubmit the task.

CMMVC6572E **The command has failed because the I/O group that manages the volume that you specified is not stable.**

Explanation

The unstable I/O group condition is typically transient, and usually occurs during I/O group failover or fail back processing.

User response

Use the **movevdisk** command to change the I/O group with which this volume is associated to recovery I/O group, and resubmit the command.

CMMVC6573E **The command has failed because the volume that you specified is a source or target of a FlashCopy mapping that is in the prepared state.**

Explanation

If the volume is the source or target of a FlashCopy mapping, the FlashCopy mapping must be in the idle_copied state or the stopped state when you change the cache mode of the volume.

User response

Either remove or stop the FlashCopy mapping and wait for the FlashCopy mapping state to become idle_copied or stopped, and resubmit the command.

CMMVC6574E **The command has failed because the volume that you specified is a source or target of a FlashCopy mapping that is in the suspended state.**

Explanation

If the volume is the source or target of a FlashCopy mapping, the FlashCopy mapping must be in the idle_copied state or the stopped state when you change the cache mode of the volume.

User response

Either remove or stop the FlashCopy mapping and wait for the FlashCopy mapping state to become idle_copied or stopped, and resubmit the command.

CMMVC6575E **The command has failed because the volume that you specified is a source or target of a FlashCopy mapping that is in the preparing state.**

Explanation

If the volume is the source or target of a FlashCopy mapping, the FlashCopy mapping must be in the idle_copied state or the stopped state when you change the cache mode of the volume.

User response

Either remove or stop the FlashCopy mapping and wait for the FlashCopy mapping state to become idle_copied or stopped, and resubmit the command.

CMMVC6576E **The command has failed because the volume that you specified is a source or target of a FlashCopy mapping that is in the stopping state.**

Explanation

If the volume is the source or target of a FlashCopy mapping, the FlashCopy mapping must be in the idle_copied state or the stopped state when you change the cache mode of the volume.

User response

Either remove or stop the FlashCopy mapping and wait for the FlashCopy mapping state to become idle_copied or stopped, and resubmit the command.

CMMVC6577E **The command has failed because the volume that you specified is a source or target of a FlashCopy mapping that is in the copying state.**

Explanation

If the volume is the source or target of a FlashCopy mapping, the FlashCopy mapping must be in the idle_copied state or the stopped state when you change the cache mode of the volume.

User response

Either remove or stop the FlashCopy mapping and wait for the FlashCopy mapping state to become idle_copied or stopped, and resubmit the command.

CMMVC6578E **The command has failed because the iSCSI name is already assigned or is not valid.**

Explanation

The cluster does not support duplicate iSCSI names. A valid iSCSI name cannot contain a comma or leading or trailing spaces.

User response

Ensure that you specify a unique and valid iSCSI name, and resubmit the command.

CMMVC6579E **The command cannot be initiated because the cluster Ethernet port 1 must always be fully configured in either the IPv4 or IPv6 format.**

Explanation

This error can be caused by an attempt to delete the only address that is configured on the primary Ethernet port on the cluster.

User response

When you delete an IP address on the primary Ethernet port, ensure that the other supported IP format is already configured on that port.

CMMVC6580E **The command cannot be initiated because the iSCSI alias that you**

specified contained either leading or trailing space characters.

Explanation

The space character cannot be the starting or ending character of an iSCSI alias name.

User response

Ensure that the iSCSI alias that you specify does not being or end with a space character, and resubmit the command.

CMMVC6581E **The command has failed because the maximum number of allowed iSCSI qualified names (IQNs) has been reached, or the IQN is already assigned or is not valid.**

Explanation

IQNs cannot exceed the maximum number allowed, must not be duplicated, must not contain a comma, and must not contain leading or trailing spaces.

User response

If the number of IQNs is within the allowed maximum, ensure that you specify a unique and valid IQN, and resubmit the command.

CMMVC6582E **The task has failed because the host that you specified is not mapped to an I/O group.**

Explanation

You cannot add a port to a host until you map the host to at least one I/O group.

User response

Map the host to at least one I/O group, and then resubmit the command.

CMMVC6583E **The command has failed because the name that you specified contains a character that is not supported for a node or cluster name.**

Explanation

A node or cluster name cannot contain any of the following characters or ASCII hexadecimal values:

- 0000-001F ASCII control characters
- 0020-002C The space character !""# $ % the ampersand character ' ( ) * + ,
- 002F /
- 003B-0040 ; the 'less than' character = > ? @
- 005B-0060 [\ ] ^ _ `
- 007B-007F { | } ~ and the DEL character

User response

Specify a valid name, and resubmit the command.

CMMVC6584E **The command cannot be initiated because it would unconfigure the remote authentication service while the service is enabled.**

Explanation

You cannot unconfigure the remote authentication service while it is enabled.

User response

Ensure that the remote authentication service is not being used, disable the service, and resubmit the task.

CMMVC6585E **The command cannot be initiated because the array that you have specified has a geometry of RAID 0, which is not a redundant geometry.**

Explanation

The array that you specify for this command must have a redundant geometry, and RAID 0 is not a redundant geometry.

User response

Ensure that you specify an array that has a redundant geometry when you submit the command.

CMMVC6586E **The command cannot be initiated because the action would cause array data loss due to the unsynchronized state of the array.**

Explanation

To avoid data loss, this command is not permitted to process an array that is not synchronized.

User response

Use the lsarraysyncprogress command to ensure that the synchronization process completes for this array, and resubmit the task.

CMMVC6587E **The command did not complete because I/O for the array was not quiesced within the allotted time period.**

Explanation

All outstanding I/O for the array must complete before the configuration can be changed. The command has failed because there is still outstanding I/O to be processed for the array, and the maximum amount of time allotted to the command has expired.

User response

Resubmit the task.

CMMVC6588E **The command cannot be initiated because a drive that you have specified has a capacity that is smaller than the minimum capacity that is required for the array that you have specified.**

Explanation

You can use the lsarraymembergoals command to identify the capacity requirement for a member of the array that you specified.

User response

Specify a drive that has sufficient capacity for the array that you specify when you submit the command.

CMMVC6589E **The command was not initiated because the drive that you have specified does not sufficiently match the array member goals and you did not specify the -balanced parameter.**

Explanation

If you do not specify the -balanced parameter, the new drive must be an exact match for the array member goals when you exchange a new drive for an existing array member. The new drive that you have specified does not match the goals. If you want to use the drive that you specified to replace an existing member of the array that you specified, you must specify the -balanced parameter, which forces the array member goals to be changed to accommodate the new drive.

User response

Either select a different drive that matches the array member goals or specify the -balanced parameter to force a change in the array member goals to accommodate the new drive, and resubmit the command.

CMMVC6590E **The command cannot be initiated because you did not specify the -allowdegraded parameter and the associated array member has insufficient spare protection.**

Explanation

This command requires that spare drives are available to assume the function of array member drives that are removed from the array. The requirement can be bypassed by using the -allowdegraded parameter.

User response

Either configure sufficient additional spare drives or specify the -allowdegraded parameter, and resubmit the command.

CMMVC6591E **The command cannot be initiated because the specified sequence number does not match the sequence number of any events in the event log.**

Explanation

The sequence number that is specified in the command must be identical to the sequence number of an event in the event log.

User response

Check the event log to verify the sequence number of the event that you want to specify, and resubmit the command using the correct sequence number.

CMMVC6592E **The command cannot be initiated because at least one parameter that was specified is not supported when submitting a command to view the details of an event log entry.**

Explanation

Filtering parameters such as '-order severity' or '-status alert' that are valid when listing multiple event log entries are not supported for the command to view the details of a single event log entry.

User response

Check the command syntax, and use supported syntax when you submit the command.

CMMVC6593E **The command cannot be initiated because the event log entry has a status that is not supported for the command.**

Explanation

Only events with a status of 'alert' or 'message' can be manually marked as fixed or unfixed. Events with a status of 'monitoring' or 'expired' are not required to be marked as fixed or unfixed.

User response

Check the event log to verify the sequence number of the event that you want to specify. Ensure that the event that you specify has a status that is supported for the command when you submit the command.

CMMVC6594E **The command cannot be initiated because a drive was specified twice in the drive list.**

Explanation

The drive list cannot contain any duplicate entries because the same drive cannot be a member of an array more than once.

User response

Ensure that the drive list that you specify does not contain any duplicate entries when you submit this task.

CMMVC6595E **The command cannot be initiated because a drive that you have specified has a technology type that is not supported for the command.**

Explanation

The command supports only certain drive technology types. You have specified at least one drive that has a technology type that is not supported for the command.

User response

Consult the command documentation to determine which drive technology types are supported for the command. Submit the lsdrive command to determine which drives are available. Specify an available drive that has a technology type that is supported for the command when you submit the command.

CMMVC6596E **The command has failed because you have specified an I/O group that does not exist.**

Explanation

You must specify an existing I/O group when you submit this command.

User response

Specify an existing I/O group, and resubmit the command.

CMMVC6597E **The command has failed because the email settings are not configured.**

Explanation

The cluster email system settings must be configured before you can submit a command for error notifications.

User response

Configure the cluster email system settings to enable error notifications and resubmit the command.

CMMVC6608E **The command cannot be initiated because Easy Tier is active on the volume copy.**

Explanation

Easy Tier is active on the volume copy, which prevents the success of the command.

User response

Disable Easy Tier on the volume copy or on the storage pool in which the volume copy resides, and resubmit the command.

CMMVC6609E **The command cannot be initiated because the size of the Mdisk is smaller than the extent size of the storage pool.**

Explanation

The sizing of the Mdisk in relation to the storage pool is not correct, which prevents the success of the command.

User response

Use a larger Mdisk or make the extent size of the storage pool smaller than the Mdisk, and resubmit the command.

CMMVC6610E **The update cannot start because one or more I/O groups are in maintenance mode.**

Explanation

Maintenance mode is used during system servicing, which prevents updates.

User response

Complete system servicing, turn off maintenance mode, and resubmit the command.

CMMVC6611E **The command has failed because the specified enclosure is offline.**

Explanation

The specified enclosure is offline, which has prevented the success of the command.

User response

Fix any errors associated with the specified enclosure, and resubmit the command.

CMMVC6612E **The command has failed because of a hardware error.**

Explanation

A hardware error has occurred, which has prevented the success of the command.

User response

Fix any errors in the specified object, and resubmit the command.

CMMVC6613E **The command has failed because the specified enclosure type is not supported.**

Explanation

You have attempted to use an enclosure of an unsupported type.

User response

Do not attempt to use the specified enclosure type.

CMMVC6614E **The command has failed because the specified canister is offline.**

Explanation

The specified canister is offline, which has prevented the success of the command.

User response

Fix any errors associated with the specified canister, and resubmit the command.

CMMVC6615E **The command cannot be initiated because nodes from this enclosure cannot be added to the specified I/O group, or another enclosure is in the process of being added.**

Explanation

The nodes in the enclosure being added are used elsewhere in the cluster, the target I/O group contains nodes of a different control enclosure, or another enclosure has not yet completed the process of being added.

User response

If a node already exists in the I/O group, add nodes from the same enclosure only. If the I/O group is empty, you can use a different control enclosure in which the nodes are not in a cluster. If you are currently adding another enclosure, wait for the completion of that process. Ensure that both nodes of an added enclosure are online and that the enclosure is listed in the output of the lsenclosure command.

CMMVC6616E **All available quorum disks are dependent on the MDisks that you have specified.**

Explanation

The list of MDisks that you have specified contains all activated quorum disks. If all of the MDisks in the list were to become inaccessible, the system would be unable to backup important data. Operating the system without any online quorum disks is not recommended.

User response

Move one or more quorum disks to MDisks that will remain online.

CMMVC6617E **All available quorum disks are dependent on the drives that you have specified.**

Explanation

The list of drives that you have specified contains all activated quorum disks. If all of the drives in the list were to become inaccessible, the system would be unable to backup important data. Operating the system without any online quorum disks is not recommended.

User response

Move one or more quorum disks to drives that will remain online.

CMMVC6618E **All available quorum disks are dependent on the enclosure that you have specified.**

Explanation

Before removing the enclosure that you have specified, the system must be configured so at least one of the drives that are allocated to hold quorum will remain online when the enclosure goes offline.

User response

Assign one or more drives in the control enclosure as a quorum drive. After you have configured the quorum drives, test for dependencies.

CMMVC6619E **All available quorum disks are dependent on the canister that you have specified.**

Explanation

Before removing the canister that you have specified, the system must be configured so at least one of the drives that are allocated to hold quorum will remain online when the canister goes offline.

User response

Assign one or more of the drives in the control enclosure as a quorum drive. After you have configured the quorum drives, test for dependencies.

CMMVC6620E **The command cannot be initiated because the drives that you have specified are in different I/O groups.**

Explanation

All of the specified drives that comprise the array must be in the same I/O group.

User response

Specify one or more drives in the same I/O group, and resubmit the command.

CMMVC6621E **The command cannot be initiated because the array member that you have specified already exists.**

Explanation

A drive has already been configured for the specified array member. You can use the lsarraymember command to display the available members of an array.

User response

Specify an array member without a corresponding drive, and resubmit the command.

CMMVC6622E **The command cannot be initiated because the drive has failed validation tests.**

Explanation

When a drive is made a candidate, the new drive is validated to ensure that adding it to the configuration will not adversely affect the existing or future array status. Either the current status of the drive has not allowed the validation to be performed, or the validation has failed.

User response

Fix any errors associated with the specified drive, or specify a different drive, and resubmit the command.

CMMVC6623E **The command cannot be initiated because the drive validation test has timed out.**

Explanation

When a drive is made a candidate, the new drive is validated to ensure that adding it to the configuration will not adversely affect the existing or future array status. The test timed out, which caused the validation to fail.

User response

Fix any errors that are associated with the specified drive, or specify a different drive, and resubmit the command.

CMMVC6624E **The command cannot be initiated because the drive is not in an**

appropriate state to perform the task.

Explanation

The drive that you have specified is offline. A format task is permitted to an offline drive only if the drive has indicated that a format is required and connectivity to the drive is available.

User response

Fix any errors that are associated with the specified drive, or specify a different drive, and resubmit the command.

CMMVC6625E **The command cannot be initiated because a task is in progress on the drive.**

Explanation

A drive can complete only one task at a time. A previous task remains incomplete. You can monitor the progress of the task by using the **lsdriveprogress** command.

User response

Wait for the previous task to complete, and then retry the command.

CMMVC6626E **The task was not initiated because a command was rejected by the drive that you have specified.**

Explanation

When attempting to initiate a task, a sequence of commands is sent to the drive. One or more of these commands was rejected by the drive that you have specified.

User response

Fix any errors that are associated with the enclosure and cabling, and resubmit the command.

CMMVC6627E **The enclosure that you have specified cannot be changed to managed mode because of a SAS configuration problem that is described in the event log.**

Explanation

The status of the enclosure that you have specified will not allow the enclosure to be managed by the cluster.

User response

Ensure that the enclosure is online and cabled correctly, and resubmit the command.

CMMVC6628E **The enclosure that you have specified cannot be changed to unmanaged mode because one or more drives are in use.**

Explanation

The status of the enclosure that you have specified will not allow the enclosure to be unmanaged by the cluster.

User response

Stop using the drives, and resubmit the command.

CMMVC6630E **A drive dump was not created because a command was rejected by the drive that you have specified.**

Explanation

When initiating a drive dump, a sequence of commands is sent to the drive. One or more of these commands was rejected by the drive that you have specified.

User response

Fix any errors associated with the drive, enclosure, and cabling, or specify a different drive, and resubmit the command.

CMMVC6631E **The task was not completed because the drive that you have specified was unavailable.**

Explanation

The drive that you have specified did not have the required connectivity to complete the task.

User response

Fix any errors associated with the drive, or specify a different drive, and resubmit the command.

CMMVC6917E **Authorization failed.**

Explanation:
This error occurs when an attempt is made to log in to the system by using invalid credentials.

User response:

Ensure that the correct username is specified with the correct password or SSH key. If the behavior continues, contact the security administrator.

CMMVC6953E **The action cannot be completed because volumes are dependent on the specified mdisk. Force is required.**

Explanation

A volume depends on a drive that was specified on the **applydrivesoftware** command.

User response

For systems that support RAID0 arrays, if the drive is a member of a RAID0 array consider whether to introduce more redundancy to protect the data on that drive. If the drive is not a member of a RAID0 array, fix any errors in the event log that relate to the array. When the drive is a member of an array with sufficient redundancy, repeat the command. Alternatively, consider running the command and specifying the `-force` option.

Warning: With any drive software update, there is a risk that the drive might become unusable. Use the **-force** option only if you accept this risk.

CMMVC6972E **The command cannot be initiated because the maximum number of extents for an MDisk would be exceeded.**

Explanation:
Each MDisk has a limited number of extents, which varies according to the extent size set by the mdiskgrp, this limit would be exceeded by this command.

User response

Create a different pool with a larger extent size by using mkmdiskgrp. Then, retry the command by using mdiskgrp.

Ensure that you are familiar with the maximum configurations of the system. Search for the term "configuration limits and restrictions" in the Search support and downloads search box at the http://www.ibm.com/support website. The relationship between extent size and the maximum MDisk capacity is shown in the Extents table.

CMMVC6974E **The command cannot be initiated because there are too many bad blocks on the MDisk.**

Explanation

This error can occur when a managed disk has close to the maximum number of bad blocks, and might reach the maximum number of bad blocks during an expansion process. If the maximum number of bad blocks is reached, the expansion is suspended.

User response

Use the **dumpmdiskbadblocks** and the **lsvdisklba** commands to find bad blocks and rewrite the areas where bad blocks are found. Then, retry the command.

CMMVC6976E **The command cannot be initiated because the MDisk has an array parity rebuild ongoing.**

Explanation

This error can occur when an attempt is made to run a command while an array is not synchronized because data is being rebuilt during a copyback process.

User response

Wait until the data is rebuilt, and then retry the command. You can use the **lsarraysyncprogress** command to view the progress of the synchronization.

CMMVC6977E **The command cannot be initiated because the MDisk has an array initialization ongoing.**

Explanation

This error can occur when an attempt is made to run a command while an array is initializing.

User response

Wait until the array is initialized, and then retry the command. You can use the **lsarrayinitprogress** command to view the progress of the initialization.

CMMVC6978E **The command cannot be initiated because the MDisk has an array member rebuild ongoing.**

Explanation

This error can occur when an attempt is made to run a command while an array member is being rebuilt.

User response

Wait until the array member is rebuilt, and then retry the command. You can use the

lsarraymemberprogress command to view the progress of the recovery process.

CMMVC6979E **The command cannot be initiated because the MDisk has an array member exchange ongoing.**

Explanation

This error can occur when an attempt is made to run a command while an array member is being exchanged (swapped).

User response

Wait until the array member is exchanged, and then retry the command.

CMMVC6988E **The command cannot be initiated because the maximum number of iSCSI qualified names (IQNs) for the cluster has been reached.**

Explanation

The specified cluster is already configured with the maximum number of IQNs.

User response

None.

CMMVC6998E **The maximum number of iSCSI qualified names (IQNs) plus WWPNs for the cluster is already configured.**

Explanation

The command cannot be initiated because the maximum number of iSCSI qualified names (IQNs) plus WWPNs for the cluster has been reached.

User response

Determine whether the action is required.

If the action is required, review the current configuration to determine whether any current iSCSI qualified name or WWPN definitions are not required. Remove at least one iSCSI qualified name or WWPN that is not required, and resubmit the command.

CMMVC6999E **The command cannot be initiated because the maximum number of iSCSI qualified names (IQNs) for the host has been reached.**

Explanation

The specified host is already configured with the maximum number of IQNs.

User response

None.

CMMVC7003E **The command cannot be initiated because the power supply unit (PSU) that you have specified is offline.**

Explanation

The power supply unit (PSU) that you specify must be online when you submit the command.

User response

Fix any errors associated with the specified PSU. Ensure that the PSU is online, and resubmit the command.

CMMVC7004E **The command cannot be initiated because the node canister port that you have specified does not exist.**

Explanation

You have submitted a command and specified a node canister port that does not exist. You can submit the lssasfabric command to show all of the available node canister ports for the associated I/O group.

User response

Specify a node canister port that exists when you submit the command.

CMMVC7005E **The command cannot be initiated because enclosures do not exist for the I/O group that you have specified.**

Explanation

You have submitted a command and specified an I/O group that is not associated with an enclosure. You can submit the lsenclosure command to show all of the existing enclosures and their associated I/O groups.

User response

Specify an I/O group that is associated with an enclosure when you submit the command.

CMMVC7006E **The command failed because no enclosures are viewable through the specified canister port.**

Explanation

There are no enclosures on which the command can be run using the parameters that you have specified.

User response

Use the lssasfabric command to check whether you have specified the correct node canister port ID and I/O group ID for the enclosures. If none of the enclosures on which you intend to run the command appear in this view, ensure that the enclosures are physically attached, not offline, and that the port that you have specified is not excluded on the node canister or the expansion canister for that enclosure.

CMMVC7007E **The task cannot be initiated because the enclosure that contains the specified drive is not supported.**

Explanation

The chdrive -use command specifies a drive for use within a system. However, the enclosure of the specified drive is not a hardware type that is supported for use with the product that you are using.

User response

Ensure that your enclosure is a product supported hardware type.

CMMVC7008E **The task cannot be initiated because the enclosure containing the specified drive is unmanaged and cannot be managed because of an enclosure cabling problem.**

Explanation

The chdrive -use specifies a drive for use within a system. However, the enclosure of the specified drive is not managed and is either cabled incorrectly or has incurred a hardware problem on the SAS network. More drives and enclosures cannot be added to the enclosure of the specified drive until the network problem is resolved.

User response

Check the event log, resolve all unfixed errors, and resubmit the command.

CMMVC7009E **The task cannot be initiated because the maximum number of managed enclosures has been reached.**

Explanation

The number of the managed enclosures that can be attached to the system is limited. This limit has been reached.

User response:
None.

CMMVC7010E **The command cannot be initiated because the MDisk mode is set to Array.**

Explanation

This command requires the selected MDisk to be a SAN MDisk (an MDisk that is not an array made from local drives). The selected MDisk has its mode set to Array.

User response

Use lsmdisk to list the MDisks, and resubmit the command against an MDisk with a mode other than Array.

CMMVC7011E **The array cannot be created because no quorum disks are currently configured.**

Explanation

When creating an array, quorum disks are required to back up metadata for the array. Creating an array while no quorum disks are configured is not permitted. Quorum disks can be assigned to drives in the control enclosure automatically, or manually by using the **chquorum** command.

User response

Manage the control enclosure, and ensure that all drives within the enclosure are online before resubmitting the command.

CMMVC7012E **The command cannot be initiated because the action will cause the array to go offline.**

Explanation

The array member drives are offline, which can result in an MDisk going offline. An offline MDisk can cause other MDisks in a storage pool to go offline.

User response

Ensure that the drives are online, and resubmit the command.

CMMVC7013E **The command cannot be initiated because a drive within the enclosure that you have specified is configured as a spare or array member.**

Explanation

Each drive within the enclosure that you have specified must not be configured as a spare or array member. One or more of the drives within the specified enclosure do not meet this requirement. Use the lsdrive command to see which drives within the enclosure have a drive use of spare or member.

User response

Ensure that no drives within the enclosure are configured as spares or array members, and resubmit the command.

CMMVC7014E **The command cannot be initiated because one or more of the drives are not supported for this RAID level.**

Explanation

Only certain RAID levels are supported in some configurations.

User response

Consult the configuration guide to determine supported RAID levels.

CMMVC7015E **The command cannot be initiated because one or more of the drives are located in the wrong node.**

Explanation

For RAID 0, all of the members must be located in the same node. For RAID 1 or RAID 10, mirrored pairs must be located in different nodes.

User response

Consult the configuration guide to determine which drives to use for the selected RAID level.

CMMVC7016E **Authorization has failed because the private key is not valid for the user name that you have specified.**

Explanation

The private key and user name that you have provided do not match what has been defined on the cluster.

User response

Ensure that the private key is valid for the specified user name, and log in again.

CMMVC7017E **Login has failed because the maximum number of concurrent CLI sessions has been reached.**

Explanation

The cluster supports up to 10 concurrent CLI sessions. The login attempt would have exceeded the supported limit.

User response

Reduce the number of open CLI sessions, and log in again.

CMMVC7018E **The command failed because the requested volume size is too large.**

Explanation

The system has a maximum size for volumes that is currently 256 TB. While creating a new volume or resizing an existing volume, you requested a volume size that exceeds the maximum.

User response

Resubmit the command with a smaller volume size.

CMMVC7019E **The command failed because the volume size is not a multiple of 512 bytes.**

Explanation

Volume capacity must be a complete number of blocks, where one block is 512 bytes. While creating a new volume or resizing an existing volume, you requested a volume size that is an incomplete number of blocks.

User response

Resubmit the command with a valid volume size.

CMMVC7020E **The command failed because the maximum number of volumes for this I/O group already exist.**

Explanation

The system has a limit of volumes per I/O group. A new volume cannot be created in an I/O group that has already reached the limit of volumes.

User response

Choose another I/O group or delete some volumes in this I/O group.

CMMVC7021E **The command failed because the maximum number of volume copies already exist.**

Explanation

The system has a limit on the number of volume copies that can be created. An additional volume copy cannot be created because the limit has been reached.

User response

Delete some existing volume copies and resubmit the command.

CMMVC7022E **The command failed because NTP is active.**

Explanation

You have attempted to manually set the cluster time while the cluster is configured to use NTP (network time protocol) to set its time.

User response

Disable NTP, and resubmit the command. If you are trying to set the time manually because the cluster time is incorrect, check the settings on the NTP server.

CMMVC7023E **The command failed because the requested node name is in use as the failover name of another node.**

Explanation

You have attempted to add a node to a cluster or rename a node that is already in the cluster. The new name that you have requested for the node is not valid because one of the nodes in the cluster has been configured with the requested new name as its failover name.

User response

Either resubmit the command specifying a different node name, or modify the configuration of the node in the cluster to change its matching failover name to a different failover name.

CMMVC7024E **The command failed because the maximum number of file systems already exist.**

Explanation

The maximum number of file systems has been reached. Additional file systems cannot be created.

User response

Remove an unused file system and reissue the command, or extend an existing file system by creating the volume there.

CMMVC7025E **The command failed because the volume is associated with a file system and cannot be removed under your current user role.**

Explanation

You are attempting to remove a volume that is associated with a file system. However, you do not possess the required role for file system actions and volume removal.

User response

Resubmit the task using the remove VDisk command.

CMMVC7026E **The command failed because volumes exist in the file system.**

Explanation

You are attempting to delete a storage pool with which volumes are associated. The storage pool cannot be removed while the associated volumes remain.

User response

Remove the file system volumes, and resubmit the command to remove the storage pool.

CMMVC7027E **The command failed because the requested action is not permitted on a volume that is in a file system.**

Explanation

The volume that you have specified is associated with a file system, which disallows the requested action.

User response

The command cannot be completed on this volume. It will only succeed with a volume that is not associated with a file system.

CMMVC7028E **The task cannot be completed because the FlashCopy target volume that you have specified is in a Metro Mirror or Global Mirror relationship, and the I/O group of the volume is different than that of the proposed FlashCopy mapping.**

Explanation

The FlashCopy map must be in the same I/O group as the target volume because the volume is a component of a remote copy relationship.

User response

Specify the I/O group of the target volume when creating the FlashCopy map.

CMMVC7029E **The task cannot be completed because one or more of the target volumes of the FlashCopy mappings is the primary of a mirroring Metro Mirror or Global Mirror relationship.**

Explanation

The target volume is part of a remote copy relationship that is active.

User response

Either force stop the FlashCopy consistency group or stop any remote copy relationships.

CMMVC7030E **The task cannot be completed because the target volume of the FlashCopy mapping is the primary of a mirroring Metro Mirror or Global Mirror relationship.**

Explanation

The target of the FlashCopy map is a component of an active FlashCopy map.

User response

Either force stop the FlashCopy map or stop the remote copy relationship.

CMMVC7031E **The task cannot be completed because the FlashCopy mapping target volume is a secondary in a Metro Mirror or Global Mirror relationship, or is the primary in an active relationship.**

Explanation

The target volume of the FlashCopy map is part of an active remote copy relationship.

User response

Stop the remote copy relationship.

CMMVC7032E **The task cannot be completed because one or more of the target volumes of the FlashCopy mappings is either a secondary in a Metro Mirror or Global Mirror relationship or the primary of an active relationship.**

Explanation

A target volume of a FlashCopy map in the consistency group is part of an active remote copy relationship.

User response

Stop any remote relationships containing a target volume of a map in the consistency group.

CMMVC7033E **The task failed because the current hardware configuration is not valid.**

Explanation

You have issued the "chnodehw" command to enable new hardware that is faulty, unsupported, or incompletely installed.

User response

Follow service procedures as prompted by the management GUI to correct the hardware configuration. Then reissue the command.

CMMVC7036E **The action failed because quorum is not permitted on the drive that you have specified.**

Explanation

Quorum is only permitted on specific drive types. The drive that you have selected will not support quorum.

User response

Reissue the command specifying a different drive.

CMMVC7037E **The action failed because the drive cannot be found.**

Explanation

You have specified a drive does not appear to exist.

User response

Reissue the command specifying a different drive.

CMMVC7038E **The action failed because the system was unable to initialize the quorum disk.**

Explanation

A sequence of SCSI commands must be sent to a quorum disk before it can become available for use. One of these SCSI commands has failed.

User response

Fix any errors associated with the disk, choose a different resource for quorum, and then reissue the command.

CMMVC7039E **The action failed because the specified drive is not online.**

Explanation

The drive that you have specified is offline possibly as a result of errors.

User response

Fix any errors associated with the drive, or choose a different resource for quorum, then reissue the command.

CMMVC7040E **The action failed because the specified MDisk is not online.**

Explanation

The MDisk that you have specified is offline possibly as a result of errors.

User response

Fix any errors associated with the MDisk, or choose a different resource for quorum, then reissue the command.

CMMVC7041E **The action failed because a better quorum candidate is available for use as quorum and override has not been enabled.**

Explanation

Quorum disks are selected automatically based on a set of selection criteria. The resource that was selected is inferior to an alternative resource.

User response

Choose a different resource for quorum, or refer to the quorum documentation before using the -override parameter.

CMMVC7042E **The action failed because the -override yes parameter was used without a specified drive or MDisk.**

Explanation

The -override yes parameter must specify a drive or MDisk.

User response

Reissue the command with the correct syntax.

CMMVC7043E **The action failed because the required extents could not be allocated.**

Explanation

When an MDisk is specified for quorum, some extents must be allocated for use by the quorum disk. Sufficient extents were not available.

User response

Reissue the command using a different MDisk or migrate data from the MDisk to free up sufficient extents.

CMMVC7044E **The action failed because the specified drive is either degraded or excluded.**

Explanation

The drive that you have specified either contains errors or is in the Excluded state.

User response

Fix any errors associated with the drive, or choose a different resource for quorum, then reissue the command.

CMMVC7045E **The action failed because the specified MDisk is either degraded or excluded.**

Explanation

The MDisk that you have specified either contains errors or is in the Excluded state.

User response

Fix any errors associated with the MDisk, or choose a different resource for quorum, then reissue the command.

CMMVC7046E **The action failed because the -rsize option must be set to auto.**

Explanation

You have run the **mkvdisk** or **addvdiskcopy** command to import a compressed volume (using -compressed and -import). The **-rsize** option must be set with a value of **auto**.

User response

Resubmit the command with **-rsize auto**.

CMMVC7047E **The action failed because the validate parameter is not supported for compressed volumes.**

Explanation

The command repairsevdiskcopy -validate was issued against a compressed volume. Unlike thin-provisioned volumes, compressed volumes do not support a validation function.

User response

None.

CMMVC7048E **The action failed because the compressed volume copies are not all corrupted.**

Explanation

You have issued the repairsevdiskcopy or recovervdisk -copy command against a compressed volume copy that is not marked as corrupt. Unlike thin-provisioned volume copies, the repair process for compressed volume copies can only be run if the system detects that they are corrupted.

User response

The issued command is not required. If the volume is offline, consult the *Troubleshooting Guide* to resolve the problem.

CMMVC7049E **The command failed because volumes are obstructing resources required by the compression function.**

Explanation

Compression could not be enabled because a volume prevented internal resources from being reassigned from cache. A volume is offline or data could not be flushed from the cache quickly enough.

User response

If any volumes are offline, follow service procedures to bring them online before resubmitting the command.

CMMVC7050E **The command failed because at least one node in the I/O group does not support compressed volumes.**

Explanation

An attempt was made to create a compressed volume in an I/O group that contains at least one node that does not meet these requirements.

User response

Resubmit the command with a different I/O group.

CMMVC7051E **The command failed because the I/O group contains compressed volumes. The node being added does not support compressed volumes.**

Explanation

An attempt was made to add a node that does not support compression to an I/O group that already contains at least one compressed volume.

User response

Add the node to a different I/O group or add a different node to the specified I/O group.

CMMVC7052E **The nested group search parameter is not valid for the target LDAP server type.**

Explanation

This error can occur when the specified LDAP server type is predefined to do nested group search.

User response

Ensure that you specify the correct LDAP server type. The following rules apply to the LDAP server type and the **-nestedgroupsearch** parameter:

- If the **-type** parameter is `itds`, the **-nestedgroupsearch** parameter is not allowed.
- If the **-type** parameter is `ad`, **-nestedgroupsearch** can be set only to `client` or `off` because there is no server support.
- If the **-type** parameter is `other`, the **-nestedgroupsearch** parameter can be set to `client`, `server`, or `off`.

After you ensure that the LDAP server type is correct, retry the command.

CMMVC7053E **The task cannot be initiated because the nested group search value (server) is not valid for the target LDAP server type.**

Explanation

The LDAP server type that you have specified only supports client-side nested group search.

User response

Reissue the task specifying client-side nested group search.

CMMVC7054E **The task cannot be initiated because the user name or password of the LDAP administrator was not specified.**

Explanation

The user name and password of the LDAP administrator are not configured on the cluster as required. Once the credentials are configured, the user name and password can be changed separately.

User response:
Reissue the task specifying both the LDAP administrator user name and password.

CMMVC7055E **The task cannot be initiated because the specified IP address, port, and base distinguished name (DN) are already configured on an LDAP server.**

Explanation

The same IP address, port, and base DN exist on more than one LDAP server.

User response:
Reissue the task specifying an different IP address, port, and base DN.

CMMVC7056E **The task cannot be initiated because the number of LDAP servers has reached the supported maximum.**

Explanation

The cluster limits the number of LDAP servers that can be configured, and the limit has been reached. To remove a configured LDAP server, the rmldapserver command can be submitted.

User response:
Remove a configured LDAP server, and resubmit the task.

CMMVC7057E **The task cannot be initiated because the specified LDAP server is the only configured LDAP server.**

Explanation

Removing the specified LDAP server would result in the failure of the remote authentication service.

User response:
Disable the remote authentication service by submitting the chauthservice command, and reissue the task.

CMMVC7058E **The task cannot be initiated because no LDAP server is configured.**

Explanation:
The LDAP remote authentication service cannot be used until at least one LDAP server has been configured. To configure LDAP servers, the mkldapserver command can be submitted.

User response:
Configure a valid LDAP server, and reissue the task.

CMMVC7059E **The task cannot be initiated because some remote users are not configured with an SSH key and password for the specified remote authentication service.**

Explanation:

An SSH key and password are required for all users of the remote authentication service. To identify remote users without an SSH key and password, the lsuser command can be submitted. To configure the user's authentication settings, you can use the chuser command.

User response:
Configure the remote users with an SSH key and password or configure the users as local.

CMMVC7060E The task cannot be initiated because the parameters that you have specified are not valid for the LDAP authentication service.

Explanation:
The authentication service URL, user name, password, and SSL certificate are not configurable for the LDAP authentication service.

User response:
Reissue the task specifying valid parameters.

CMMVC7061E The task cannot be initiated because the LDAP administrator user name that you have specified is not valid.

Explanation

LDAP administrator user names must be a valid Distinguished Name, NT login, or User Principal Name.

- Distinguished Names must be a sequence of attribute=value pairs separated by a comma (,), semicolon (;), or plus sign (+), and include special characters and UTF-8 characters that are appropriately escaped with a backslash (\).
- NT logins are valid for Active Directory only and should be in the format DOMAIN\user. They must not start or end with a period (.) and both DOMAIN and user must exclude characters in the set: \ / : ? " < > |
- UPN logins are valid for Active Directory only and must be in the format user@suffix. Both user and suffix must exclude spaces and the following characters: ( ) < > , ; : \ " [ ] @

User response:
Reissue the task specifying a valid Distinguished Name, NT login, or User Principal Name.

CMMVC7062E The task cannot be initiated because you have specified an LDAP attribute that is not valid.

Explanation:
An LDAP attribute name can contain only alphanumeric characters and hyphens, and the name must begin with a letter.

User response:
Reissue the task specifying a valid LDAP attribute name.

CMMVC7063E The task cannot be initiated because the Distinguished Name that you have specified is not valid.

Explanation:
A Distinguished Name must be a sequence of attribute=value pairs separated by a comma (,), semicolon (;), or plus sign (+) that includes special characters and UTF-8 characters escaped with a backslash (\).

User response:
Reissue the task specifying a valid Distinguished Name.

CMMVC7064E User authentication failed because one or more LDAP servers could not be contacted.

Explanation:
The LDAP server is not operating correctly or an incorrect IP address and port are defined for the LDAP authentication service. The event log has been logged and the corresponding service procedure can be used to resolve this issue. To change the IP address and port of an LDAP server, a security administrator role can submit the chldapserver command.

User response:
Ensure that the LDAP servers are operating correctly. Ensure that the IP address and port defined for each LDAP server is correct, and reissue the task.

CMMVC7065E User authentication failed because a timeout occurred while communicating with one or more LDAP servers.

Explanation:
A timeout occurred while the cluster was attempting to contact the LDAP servers. This timeout might be caused by a TCP/IP network problem, the LDAP servers not operating correctly, or by an incorrect IP address and port being defined for the LDAP servers. The event has been logged and the corresponding service procedure can be used to resolve this issue. To change the IP address and port of an LDAP server, a security administrator can use the chldapserver command.

User response:
Ensure that LDAP servers and the TCP/IP network between them and the cluster are functional. Ensure that the IP address and port defined for each LDAP server is correct, and reissue the task.

CMMVC7066E User authentication failed because an SSL connection could not be

established with one or more LDAP servers.

Explanation:
An incorrect LDAP security configuration exists on the cluster or an SSL certificate on the cluster is not valid. The event was logged and the corresponding service procedure is available to resolve this issue. To turn off transport layer security, a security administrator can submit the chldap command or submit the chldapserver command to set the SSL certificate for an LDAP server.

User response:
Ensure that the SSL configuration on each LDAP server is correct and that the SSL certificate defined in the cluster for each LDAP server is correct, or ensure that Transport Layer Security is disabled. Then reissue the task.

CMMVC7067E **User authentication failed because one or more LDAP servers rejected an anonymous bind attempt.**

Explanation:
A user name and password were not specified on the cluster for LDAP authentication and the LDAP server has refused an attempt to bind anonymously. The event has been logged and the corresponding service procedure can be used to resolve this issue. To configure a user name and the password for LDAP authentication, a security administrator can submit the chldap command.

User response:
Ensure that all LDAP servers are configured to allow anonymous bind, or configure a user name and password for LDAP authentication. Then reissue the task.

CMMVC7068E **User authentication failed because one or more LDAP servers rejected an attempt to bind with the LDAP administrator credentials configured on the cluster.**

Explanation:
A user name and password were configured on the cluster for LDAP authentication and an LDAP server has refused an attempt to bind with these credentials. The event has been logged and the corresponding service procedure can be used to resolve this issue. To change the user name and password defined on the cluster, a security administrator can submit the chldap command.

User response:
Ensure that the LDAP credentials configured on the cluster match the credentials that are configured on all LDAP servers, and reissue the task.

CMMVC7069E **User authentication failed because one or more LDAP servers report an incorrect user name or password.**

Explanation:
The user name and password that you have provided do not match any user name and password on the configured LDAP servers. If the password for the user name has recently changed on the configured LDAP servers, it may be necessary to force the cluster to refresh its authentication cache. To force a refresh, a security administrator can submit the chauthservice -refresh command.

User response:
Ensure that the user name and password are correct. Ensure that any recently changed passwords are flushed from the cache of the cluster, and reissue the task.

CMMVC7070E **User authentication failed because the LDAP user attribute is incorrectly configured on one or more LDAP servers.**

Explanation:
The LDAP configuration on the cluster specifies an LDAP user attribute that does not exist on the LDAP server. Users cannot be identified by user name because the attribute is incorrectly configured. The event has been logged and the corresponding service procedure is available to resolve this issue. To specify a different user attribute, a security administrator can submit the chldap command.

User response:
Ensure that the LDAP user attribute specified on the cluster is correct. Ensure that the schema on the configured LDAP servers includes the specified attribute, and reissue the task.

CMMVC7071E **User authentication failed because the LDAP group attribute is incorrectly configured on one or more LDAP servers.**

Explanation:
The LDAP configuration on the cluster specifies an LDAP group attribute that does not exist on the LDAP server. The groups to which a user belongs cannot be identified because the attribute is incorrectly configured. The event has been logged and the corresponding service procedure can be used to resolve this issue. To specify a different group attribute, a security administrator can submit the chldap command.

User response:
Ensure that the LDAP group attribute specified on the cluster is correct. Ensure that the schema on

the configured LDAP servers includes the specified attribute, and reissue the task.

CMMVC7072E User authentication failed because the LDAP group attribute is not in a valid format on one or more LDAP servers.

Explanation:
The LDAP group attribute in the user entry on the configured LDAP servers is in an invalid format. The event has been logged and the corresponding service procedure can be used to resolve this issue. The attribute must be a multivalued attribute containing the distinguished names of groups, or a colon-separated list of up to eight user group names.

User response:
Ensure that the LDAP group attribute is correctly formatted on the LDAP servers, and reissue the task.

CMMVC7073E User authentication failed because the LDAP audit log attribute is not configured correctly on one or more LDAP servers.

Explanation:
The LDAP configuration on the cluster specifies an LDAP audit log attribute that does not exist on the LDAP server. The string to use in the audit log cannot be identified because this attribute is incorrectly configured. The event has been logged and the corresponding service procedure can be used to resolve this issue. To specify a different audit log attribute, a security administrator can issue the chldap command.

User response:
Ensure that the LDAP audit log attribute is correctly specified on the cluster. Ensure that the schema on the LDAP servers includes the specified attribute, and reissue the task.

CMMVC7074E The task cannot be initiated because the user could not be found on any of the configured LDAP servers.

Explanation:
The remote user is configured but either no entry for the user exists on the configured LDAP servers, or more than one entry was found. The event has been logged and the corresponding service procedure can be used to resolve this issue.

User response:
Ensure that the user name is unique on the LDAP servers. Ensure that the LDAP bind credentials allow the LDAP server to be searched, and reissue the task.

CMMVC7075I The LDAP task completed successfully.

Explanation

The LDAP task completed successfully.

User response

None.

CMMVC7076E *VOLUME* cannot be created with *VALUE* without *VALUE*.

Explanation:
You are attempting to create a thin provisioned file system volume without compression. Thin provisioned file system volumes must include compression.

User response

Create a thin provisioned file system volume with compression or create a file system volume without thin provisioning.

CMMVC7077E The command failed because adding a thin provisioned copy to a file system volume is not allowed.

Explanation:
You are attempting to add a volume copy to a file system volume that is not compressed but is thin provisioned. Only copies with compression or without thin provisioning can be added to file system volumes.

User response

Add a copy either with compression or without thin provisioning to the file system volume.

CMMVC7078E The command cannot be initiated because adding a copy to the storage pool of file system volumes is not allowed.

Explanation:
You are attempting to add a volume copy to a file system volume from a different storage pool. Only copies from the same storage pool can be added to file system volumes.

User response

Add a volume copy to a storage pool within the same file system volume, only.

CMMVC7079E The command failed because a volume copy must be different when added to a file system volume.

Explanation:
You are only allowed to add a different volume copy to perform conversions between uncompressed and compressed.

User response

Add a compressed copy to a file system volume with an uncompressed copy or an uncompressed copy to a file system volume with a compressed copy.

CMMVC7080W **The compressed storage used by the cluster is approaching the capacity that is licensed.**

Explanation:
You are being informed that compressed storage in use by the cluster has nearly reached the total licensed capacity.

User response

Compare your actual and planned usage of compression.

CMMVC7081W **The compressed storage used by the cluster has exceeded the capacity that is licensed.**

Explanation:
You are being informed that compressed storage in use by the clustered system exceeds the total licensed capacity.

User response

Reduce the use of compressed storage or purchase additional licensing.

CMMVC7082W **The number of control enclosures with compressed volumes exceeds the number that are licensed.**

Explanation:
You are being informed that the licensed number of control enclosures that can contain compressed volumes has been exceeded.

User response

Reduce or consolidate compressed volumes or purchase additional licensing.

CMMVC7083E **The specified number of control enclosures is not valid.**

Explanation:
The valid range of values for licensed control enclosures is 0–4. The value that you specify must be within this range.

User response

Specify a value between 0 and 4.

CMMVC7084E **The action failed because the command is not permitted for compressed volumes.**

Explanation:
The command that you have submitted is not valid on compressed volumes.

User response

Do not submit this command for a compressed volume.

CMMVC7102E **The action could not be performed because one or more of the requested secondary volume is the target of an active FlashCopy mapping.**

Explanation:
An auxiliary volume cannot be selected as a change volume of a Remote Copy relationship while currently defined for a different relationship.

User response

Choose a different auxiliary volume.

CMMVC7143E **The command cannot be initiated because nodes from another cluster are visible.**

Explanation:
The system layer can be changed only when no other systems are visible on the fabric.

User response

Change the Fibre Channel SAN zoning to remove connectivity between the nodes in the local system and nodes in remote systems, issue 'svctask detectmdisk', then retry the command.

CMMVC7144E **The command cannot be initiated because it is not supported by this hardware type.**

Explanation:
This feature is not available on this hardware

User response

For more information, refer to the Software feature limits page in the relevant Hardware Guide.

CMMVC7145E **The command cannot be initiated because one or more partnerships are defined.**

Explanation:
The system layer can only be changed when there are no partnerships defined with remote systems.

User response

Remove all partnerships to remote systems, first removing any associated remote copy relationships and consistency groups using those partnerships, then retry the command.

CMMVC7146E **The command cannot be initiated because a host object is associated with SAN Volume Controller ports.**

Explanation:
The system layer cannot be changed when there are host objects that contain Fibre Channel ports from SAN Volume Controller nodessystem.

User response

Remove all host objects that contain node ports, then retry the command.

CMMVC7147E **The command cannot be initiated because one or more MDisks are provided by a storage system that does not support the system layer changing.**

Explanation:
The system layer cannot be changed if MDisks are currently being provided by a system.

User response

For each MDisk provided by a system, remove that MDisk from its storage pool. Once all such MDisks have been removed, change the Fibre Channel SAN zoning to remove connectivity between the nodes in the local system, and the nodes in the remote system. Finally, execute the **detectmdisk** command, and then change the system layer.

CMMVC7154E **The task cannot be completed because the specified FlashCopy mapping is controlled by a Metro Mirror or Global Mirror relationship.**

Explanation:
The Metro Mirror or Global Mirror relationship under which the specified FlashCopy mapping is controlled prevents completion of the task.

User response

Check whether the specified task is permitted under the configuration of the Metro Mirror or Global Mirror relationship.

CMMVC7155E **The Create FlashCopy mapping task cannot be initiated because the source or target volume is being used as a change volume for a Metro Mirror or Global Mirror relationship.**

Explanation:
A volume cannot become the source or target of a FlashCopy mapping while in use as a change volume in a Metro Mirror or Global Mirror relationship.

User response

Specify a volume other than the source or target that is currently in use.

CMMVC7156E **The change volume could not be associated because it is already a source or target volume in an existing FlashCopy mapping.**

Explanation:
A change volume cannot be associated if the same volume is the source or target of a FlashCopy mapping.

User response

Specify a volume other than the source or target that is currently in use.

CMMVC7157E **The change volume could not be associated because the Metro Mirror or Global Mirror relationship has a volume on this cluster that is the target of a FlashCopy mapping in a different I/O group.**

Explanation:
The I/O group of the change volume conflicts with an I/O group in the relationship with which an association was attempted.

User response

Ensure that no conflicting I/O groups exist.

CMMVC7158E **The change volume could not be associated because the Metro Mirror or Global Mirror relationship has a volume on this cluster that is already involved in the maximum number of FlashCopy mappings.**

Explanation:
The change volume cannot cause another volume to exceed the number of FlashCopy mappings allowed.

User response

Reduce the number of FlashCopy mappings in the volume that has reached its maximum.

CMMVC7159E **The change volume could not be associated because the relationship has a volume on this cluster that is in an I/O group with no online nodes or because there are unrecovered FlashCopy mappings in the I/O group.**

Explanation:
FlashCopy metadata was lost and cannot be recovered.

User response

The nodes of the I/O group must be brought online before the change volume can be associated. If the nodes of the I/O group were removed, then delete any FlashCopy mappings that existed prior to removing the nodes. .

CMMVC7160E **The change volume could not be associated because the I/O group has insufficient free bitmap space.**

Explanation:
The I/O group must have additional bitmap space to allow the change volume to be associated.

User response

Increase the total bitmap space of the I/O group.

CMMVC7161E **The change volume could not be associated because the master change volume can only be configured from the master cluster, and the auxiliary change volume from the auxiliary cluster. The change volume must be configured from the remote cluster.**

Explanation:
A change volume must be associated from a cluster of the same type (master or auxiliary).

User response

Configure the change volume from the remote cluster.

CMMVC7162E **The change volume could not be associated because one is already configured for the specified Metro Mirror or Global Mirror relationship.**

Explanation:
A change volume has been previously configured for the specified Metro Mirror or Global Mirror relationship.

User response

Ensure that the change volume is associated where a change volume has not been configured.

CMMVC7163E **The change volume could not be associated because it is already involved in a Metro Mirror or Global Mirror relationship.**

Explanation:
The change volume is currently associated with a Metro Mirror or Global Mirror relationship.

User response

Choose an unassociated change volume for the specified Metro Mirror or Global Mirror relationship.

CMMVC7164E **The change volume could not be associated because its size is different from those in the Metro Mirror or Global Mirror relationship.**

Explanation:
A change volume cannot be associated with volumes of a different size.

User response

Choose a change volume with a size that matches those with which it is being associated.

CMMVC7165E **The change volume could not be disassociated because the Metro Mirror or Global Mirror relationship does not have one configured.**

Explanation:
An attempt was made to disassociate a change volume where one does not currently exist.

User response

Verify whether the intended change volume was specified.

CMMVC7166E **The change volume could not be disassociated because it is currently in use by the Metro Mirror or Global Mirror relationship.**

Explanation:
An attempt was made to disassociate a change volume that is currently in use.

User response

Verify whether the intended change volume was specified.

CMMVC7167E **The change volume could not be associated because it is mapped to a host.**

Explanation:
A change volume cannot be associated if it is mapped to a host.

User response

Unmap the change volume from its host, or choose a different change volume.

CMMVC7168E **The host mapping was not created because the volume is a change volume for a Metro Mirror or Global Mirror relationship.**

Explanation:
A volume cannot be mapped to a host if it is associated as a change volume.

User response

Choose a different change volume.

CMMVC7169E **The Remote Copy relationship could not be deleted because this would corrupt the secondary volume.**

Explanation:
The deletion of the relationship is being prevented as a safeguard against corrupting the secondary. The result can be prevented by allowing resynchronization, or by overriding the safeguard.

User response

Allow the relationship to become synchronized before deleting, or reissue the command with the -force flag to allow corruption of the secondary.

CMMVC7170E **The Remote Copy relationship could not be created because the specified master volume is already a change volume for a different relationship.**

Explanation:
A master volume cannot be selected as a change volume of a Remote Copy relationship while currently defined for a different relationship.

User response

Choose a different master volume.

CMMVC7171E **The Remote Copy relationship could not be created because the specified auxiliary volume is already a change volume for a different relationship.**

Explanation:
An auxiliary volume cannot be selected as a change volume of a Remote Copy relationship while currently defined for a different relationship.

User response

Choose a different auxiliary volume.

CMMVC7172E **Enabling access to the secondary volume of the Remote Copy relationship could not be completed in a reasonable time.**

Explanation:
A timeout occurred before the task could be completed. The relationship is continuing to enable access, and will have a state of idling when access is enabled.

User response

Check the event log for any events to be resolved, and resubmit the task.

CMMVC7173E **Enabling access to the secondary volumes of the Remote Copy consistency group could not be completed in a reasonable time.**

Explanation:
A timeout occurred before the task could be completed. The consistency group is continuing to enable access, and will have a state of idling when access is enabled.

User response

Check the event log for any problems that need to be resolved, and resubmit the task.

CMMVC7174E **The task cannot be completed because the other cluster is running a software version that is not recent enough.**

Explanation:
The software version of the one of the clusters is not supported.

User response

update the software version of the cluster.

CMMVC7175E **Enabling access to the secondary volumes of the Remote Copy consistency group could not be completed because the**

relationships in the group are not mutually consistent.

Explanation:
The relationships in the consistency group must be mutually consistent before access to the secondary volumes can be enabled.

User response

Ensure that the relationships in the Remote Copy consistency group are mutually consistent.

CMMVC7176E **The Remote Copy relationship could not be added to the consistency group because the cycling modes do not match.**

Explanation:
The cycling modes of the Remote Copy relationship and the consistency group to which it is being added must match.

User response

Ensure that the cycling modes match.

CMMVC7177E **The Remote Copy relationship could not be added to the consistency group because the cycling periods do not match.**

Explanation:
The cycling periods of the Remote Copy relationship and the consistency group to which it is being added must match.

User response

Ensure that the cycling periods match.

CMMVC7178E **The Remote Copy relationship could not be started in a reasonable time. It is now stopped.**

Explanation:
A timeout occurred before the task could be completed.

User response

Check the event log for any problems that need to be resolved, and resubmit the task.

CMMVC7179E **The Remote Copy consistency group could not be started in a reasonable time. It is now stopped.**

Explanation:
A timeout occurred before the task could be completed.

User response

Check the event log for any problems that need to be resolved, and resubmit the task.

CMMVC7180E **The Remote Copy relationship could not be started because no master change volume is defined.**

Explanation:
A master change volume must be defined for the Remote Copy relationship.

User response

Define a master change volume.

CMMVC7181E **The Remote Copy relationship could not be started because no auxiliary change VDisk is defined.**

Explanation:
An auxiliary change volume must be defined for the Remote Copy relationship.

User response

Define an auxiliary change volume.

CMMVC7182E **The Remote Copy consistency group could not be started because no master change volume is defined.**

Explanation:
A master change volume must be defined for the Remote Copy consistency group.

User response

Define a master change volume.

CMMVC7183E **The Remote Copy consistency group could not be started because no auxiliary change volume is defined.**

Explanation:
An auxiliary change volume must be defined for the Remote Copy consistency group.

User response

Define an auxiliary change volume.

CMMVC7184E **The task cannot be completed as the Remote Copy object is not stopped.**

Explanation:
The task cannot be completed as the Remote Copy object is not stopped.

User response

Stop the Remote Copy object.

CMMVC7185E **The change volume could not be associated because the Metro Mirror or Global Mirror relationship has a volume on this cluster that is in a different I/O group.**

Explanation:
The I/O group of the change volume conflicts with an I/O group in the relationship with which an association was attempted.

User response

Ensure that no conflicting I/O groups exist.

CMMVC7186E **The Remote Copy relationship was not created because the master VDisk is owned and has restricted use.**

Explanation:
The specified task cannot be performed while the master volume is in a file system or owned.

User response

Choose a different master volume, if the specified volume cannot be removed from the file system.

CMMVC7187E **The Remote Copy relationship was not created because the auxiliary VDisk is owned and has restricted use.**

Explanation:
The specified task cannot be performed while the auxiliary volume is in a file system or owned.

User response

Choose a different auxiliary volume, if the specified volume cannot be removed from the file system.

CMMVC7188E **The command failed because the master volume is in a file system.**

Explanation:
The specified task cannot be performed on a master volume while it is in a file system.

User response

Choose a different master volume, if the specified volume cannot be removed from the file system.

CMMVC7189E **The change volume could not be associated because it is in a file system.**

Explanation:
The specified change volume cannot be associated while it is in a file system.

User response

Choose a different change volume, if the specified volume cannot be removed from the file system.

CMMVC7203E **The command failed because the hardware configuration of the local cluster is not compatible with the code of a partnered cluster.**

Explanation:
The hardware configuration of the local cluster is not compatible with the code of a partnered cluster. See the chnodehw explanation for more information.

User response

Make sure the hardware configuration and code level of all the clusters in a partnership are compatible before you create the partnership. Run chnodehw for diagnostic information.

CMMVC7204E **Cannot remove the last I/O group from the volume access set.**

Explanation

A volume must have at least one I/O group in an access set. All of the I/O groups cannot be removed from a volume access set.

User response

Before you run the command again, complete one of the following steps:

- Modify the I/O group list so that the list does not include all of the I/O groups that provide access to the volume.
- Add I/O groups to the access set.

CMMVC7205E **The command failed because it is not supported.**

Explanation:
The command fails because it is not supported by the product.

User response

Review the documentation for the product to select an appropriate command.

CMMVC7206E The command failed because a parameter is not supported.

Explanation:
The user entered a parameter which is not supported by the product they are using.

User response

Review the documentation and select appropriate parameters for the product.

CMMVC7210E The command failed because it would create too many compressed volume copies that are contained within regular pools in the I/O group.

Explanation

The number of compressed copies within regular pools in an I/O group is limited to 200 or 512, depending on the platform type. This limitation does not apply to compressed copies in data reduction pools.

User response

Complete one of the following actions:

- Delete compressed copies from regular pools in the I/O group, then retry the command.
- Retry the command and specify a data reduction pool for the compressed copies.
- Retry the command for a different I/O group.

CMMVC7211E The command failed because it is not supported for image-mode MDisks.

Explanation

This error is returned by the remove mdisk **rmmdisk** command when issued against an MDisk that is backing an image-mode volume on a platform that does not support migration for image-mode volumes.

User response

If the image-mode volume is not required then use rmvdisk to delete the volume. This deletes the MDisk as well. If the user wants to migrate the image-mode volume's data to internal storage, do that using Volume Mirroring and then delete the image-mode volume copy.

CMMVC7218E The task cannot be initiated because an invalid license key was specified.

Explanation:
The license key specified has not been recognized as a valid key.

User response

Check for typing errors and try again.

CMMVC7219E The task cannot be initiated because an invalid function id was specified.

Explanation:
The function id specified in the command is invalid. There are a set number of features that can be activated. This message is shown when the user specifies and invalid function id.

User response

Use lsfeature to check which function id to specify.

CMMVC7220E The task cannot be initiated because an invalid function license key file path was specified.

Explanation:
The file path specified for the license key file is invalid.

User response

Check for typing errors and try again.

CMMVC7221E The command has failed because the machine signature in the license does not match this machine.

Explanation:
The license key provided is not valid for this storage system.

User response

Use a license key that has been generated using this enclosure's machine signature.

CMMVC7222E The command has failed because the function specified by the license key has not been recognized.

Explanation:
The license key specified is for activating a function that is not supported on this level of firmware.

User response

update the storage system to a firmware level that supports the function and try again.

CMMVC7223E The command has failed because the trial specified has previously been used.

Explanation:
Each function has a trial period that can only be redeemed once. The trial period on this machine has already been used.

User response

Purchase the full license in order to continue using this function.

CMMVC7224E The command has failed because the specified function is currently in use.

Explanation:
A function can only be deactivated if is no longer in use.

User response

Ensure that the function that you need to deactivate is no longer in use, and try again.

Reduce the number of FlashCopy targets to 64 or fewer before deactivating the function.

CMMVC7226E The command has failed because the license key file specified is not in a supported format

Explanation

The command has failed because the license key file specified is not in a supported format.

User response

Check that the correct file was uploaded and try again.

CMMVC7233W Easy Tier is active without a license for each enclosure.

Explanation:
The system does not have enough licenses for Easy Tier.

User response

The user should correct their Easy Tier license entitlement.

CMMVC7234W FlashCopy is enabled without a license for each enclosure.

Explanation:
The system does not have enough licenses for FlashCopy.

User response

The user should correct theirFlashCopy license entitlement.

CMMVC7235W Remote copy is enabled without a license for each enclosure.

Explanation:
The system does not have enough licenses for remote copy.

User response

The user should correct their Remote copy license entitlement.

CMMVC7236W Multiple functions are enabled without a license for each enclosure.

Explanation:
The system does not have enough licenses for multiple functions.

User response

The user should correct their license entitlement.

CMMVC7238E The cycling mode may only be changed for Global Mirror relationships or groups.

Explanation

An attempt was made to change to cycling mode for an active-active relationship. This action is not permitted.

User response

If you must use cycling mode, you must change the relationship to Global Mirror by using the **chrcrelationship** command. This scenario is not common.

CMMVC7239E The node could not be added because the number of configured hosts exceeds the supported limit for the node type being added.

Explanation

One (or more) I/O groups has more hosts configured than is supported by the node hardware type being added.

User response

Either reduce the number of configured hosts or choose a different node hardware type to add.

CMMVC7240E **The node could not be added because at least one host has more volumes mapped than is supported for the node type being added.**

Explanation

One (or more) hosts has more volumes mapped than is supported by the node hardware type being added.

User response

Either reduce the number of volumes mapped or choose a different node hardware type to add.

CMMVC7241E **System code update cannot start because a component firmware update is in progress.**

Explanation:
An attempt was made to initiate either a system code update or a firmware update while the system was updating the firmware of various hardware components. This update cannot be completed while a firmware download is in progress, so the request failed.

User response

The firmware download must complete before you can perform another update. Due to the dynamic nature of firmware downloads, you cannot follow the progress of the download. Wait approximately 10 minutes and retry the command. You might need to repeat this step several times. You can use the **svcinfo lsupdate** command to see whether the firmware download completed.

CMMVC7242E **No help available for [%1].**

Explanation:
There is no help available for this command. [%1] shows the command for which help is not available.

User response

None.

CMMVC7243E **The specified port mask cannot be applied because insufficient paths would exist for node communication.**

Explanation:
The localfcportmask port mask value that was specified would cause one or more nodes to lose contact with the system.

User response

Check zoning. Fix any port errors in the event log. Use the lsfabric CLI command to ensure that when the correct port mask is specified and applied, all nodes would still have two paths to contact every other node in the system.

CMMVC7248E **You must use an 0x parameter for the -drivelba parameter.**

Explanation:
The parameter format for the **lsmdisklba** might be in error. See the description of the command to find an acceptable format.

User response:
Try the command again using the parameter format described in the command documentation.

CMMVC7249E **The name that you have entered is not valid. The name can contain letters, numbers, spaces, periods, dashes, and underscores. The name must not begin or end with a space. The name must not begin with a period.**

Explanation

A period or space cannot be the first character and a space cannot be the last character in the name that you enter. Also, the following characters are not supported anywhere in the name: * : , "" ' % #

User response

Ensure that the name you enter does not begin with a period, does not begin or end with a space character, and does not contain any of the unsupported characters listed above, and resubmit the task.

CMMVC7273E **[-rpoalert] is only valid when [-topology] does not have the value of [2-site-ha].**

Explanation:
A 2-site-ha topology must not have the **-rpoalert** parameter specified, as it cannot be changed for this topology.

User response:
Run the command again without the **-rpoalert** parameter.

CMMVC7300E **The maximum number of volumes already exist.**

Explanation:
The maximum number of volumes have been created and one or more must be deleted.

User response

Destroy one or most of the volumes before more can be created.

CMMVC7301E **The command failed because the volume would be smaller than the minimum size.**

Explanation:
A volume must be 1 MB or greater to be successfully created.

User response

Specify a capacity which is greater than the minimum supported capacity (1MB).

CMMVC7302E **The command failed because insufficient extents are available.**

Explanation:
Too many volumes have been created. Contact IBM support for assistance.

User response

If possible, delete unused volumes and try again. If the error persists, the extent mapping must be defragmented.

Contact IBM support for assistance.

CMMVC7304E **The command cannot be initiated because a spare is already configured.**

Explanation:
A configured spare already exists.

User response:
Only new drives can be formatted. If a spare already exists, the command cannot be executed.

CMMVC7305E **The command cannot be initiated because the drive fault is unrecoverable.**

Explanation:
A drive fault is preventing further actions.

User response:
Replace the drive.

CMMVC7306E **The command cannot be initiated because no array currently exists.**

Explanation:
No array has been created. No further action is possible.

User response:
Create an array before using this command.

CMMVC7307E **The command cannot be initiated because the battery slot number specified is invalid.**

Explanation:
Batteries occupy slot numbers 1 and 2 only for batteries.

User response:
Select Slot 1 or 2.

CMMVC7308E **The command cannot be initiated because an array already exists.**

Explanation:
An array exists that is prohibiting the desired action.

User response:
Only a single array can be created. Either create new volumes on the same array, or remove all volumes and arrays, then recreate the array.

CMMVC7309E **The command cannot be initiated because the specified raid level is unsupported on this platform.**

Explanation:
This error message is displayed when the RAID level that is specified in a command is not supported by the system.

User response:
Retry the command, specifying a supported RAID level.

CMMVC7310E **The command cannot be initiated because one or more drives is in the failed state.**

Explanation:
One or more drives has failed and is preventing most commands.

User response:
Work through Directed Maintenance Procedures (DMP) associated with failed drives.

CMMVC7311E **The command cannot be initiated because the number of drives is unsupported.**

Explanation

The number of installed drives is incorrect for the array configuration, or the capacity of the drives is not uniform.

- RAID 0 requires at least one drive.
- RAID 5 requires at least three drives.
- All the drives in an array must have the same capacity.

User response:
Remove or insert drives to obtain a supported configuration.

CMMVC7312E **The command could not be initiated because one or more volumes are using the array.**

Explanation:
One or more volumes are already using the array. All volumes must be removed before removing the array.

User response:
Remove all volumes before removing the array.

CMMVC7313E **The specified array does not require recovery as it is not corrupt.**

Explanation:
The specified array is not corrupt and does not require repair.

User response:
This command is only supported for corrupt arrays.

CMMVC7314E **The specified array cannot be recovered due to failed drives.**

Explanation:
One or more drives in an array have failed and this is preventing recovery.

User response:
Put the drives back in if they are missing.

CMMVC7315E **The command cannot be initiated because the specified slot does not exist.**

Explanation:
The command was attempted on an invalid slot.

User response

Batteries have slots 1 and 2. Enclosures have slots 3-12.

Retry command with a slot that exists.

CMMVC7316E **The specified array cannot be recovered due to failed drives.**

Explanation:
One or more failed drives are preventing the desired action. Too many errors.

User response:
Change the volume SCSI Drive ID by enabling volume Open Access.

CMMVC7317E **The command cannot be initiated because the specified PSU does not exist.**

Explanation:
A requested PSU does not exist.

User response:
Select another PSU or install the requested PSU.

CMMVC7318E **Open access is not enabled.**

Explanation:
Host mappings or logical host already exist.

User response:
Delete host mappings or the logical host.

CMMVC7319E **The air filter must be enabled to change the period.**

Explanation:
To change the timer for the air filter, the filter must currently be in use.

User response:
Enable the filter timer, then try the command again.

CMMVC7320E **Invalid port mask supplied.**

Explanation:
An invalid port mask was supplied.

User response:
Review the command parameters and try again with corrected value.

CMMVC7321E **Mask cannot be applied because insufficient paths exist for node communication.**

Explanation:
An invalid port mask was supplied.

User response:
Review the command parameters and try again with corrected value.

CMMVC7322E **Another volume uses this SCSI ID.**

Explanation:
The SCSI ID is already in use.

User response:
Use a different SCSI ID or remove the other volume.

CMMVC7323E **The command cannot be initiated due to a hardware fault.**

Explanation:
A hardware fault has occurred. Follow the Directed Maintenance Procedures (DMP).

User response

An unknown hardware fault exists.

Follow the DMP to resolve the hardware fault.

CMMVC7324E **The command cannot be initiated because the drive task is not supported.**

Explanation:
The drive does not support the command.

User response:
None.

CMMVC7325E **Open Access setting not changed because hosts exist.**

Explanation:
You cannot change the state of Open Access when hosts are defined.

User response:
Remove the hosts and try the command again.

CMMVC7326E **Cannot create logical host because open access is enabled.**

Explanation:
Open Access is preventing the creation of a logical host.

User response:
Disable Open Access if host mapping access is desired

CMMVC7329E **The maximum number of Infiniband GIDs for the host is already configured.**

Explanation:
The command has failed because the maximum number of Infiniband addresses has been exceeded.

User response:
Remove unused hosts addresses and try again.

CMMVC7330E **An invalid Infiniband GID was entered.**

Explanation:
The Infiniband address entered is not a valid address.

User response:
Check your input and try again.

CMMVC7331E **A supplied Infiniband GID is already assigned to another host.**

Explanation:
The Infiniband address configured is already assigned to a configured host.

User response:
Check your input and try again.

CMMVC7332E **The array does not exist.**

Explanation:
A flash memory array was not found by the Command Console LUN (CCL).

User response:
Call IBM support for instructions to resume your failed upgrade.

CMMVC7334E **The array could not be removed because it is initializing.**

Explanation:
A Remove Array operation was attempted while a newly created array was initializing.

User response:
Allow the array to complete initialization before attempting to remove it.

CMMVC7335E **The encryption state cannot be changed when an array exists.**

Explanation

An attempt was made to change the encryption state of the system when an array was present.

User response

Remove the array by using the **rmarray** command, then retry the **chencryption** command.

CMMVC7336E **The topology and speed are not compatible with the specified port.**

Explanation:
The user has entered a combination of speed and topology which is not compatible.

User response:
Check your input and try again.

CMMVC7337E **The command cannot be initiated because the encryption key was not found.**

Explanation

The system cannot locate a correct encryption key.

User response

Provide the current master encryption key to the system before you retry the command. If you use USB flash drives to manage the master encryption key of this system, connect one of the USB flash drives containing the current key to each node. If you are using a key server to manage encryption keys of this system, ensure that the system, key servers, and network are configured correctly so that a key server that contains the current master encryption key is online to the system.

CMMVC7338E **The command cannot be initiated because the encryption key is not valid.**

Explanation

The encryption key that was provided cannot be used.

User response

Provide the current master encryption key to the system before you retry the command. If you use USB flash drives to manage the master encryption key of this system, connect one of the USB flash drives containing the current key to each node. If you are using a key server to manage encryption keys of this system, ensure that the system, key servers, and network are configured correctly so that a key server that contains the current master encryption key is online to the system.

CMMVC7339E The command cannot be initiated because the drive did not unlock.

Explanation

The system was unable to unlock the drive by using the encryption key that was provided.

User response

Provide the current master encryption key to the system before you retry the command. If you use USB flash drives to manage the master encryption key of this system, connect one of the USB flash drives containing the current key to each node. If you are using a key server to manage encryption keys of this system, ensure that the system, key servers, and network are configured correctly so that a key server that contains the current master encryption key is online to the system.

CMMVC7340E The command cannot be initiated because the array is offline.

Explanation:
Failure to create a volume because the array specified is offline.

User response

Bring the array online.

CMMVC7341E The update cannot be resumed at this time.

Explanation:
The update cannot proceed because of hardware errors.

User response

Contact IBM support.

A service mode update may be necessary.

CMMVC7342E The array is already encrypted.

Explanation:
The array is already encrypted.

User response:
No response required.

CMMVC7343E A software upgrade is in progress and must finish before battery reconditioning can be started.

Explanation:
Battery reconditioning cannot be performed while a software upgrade is in progress.

User response:
Wait until the software upgrade has finished, and then restart battery reconditioning.

CMMVC7344E Battery reconditioning is not currently running.

Explanation:
Battery reconditioning is not running, therefore it cannot be cancelled.

User response:
No response required.

CMMVC7348E One of the batteries is currently reconditioning. Wait until reconditioning finishes.

Explanation:
Battery reconditioning is already running on one of the batteries. Only one battery can be reconditioned at a time.

User response:
Wait until the current battery reconditioning finishes before reconditioning another battery.

CMMVC7349E The other battery in the enclosure is not in a good and charged state.

Explanation:
Both batteries in the enclosure must be in a good and charged state to maintain redundancy during the battery reconditioning process. The selected battery cannot be reconditioned because the other battery in the storage enclosure is not charged, not installed, or it has an unresolved error condition.

User response:
Determine the condition of the other battery in the storage enclosure and take appropriate action. For example, install the battery if it is missing. If the battery has an error condition, check the event log and follow the recommended Directed Maintenance Procedure (DMP).

CMMVC7350E The selected battery cannot be reconditioned.

Explanation:
The selected battery is not in a state in which it can be reconditioned. The battery is either not charged, not installed, or it has an unresolved error condition.

User response:
Determine the condition of the selected battery and take appropriate action. For example, install the battery if it is missing. If the battery has an error condition, check the event log and follow the recommended Directed Maintenance Procedure (DMP).

CMMVC7351E Encryption cannot be disabled while an encrypted array exists.

Explanation:
System encryption cannot be disabled if there is an encrypted array.

User response:
Delete the encrypted array and then disable encryption.

CMMVC7352E The array cannot be encrypted because system encryption is disabled.

Explanation:
The array cannot be encrypted because system encryption is disabled.

User response:
Enable system encryption using the chencryption command and retry.

CMMVC7353E The command cannot be initiated because a rekey operation is in progress.

Explanation:
The command cannot execute because there is a rekey operation in progress.

User response:
Wait for the rekey operation to complete, and then re-enter the command. Use the lsencryption command to verify the status of the rekey operation.

CMMVC7354E Rekey operation failed.

Explanation:
The rekey operation failed, possibly due to hardware errors or a missing USB drive.

User response

Provide the current master encryption key to the system before you retry the command. If you use USB flash drives to manage the master encryption key of this system, connect one of the USB flash drives containing the current key to each node. If you are using a key server to manage encryption keys of this system, ensure that the system, key servers, and network are configured correctly so that a key server that contains the current master encryption key is online to the system.

Examine the event log to check for hardware errors. Contact IBM Support if you are unable to determine the cause.

CMMVC7355E The command cannot be initiated because system encryption is not enabled.

Explanation:
The command entered by the user requires system encryption to be enabled.

User response:
Enable system encryption and then re-enter the command.

CMMVC7356E Unable to validate the key on the current USB drive.

Explanation:
Validation failed. Insert the USB drive that contains the correct key for this system.

User response:
Review the output of the lseventlog and lsencryption commands for additional information, and take appropriate corrective action.

CMMVC7357E The command cannot be initiated because a drive has failed.

Explanation:
The command cannot be initiated because a drive has failed.

User response:
Resolve the problem that caused the drive to fail before retrying the operation. Review the event log for possible causes.

CMMVC7358E Rekey commit not allowed until keys have been copied.

Explanation:
An attempt was made to commit a new key without making sufficient copies of the new key.

User response:
Contact IBM Support.

CMMVC7359E Operation not applicable when encryption is disabled.

Explanation:
Either key validation was attempted while system encryption was disabled, or an attempt was made to change an array from unencrypted to encrypted while system encryption was disabled.

User response:
Enable system encryption and retry the command.

CMMVC7360E The command cannot be executed because the array is not online.

Explanation:

The array has to be online to run recoverarray -validate.

User response:
Use the lsarray command to check the value of the attribute parameter raid_status for the array. The command cannot be run if the array does not exist, or if the raid_status is **offline**, **degraded**, **syncing**, or **initiating**.

CMMVC7361E **The command cannot be initiated because another operation is in progress on the array.**

Explanation:
Another operation on the array is already in progress.

User response:
Wait until the current operation is completed before running another command.

CMMVC7362I **The encryption copy tool should only be used with the assistance of support.**

Explanation:
The command should only be run with the assistance of a support representative.

User response

Contact your support representative.

Attention: Do not proceed without the assistance of a support representative.

CMMVC7363E **The system is unable to mount/unmount the USB drive.**

Explanation:
The command was unable to either mount or dismount the USB drive.

User response

Contact your support representative.

Attention: Do not proceed without the assistance of a support representative.

CMMVC7364E **The system is unable to collect the information needed to provide the proper output from the lsencryption command.**

Explanation:
The command is unable to collect the output of the lsencryption command.

User response

Contact your support representative.

Attention: Do not proceed without the assistance of a support representative.

CMMVC7365E **The system could not complete the key copy operation because it could not read the new key file.**

Explanation:
The new key file was not found on the USB drive or the key server.

User response:
Provide the current master encryption key to the system before you retry the command. If you use USB flash drives to manage the master encryption key of this system, connect one of the USB flash drives containing the current key to each node. If you are using a key server to manage encryption keys of this system, ensure that the system, key servers, and network are configured correctly so that a key server that contains the current master encryption key is online to the system. Contact IBM Support if you are unable to resolve the problem.

CMMVC7366E **The system could not complete the key copy operation because it could not read the current key file.**

Explanation:
The current key file was not found on the USB drive or the key server.

User response:
If a key copy operation is running on a new system for the first time, then no response is necessary. Otherwise, provide the current master encryption key to the system before you retry the command. If you use USB flash drives to manage the master encryption key of this system, connect one of the USB flash drives containing the current key to each node. If you are using a key server to manage encryption keys of this system, ensure that the system, key servers, and network are configured correctly so that a key server that contains the current master encryption key is online to the system. Contact IBM Support if you are unable to resolve the problem.

CMMVC7367E **The key file is not valid.**

Explanation:
The key file is not valid.

User response

Ensure that the correct USB drive that contains the current master encryption key is inserted. You might need to insert a USB drive that contains a backup copy of the current master encryption key. Contact IBM Support if you are unable to resolve the problem.

If a USB flash drive fails, securely discard it and replace it with a new USB flash drive. Copy the key file from a known good USB flash drive to the new USB flash drive.

CMMVC7368I **%1 additional copy has been made.**

Explanation:
Copy # %1 of the new key was made.

User response
No response required.

This is an informational message that states how many copies of the new key have been made.

CMMVC7369E **No additional key copies were written to the USB drive; all required keys already exist on the drive.**

Explanation:
No additional keys were written to the USB drive because both the new and current keys already exist on the drive.

User response:
Insert a new USB drive that does not contain any key files and retry the command, or call IBM Support to continue.

CMMVC7370W **Only the new key file has been copied to the USB drive; the current key was already on the drive.**

Explanation:
Only the new key file has been written to the USB drive; the current key was already on the drive.

User response:
The USB drive already contains the current key file. If the USB drive was expected to be empty, examine the files on the drive to determine the problem.

CMMVC7371E **Only the current key file has been copied to the USB drive; the new key was already on the drive.**

Explanation:
Only the current key has been written to the USB drive; the new key was already on the drive. Call IBM Support to continue.

User response:
Call IBM Support to continue. Do not proceed without IBM Support assistance.

CMMVC7372E **The current key on the USB drive does not match the key that was previously read.**

Explanation:
The current key on the USB drive does not match the key that was previously read.

User response:
Call IBM Support to continue. Do not proceed without IBM Support assistance.

CMMVC7373E **The attempted write of the new key to the USB drive has failed.**

Explanation:
The attempted write of the new key to the USB drive has failed.

User response:
Ensure that the correct USB drive is inserted. Call IBM Support if unable to resolve the problem.

CMMVC7374E **The attempted write of the current key to the USB drive has failed.**

Explanation:
The attempted write of the current key to the USB drive failed.

User response:
Ensure that the correct USB drive is inserted. Call Support if unable to resolve the problem.

CMMVC7375E **No keys were written to the USB drive. Do not proceed without Support assistance.**

Explanation:
No keys were written to the USB drive. Contact your support representative.

User response
Contact your support representative.



Attention: Do not proceed without the assistance of a support representative.

CMMVC7376E **The command cannot be run because an encryption key has not been created.**

Explanation:
This command requires an encryption key, but the encryption state has not been properly set up.

User response:
Use the encryption enablement procedure to create encryption keys, or contact IBM Support for assistance.

CMMVC7377E **The command cannot be executed because the array does not exist.**

Explanation:
This command requires an encryption key, but the encryption state has not been properly set up.

User response:
Create an array and then reattempt the command.

CMMVC7378E **The command cannot be executed because the array is not encrypted.**

Explanation:
The array must be encrypted before running the command.

User response:
Encrypt the array and then reattempt the command.

CMMVC7379E **System encryption cannot be disabled because an encrypted array already exists.**

Explanation:
System encryption cannot be disabled while an encrypted array exists. To disable encryption, the encrypted array must be deleted, which will result in the loss of data.

User response:
To disable encryption, delete the encrypted array, and then reattempt the command. Contact IBM Support for assistance before performing this step.

CMMVC7380E **The command cannot be executed because an encrypted array exists.**

Explanation:
The command cannot be run because the array encryption cannot be disabled.

User response:
No response is possible. The command is not allowed in this context.

CMMVC7381E **The rekey failed because the current encryption key was not found on the USB drive.**

Explanation:
Explanation: The rekey operation could not be completed because the current encryption key was not found on the USB drive.

User response:
Cancel the failed rekey operation, restore the current encryption key file to the USB drive, and then reattempt the operation.

CMMVC7382E **The rekey failed because the new encryption key was not found on the USB drive.**

Explanation:
The rekey operation could not be completed because the new encryption key was not found on the USB drive.

User response:
Cancel the failed rekey operation, restore the new encryption key file to the USB drive, and then reattempt the operation.

CMMVC7383E **The rekey failed because the proposed new encryption key could not be generated.**

Explanation:
The rekey operation failed because the new encryption key could not be generated. The USB flash drive might be unusable.

User response

Cancel the failed rekey operation and replace the USB flash drive before you attempt a new rekey operation.

If the `chencryption -usb newkey -key prepare` command was issued, then the encryption key file might already exist on the USB flash drive.

Replace the USB flash drive with a blank USB flash drive and retry the command. Securely format or destroy the unusable USB flash drive.

CMMVC7384E **The rekey operation failed because one of the drives failed to prepare for the rekey.**

Explanation:
The rekey operation failed because the -key prepare option of the chencryption command failed on one of the drives.

User response:
Review the event log and resolve the corresponding event.

CMMVC7385E **The rekey operation failed because one of the drives failed to commit to the rekey.**

Explanation:
The rekey operation failed because the -key commit option to the chencryption command failed on one of the drives, causing the drive to fail to commit.

User response:
Review the event log and resolve the corresponding event.

CMMVC7386E **The rekey operation failed because one of the drives failed to cancel the rekey.**

Explanation:
The rekey operation failed because the -key cancel option to the chencryption command failed on one of the drives.

User response:
Review the event log and resolve the corresponding event.

CMMVC7387E **The command cannot be initiated because the rekey is not in an expected state.**

Explanation:
The command cannot be initiated because the rekey is not in an expected state.

User response

Issue the appropriate command based on the current rekey state.

A commit operation can only be performed when the rekey state is prepare or commit_failed.

A cancel operation can only be performed when the rekey state is prepare_complete, prepare_failed or cancel_failed.

CMMVC7388E **Rebuild options are not supported for this array RAID level.**

Explanation:
The command cannot be initiated because the rekey is not in an expected state.

User response:
Issue the command against a RAID 5 array. Do not target a RAID 0 array.

CMMVC7391E **Specified drive is in the incorrect enclosure.**

Explanation

An attempt was made to swap an array member for a spare. Arrays for this product must be contained within a single enclosure. The specified drive is in the wrong enclosure.

User response

Retry the command and specify a new array member that is in of the same enclosure as the rest of the array.

CMMVC7392E **The enclosure is already managed by a different cluster.**

Explanation

An enclosure cannot be managed by more than one cluster at a time.

User response

If you want to change the cluster that manages the enclosure, complete the following steps:

1. Log in to the cluster that is currently managing the enclosure.
2. Delete the array on the enclosure by using the management GUI or by running the **rmarray** command.
3. Stop managing the enclosure by using the management GUI or by running the **chenclosure -managed no** command.

If you do not have access to the cluster that is currently managing the enclosure, you can log on to the enclosure and run the **satask leavecluster -force** command. Using the **-force** parameter can lead to data loss if the enclosure is in use.

After you remove the enclosure from the management of the previous cluster, you can retry the original **chenclosure -managed yes** command.

CMMVC7393E **The enclosure cannot be unmanaged as there is an array present.**

Explanation

You can remove management from an expansion enclosure only when no arrays are configured.

User response

Remove the array from the enclosure by using the management GUI or by running the **rmarray** command. Then, retry the original command.

CMMVC7394W **Encryption feature is being used without system license entitlement.**

Explanation

An encryption license was not available for the command that was run. The command completed, but you must still obtain a valid license.

User response

Contact your support representative to obtain the necessary encryption license.

CMMVC7395E **The command cannot be executed because the enclosure is not a part of the cluster.**

Explanation

You tried to create an array with a flash enclosure that is not managed.

User response

Change the flash enclosure to a managed type by using the following command:

```
chenclosure -managed yes enclosureId
```

where *enclosureId* is the ID of the enclosure as displayed in the **lsenclosure** command. Then, retry the command.

CMMVC7396E The command failed because the target managed disk (MDisk) is in offline state.

Explanation

An MDisk cannot be added to a storage pool when that MDisk is offline.

User response

Make sure that you specified the correct MDisk. If so, make sure the specified MDisk is online and accessible before you retry the command.

CMMVC7399E The command cannot be run as the enclosure is not yet managed.

Explanation

This command requires the enclosure to be in the managed state.

User response

Manage the enclosure by using the management GUI or by running the **chenclosure -managed yes** command. Then, retry the original command.

CMMVC7402E The command has failed because the IP address is not valid.

Explanation

The command specified an IP address with an invalid format.

User response

Retry the command with a valid IPv4 or IPv6 address.

CMMVC7403E The command cannot be initiated because it is not supported for the specified enclosure type

Explanation

The specified enclosure was not the correct type for this command.

User response

Review the command documentation or specify an alternative enclosure.

CMMVC7404E The command cannot be initiated because it is not supported for the specified drive type

Explanation

The specified drive was not the correct type for this command.

For systems that support NVMe drives, the **chquorum** command is not supported on NVMe drives.

User response

Review the command documentation or specify an alternative drive.

For systems that support NVMe drives, specify a SAS drive or managed disk when you use the **chquorum** command.

CMMVC7405E The command cannot be initiated because it is not supported for the specified array type

Explanation

The specified array was not the correct type for this command.

User response

Review the command documentation or specify an alternative array.

CMMVC7406E The command failed because the requested volume size is too large. Reduce the array's reserved capacity to free space.

Explanation:
The **mkvdisk -size** and **chvdisk -size** commands cannot encroach on the array space that is reserved for performance enhancement.

User response

- Specify a smaller volume size.
- Reduce the amount of space in the array that is reserved to increase performance.

CMMVC7407E The array has insufficient free space to reserve for performance.

Explanation:
The array does not have sufficient free space reserved for performance.

User response

Complete one of the following tasks:

- Use the **charray -reservesize** command to reserve a size that is smaller than that of the array.
- Delete volumes to increase available space in the array.

CMMVC7408E **The command has failed because one or more of the parameters entered is not valid for the specified enclosure type.**

Explanation

Some parameters are only valid when used with certain types of enclosure.

User response

Review the command syntax for the type of enclosure specified or specify an alternative enclosure.

CMMVC7409E **The command cannot be initiated because the port is not online.**

Explanation:
chportip must be used on a port that is online.

User response:
Use lsportip to see the state of the port. If the state is listed as**offline**, the port is offline.

CMMVC7410E **The command has failed because one or more of the parameters entered is not valid for the specified drive type.**

Explanation

Some parameters are only valid when used with certain types of drives.

User response

Review the command syntax for the type of drive specified or specify an alternative drive.

CMMVC7411E **The command has failed because one or more of the parameters entered is not valid for the specified array type.**

Explanation

Some parameters are only valid when used with a certain type of array.

User response

Review the command syntax for the type of array specified or specify an alternative array.

CMMVC8000E **Cannot execute on an active node.**

Explanation

This operation cannot be executed on an active node.

User response

Choose an inactive node, or deactivate the node before trying the operation again.

CMMVC8001E **Cannot execute on a candidate node.**

Explanation:
This operation cannot be executed on a candidate node.

User response

Choose another node and try the operation again.

CMMVC8002E **Cannot execute on a service-state node.**

Explanation:
This operation cannot be executed on a service-state node.

User response

Choose another node and try the operation again.

CMMVC8003E **Cannot execute on a node in a cluster-recovery state.**

Explanation

This operation cannot be executed on a node in a cluster-recovery state.

User response

Choose another node and try the operation again.

CMMVC8004E **Cannot execute on a node with a location error.**

Explanation:
You cannot process this operation on a node with a location error.

User response:
Fix the error or choose another node and try the operation again.

CMMVC8005E **Cannot execute on a node displaying hardware errors.**

Explanation:
You cannot execute this operation on a node displaying hardware errors.

User response:
Fix the errors or choose another node and try the operation again.

CMMVC8006E **Cannot execute on a node displaying errors.**

Explanation:
You cannot execute this operation on a node displaying errors.

User response:
Fix the errors or choose another node and try the operation again.

CMMVC8007E Cannot execute on a node which is charging.

Explanation

You cannot execute this operation on a node that is charging its batteries.

User response

Wait until the batteries have finished charging and try the operation again.

CMMVC8008E Command not supported on this hardware platform

Explanation

The command is not supported on this system or cannot be run on this node.

User response

Ensure that the correct service actions are being attempted for the hardware, and ensure that the appropriate command is run on the correct node, or choose another node and try the operation again.

CMMVC8009E Cannot execute on a node canister.

Explanation:
You cannot execute this operation on a node canister.

User response

Choose an appropriate target and try the operation again.

CMMVC8010E Not from a USB stick.

Explanation:
You cannot execute this operation from a USB stick.

User response

Change location to an appropriate place and try the operation again.

CMMVC8011E Version too high for this client.

Explanation:
You cannot complete this operation during a manual update.

User response:
Wait until the manual update finishes and try the operation again.

CMMVC8012E The operation did not complete in the time allowed.

Explanation:
The operation did not complete in the time allowed.

User response

Set the time allowed for operations to complete appropriately. Determine if another operation or error caused the problem.

CMMVC8013E Incompatible parameters set.

Explanation:
Parameters provided are mutually exclusive.

User response

Set appropriate parameters and try the operation again.

CMMVC8014E Bad parameter value.

Explanation:
This error occurs when a command is specified with an unparseable IP address, worldwide node name (WWNN), or other unknown parameter value. This error can also occur when a specified password contains too few or too many characters. Passwords must contain more than 5 and fewer than 65 ASCII characters.

User response:
Retry the command, specifying appropriate parameters.

CMMVC8015E Failed.

Explanation:
The enclosure midplane has a cluster ID set or the next cluster ID on the midplane is corrupt or invalid.

User response

Fix the problem and try the operation again.

CMMVC8016E Node would be clustered if not in service state.

Explanation:
Cannot run because the node will be in cluster when it exits service.

User response

Fix the problem and try the operation again.

CMMVC8017E Info value not recognised.

Explanation:
Info value not recognized.

User response
Use a valid info value and try the operation again.

CMMVC8018E Provided buffer was too small.

Explanation:
Provided buffer was too small.

User response
Increase the buffer size.

CMMVC8019E Task could interrupt IO and force flag not set.

Explanation:
Running on an active node could impact I/O.

User response
Wait until the node is inactive before performing this task.

CMMVC8020E Attempting to create cluster while there is a stored cluster ID.

Explanation:
An attempt was made to create a cluster while the control enclosure or the node has a stored cluster ID.

User response
Change the cluster ID or chose a different control enclosure or node.

CMMVC8021E Invalid panel name.

Explanation:
Invalid panel name given in the parameter.

User response
Use a valid panel name.

CMMVC8022E New cluster created but node cannot leave service state.

Explanation:
New cluster created but node cannot leave service state. The battery might be charging or some other service task running.

User response
Wait until all service tasks complete.

CMMVC8023E Partner node is clustered.

Explanation:
Cannot execute because the partner node is clustered.

User response
Take the partner node out of the cluster, or choose a different node.

CMMVC8024E Gateway or subnet/prefix required.

Explanation:
A Gateway or subnet/prefix required.

User response
Use a Gateway or subnet/prefix.

CMMVC8025E DHCP Failed.

Explanation:
DHCP Failed.

User response
Try operation again.

CMMVC8026E No suitable donor.

Explanation:
No suitable donor.

User response
[User response needed].

CMMVC8027E T3 prepare failed.

Explanation:
T3 prepare failed.

User response
[User response needed].

CMMVC8028E T3 prepare incomplete.

Explanation:
T3 prepare not complete.

User response
[User response needed].

CMMVC8029E T3 execute failed.

Explanation:
T3 execute failed.

User response
[User response needed].

CMMVC8030E Another instance of this command is already running.

Explanation:
Cannot execute because another instance of this command is already running.

User response
Wait for the other instance of the command to complete.

CMMVC8031E file not found.

Explanation:
Required/provided file is not on the file system in the expected location.

User response
Locate the missing file.

CMMVC8032E The specific update package cannot be installed over the current version.

Explanation:
Cannot install this code over current version; cannot update to this version; code is already at this level.

User response
Make sure you are installing the right version.

CMMVC8033E Password reset disabled.

Explanation:
The password reset function is disabled.

User response
Enable the password reset function or contact your system administrator.

CMMVC8034E A compulsory parameter is missing.

Explanation:
Cannot execute because required parameter has not been provided.

User response
Provide the required parameter.

CMMVC8035E The service assistant CLI is not ready - try again.

Explanation:
The Service CLI interface is not ready/running yet.

User response
Wait a few moments and try again.

CMMVC8036E No help available.

Explanation:
There is no available help.

User response
Contact IBM Support.

CMMVC8037E A required parameter is missing.

Explanation:
A required parameter is missing.

User response
Rerun the operation with the required parameter set.

CMMVC8038E Required parameters are missing.

Explanation:
Required parameters are missing.

User response
Rerun the operation with the required parameters set.

CMMVC8039E The [%1] parameter is missing its associated arguments.

Explanation:
The parameter is missing an argument.

User response
Rerun the operation with the required arguments for the parameter.

CMMVC8040E [%1] is not a supported parameter.

Explanation:
The parameter is not supported.

User response
Rerun the operation using a supported parameter.

CMMVC8041E [%1] is not a valid command line option.

Explanation:
The command given does not exist.

User response
Use an existing command.

CMMVC8042E Invalid or inconsistent arguments.

Explanation:
Invalid or inconsistent arguments. For example, the argument at the end is not a recognized panel id.

User response
Use valid and consistent arguments.

CMMVC8043E This command can only be run by the superuser.

Explanation:
Cannot execute because user is not superuser.

User response
Have a superuser run the command.

CMMVC8044E Command completed successfully.

Explanation:
Command completed successfully. This message is only used within lscmdstatus.

User response
None.

CMMVC8045E Authentication failed.

Explanation:
This error occurs when an attempt is made to log in to the system by using invalid credentials.

User response:
Ensure that the correct username is specified with the correct password or SSH key. If the behavior continues, contact the security administrator.

CMMVC8046E Partner node has lost cluster data.

Explanation:
Partner node has lost cluster data.

User response
[Need user response].

CMMVC8047E Not a valid ssh key.

Explanation:
Argument is not a valid ssh key.

User response
Use a valid ssh key for the argument.

CMMVC8048E Invalid file permissions.

Explanation:
Cannot execute the argument. The file has invalid file permissions.

User response
Set valid file permissions.

CMMVC8049E Invalid cluster name.

Explanation:
User provided an invalid cluster name.

User response
Use a valid cluster name.

CMMVC8050E Files cannot be unpacked from the update package. The system must unpack them.

Explanation
Possible causes are:

1. A bad boot drive or sector.
2. Full `/upgrade`, `/tmp`, or `/upgrade`.
3. An invalid package.

User response

1. Clear all dumps and retry the installation.
2. Reboot the node and retry the installation.

CMMVC8051E Utilities package installed.

Explanation:
The utilities package installed successfully.

User response
None.

CMMVC8052E The signature of the utility package has failed verification.

Explanation
This can be caused by the following issues:

1. The package has become corrupted.
2. The package is not a valid IBM utility.
3. The system clock on the node is long out-of-date and the package signature is too far into the future.

User response

1. Ensure that the installation package is complete and retry the installation.
2. Verify that the utility was provided by IBM support personnel.
3. Use chsystemtime Change the system clock to accurately reflect the date.

CMMVC8053E The specific update package cannot be installed on this hardware.

Explanation:
The software is incompatible with the hardware level.

User response

Ensure that you have downloaded the appropriate package for the hardware you are updating.

CMMVC8054E The update failed as the package is missing files.

Explanation:
The package can be missing files due to a bad upload.

User response:
Validate that the package was downloaded or uploaded properly and try to run the update again.

CMMVC8055E The command cannot be run because the node is busy.

Explanation:
The node or enclosure can run only one task command at a time, or the firmware is being updated so the command cannot be run currently.

User response

Wait for the task to complete and run the command again.

CMMVC8056E An error occurred in communicating with the target node.

Explanation

The error can be caused by:

1. A fault in one partner node is preventing another partner node from seeing it.
2. The Fibre Channel network is congested or faulty, and packages are failing to transfer.

User response

1. If you are using SAN connectivity, check your fabric to ensure that all the nodes in a cluster have clear paths to one another.
2. Ensure that the target node is online and does not have a hardware or location error (if it is Enclosure based).
3. Ensure that both the source and target nodes see each other in lsservicenodes. If they do not see each other, a path is missing.
4. If trying the first three steps does not correct the problem, work directly on the node in question rather then from the partner node. Log onto the target node's service assistant and run the operation locally.

CMMVC8057E Files cannot be copied to a node which is not part of the source cluster.

Explanation:
Files cannot be copied to a node which is not part of the source cluster using cpfiles.

User response

Choose a node that is part of the source cluster.

CMMVC8058E Cannot create file, file already exists.

Explanation:
The ssh key file you are trying to create already exists.

User response

None.

CMMVC8059E The update package supplied cannot be installed using service state whilst maintaining the cluster configuration on this node. To maintain the cluster configuration on this node, this package can only be installed with applysoftware or pacedccu mode. This package can be installed in service state using the -ignore flag, however the cluster state will be destroyed and cluster configuration will be lost from the node.

Explanation

A software update to this level cannot be issued when it uses service state without using the **-ignore** flag. The **-ignore** flag removes the cluster configuration from the node. If you want to maintain the cluster configuration, you must install the package by using the automated **applysoftware** command or by manual update.



Attention: Be careful when you use the **-ignore** flag. You might see adverse consequences to the data that is processed.

User response

Use the correct procedure to update the code package.

CMMVC8060E DHCP fallback is not supported on this platform.

Explanation:
You tried to set the service IP via DHCP with fallback enabled. This platform does not support the fallback option.

User response

Set the service IP via DHCP without fallback enabled. This product does not support the fallback option.

CMMVC8061E The enclosure does not support setting a machine part number.

Explanation:
Cannot execute because there is no -machinepartnum on this system.

User response

None.

CMMVC8062E The machine type and model is not valid for this enclosure.

Explanation:
The machine type and model is not valid for this enclosure.

User response

Use a valid machine type and model with the command.

CMMVC8063E The machine part number is not valid for this enclosure.

Explanation:
The machine part number is not valid for this enclosure.

User response

Use a valid machine part number with the command.

CMMVC8064E The machine part number and machine type and model are not valid for this enclosure.

Explanation:
The machine part number and machine type and model are not valid for this enclosure.

User response

Use a valid machine part number and machine type and model with the command.

CMMVC8065E An enclosure property has already been set and cannot be modified.

Explanation:
An enclosure property has already been set and cannot be modified. A valid VPD exists and will not be overwritten.

User response

If the machine serial number is not 00000000 or is not the same as the serial number stored on one of the boot drives, then the system board must be replaced again with a clean FRU that has a machine serial number of 0000000 stored on it. .

CMMVC8066E A new enclosure VPD field does not match the node copy.

Explanation:
A new enclosure VPD field does not match the node copy. An unexpected value was given.

User response

Rerun the command with the correct value in the enclosure VPD field.

CMMVC8070E The specified IP addresses are not in the same sub-network.

Explanation:
The IP address parameter values in satask.txt and cfgtask.txt that must be in the same sub-network are not all in the same sub-network. i.e. The bit wise AND of an IP address with the subnet mask is not the same as the bit wise AND of another IP address with the subnet mask.

User response

Check and correct the parameters in satask.txt and cfgtask.txt or use InitTool to create new valid satask.txt and cfgtask.txt files on the USB flash drive to start the initial setup of the system.

CMMVC8071E The specified IP address is already in use.

Explanation:
An arping of the IP address has received a reply from an IP host on the network that is already using this IP address from satask.txt.

User response

If you cannot stop the other IP host from using that IP address then allocate another IP address and use InitTool to create new valid satask.txt and cfgtask.txt files on the USB flash drive to start the initial setup of the system.

CMMVC8072E The above parameters are invalid or missing from cfgtask.txt.

Explanation:
The required parameters shown above this error message are invalid or missing from cfgtask.txt.

User response

Check and correct the parameters in cfgtask.txt or use InitTool to create new valid satask.txt and cfgtask.txt files on the USB flash drive to start the initial setup of the system.

CMMVC8085E The command failed because the node does not support technician port functionality.

Explanation

Some older models do not provide a technician port.

User response

To connect a workstation directly to a node without a technician port, connect the workstation using Ethernet port 1 or 2 and an IP address compatible with the system IP or service IP of the node. If the IP configuration of these ports is not known, use a USB flash drive inserted into the USB port of the node.

CMMVC8086E The command failed because the node already has a dedicated technician port.

Explanation

When a node has a dedicated technician port, it is not possible to configure another port as a technician port. The dedicated technician port is permanently enabled.

User response

If you need to connect a workstation directly to the node for maintenance, use the dedicated technician port.

CMMVC8087E The command failed because the file specified is a valid USVNID file but it is for a different node.

Explanation

An activation file was provided that is in the correct format but has an incorrect unique ID (obtained from the node during the installation process).

User response

Verify that the correct file was downloaded. You might need to repeat the steps that you took to generate this file, making sure that you use the correct node ID.

CMMVC8088E The command failed because the file specified is not a valid USVNID file.

Explanation

An activation file was provided that is in the wrong format.

User response

Verify that the correct file was downloaded. You might need to repeat the steps that you took to generate this file.

CMMVC8091E No upload in progress.

Explanation

An attempt was made to cancel an upload when no upload was in progress.

User response

Only active uploads can be canceled.

CMMVC8092E An error occurred in communicating with the remote server.

Explanation

An attempt was made to connect to the remote server through the **supportupload** or the **downloadsoftware** command. The connection was not completed.

User response

See the prerequisites for the **supportupload** or the **downloadsoftware** command. After you meet those conditions, retry the command.

CMMVC8093E No download in progress.

Explanation

An attempt was made to cancel a download when no download was in progress.

User response

Only active downloads can be canceled.

CMMVC8094E The command failed because the available space is inadequate for the download.

Explanation

An attempt was made to download files from the Fix Central server by using the **downloadsoftware** command. Not enough space was available for the download to complete.

User response

Remove unwanted files from the target system and retry the command.

CMMVC8095E **The command failed because the target Ethernet port does not exist.**

Explanation

Ethernet ports are numbered sequentially beginning with 1. The specified Ethernet port has a number that is larger than the total number of Ethernet ports on the system.

User response

Retry the command with a valid Ethernet port number. To display the available Ethernet ports, enter one of the following commands:

- **lsportip**
- **sainfo lsnodeip**
- **sainfo lsservicestatus**

CMMVC8096E **The command failed because insufficient paths would exist for the IP connection between the nodes in the same IP discovery zone.**

Explanation

An attempt was made to remove or change the IP address of a non-redundant connection. Removing or changing this address can cause the I/O group to lose redundancy, or can cause a degraded I/O group to become unavailable.

User response

Complete one of the following actions:

- Use the **satask chnodeip** command to add a node IP address to the non-redundant connection nodes in the same IP discovery zone. You can then retry the command.
- You can retry the command with the **-force** parameter. Use of this parameter can cause the I/O group to lose redundancy, or can cause a degraded I/O group to become unavailable. Because the **-force** parameter can have unforeseen consequences, it is typically used in test situations only.

CMMVC8097E **The command failed because a duplicate WWNN has been detected.**

Explanation

This error can occur if a duplicate worldwide node name (WWNN) is detected when you run the **chnodeip** command.

User response

Issue the `satask chvpd -wwnn` *WWNN* command to change the duplicate WWNN. For more information about correcting duplicate WWNNs, see the documentation for node error 556.

CMMVC8098E **The command failed because the specified ethernet port is not an RDMA capable port.**

Explanation

This error occurs when an attempt is made to configure a port for Remote Direct Memory Access (RDMA) clustering, but the port does not support RDMA.

User response

Retry the command, specifying a port that supports RDMA.

CMMVC8099E **The command failed because the target Ethernet port is set to non default MTU value.**

Explanation

This error occurs when an attempt is made to configure a port for Remote Direct Memory Access (RDMA) clustering, but the MTU value of the port is not the default setting (1500).

User response

Use the **chportethernet** command to change the MTU value of the port to 1500, and then retry the command.

CMMVC8202E **To perform this task, a system board mismatch node error must be active.**

Explanation

This error occurs when an attempt is made to run the **chbootdrive -replacecanister** command but no mismatch between the boot drives and system board or canister can be detected.

User response

Use the **lsbootdrive** command to determine the current state of the boot drives. If all drives are online, no further action is required. If one or more drives are offline, use the **lsservicenodes** command to determine the active node error, and the follow the suggested procedure for that error.

CMMVC8203E To perform this task, the number of boot drives present must be correct.

Explanation

This error occurs when an incorrect number of boot drives are detected. The node canisters of a model AF7 control enclosure must have one boot disk. The node canisters of a model AF8 control enclosure must have two boot disks.

User response

Depending on the control enclosure, either attach the missing drive or remove the additional drive.

CMMVC8204E To perform this task, the boot drive owners (actual_node_sn) must be the same.

Explanation

This error occurs when the `actual_node_sn` values displayed by the **sainfo lsbootdrive** command do not match.

User response

The `actual_node_sn` values for both boot drives must match the canister serial number of the replaced canister. The serial number can be read from the sticker at the rear of the canister. Shut down the node, locate and install the missing boot drives, restart the node, and then run the command again.

CMMVC8206E The command failed because the IP address is a reserved/special IP.

Explanation

The IP address that was specified in the **chnodeip** command is special or reserved.

User response

Specify an IP address that is not special or reserved to configure a node.

CMMVC8207E Backup data for the midplane is unavailable.

Explanation

The midplane backup data is not available and must be entered manually.

User response

Provide the necessary information by using the following command:

```
satask chvpd -serial serial_number -mtm mtm
```

CMMVC8208E The maximum number of discovery subnets are already configured on the node.

Explanation

This error occurs if an attempt is made to configure more than 16 discovery subnets on a node.

User response

Remove an unused discovery subnet from the node, and then try the **addnodediscoverysubnet** command again.

CMMVC8209E This subnet conflicts with one that is already configured on the node.

Explanation

This error occurs if a discovery subnet is specified that overlaps with the address range of a discovery subnet that is already configured on the node.

User response

Retry the command, specifying a discovery subnet that does not overlap with the address ranges of any discovery subnets that are already configured on the node.

CMMVC8210E This discovery subnet conflicts with the one reserved by IANA.

Explanation

This error occurs when an attempt is made to specify a discovery subnet that uses an address range that is reserved by the Internet Assigned Numbers Authority (IANA).

User response

Retry the command, specifying a discovery subnet that does not use an address range that is reserved by the IANA.

CMMVC8211E **This subnet is too big to be set as a discovery subnet.**

Explanation

This error occurs when an attempt is made to specify a discovery subnet that includes more than 16 IP addresses. For an IPv4 address, this limit means that the network prefix must be 28 bits or larger.

User response

Retry the command, specifying a discovery subnet that includes 16 or fewer IP addresses.

CMMVC8212E **Node rescue is not supported in this platform.**

Explanation

This error can occur when an attempt is made to run the **satask rescuenode** command on a cloud node.

User response

Remove the cloud node from the system, reinstall the software on the cloud node, and then add the node back to the system.

CMMVC8216E **An insecure password was chosen.**

Explanation

This error can occur when an attempt is made to reset the password to the default password, or to use a password that does not comply with the password policy.

User response

Specify a password that complies with the password policy, and is not the default password.

CMMVC8217E **Cannot execute on a spare node.**

Explanation

This error occurs when an attempt is made to unlock the superuser account on a node that is a spare node for a clustered system. Unlocking the superuser account is not support on spare nodes.

User response

Retry the command, specifying a node that is an active member of a clustered system.

CMMVC8219E **The command failed because a maximum of 8 Ethernet clustering ports can be configured per node.**

Explanation:
This error occurs when you try to configure more than the maximum number of ports that are allowed. A maximum of 8 Ethernet clustering ports can be configured per node.

User response:
Remove the existing Ethernet clustering port from the node configuration by using the command **satask chnodeip** with **-noip** option. Then, retry the command to configure the specified port.

CMMVC8222I **Patch operation succeeded.**

Explanation

This message is informational only.

User response

No action is required.

CMMVC8232E **Patch not allowed for this configuration.**

Explanation

This error occurs when the patch is not compatible with the current software level (including previously applied patches) of the node or system.

User response

Retry installing the correct patch for the software level of the system or node. If the problem persists, contact IBM support for assistance.

CMMVC8261E **The command failed because the hardware configuration of the local cluster is not compatible with the software of a partnered cluster.**

Explanation:
The software version of the local cluster is newer than the software version of a partnered cluster, and additional hardware has been enabled that is not supported by the older software on the partnered cluster.

User response

Either update the software on the partnered cluster, turn off the new hardware on the local cluster, or stop the remote copy relationship with the remote cluster. Use the CLI command chnodehardware -legacy to disable hardware that is not supported by older software versions.

CMMVC8262E **Cannot remove the host from the specified I/O groups because host mappings exist in multiple I/O groups.**

Explanation:
If a host is defined in multiple I/O groups, and has access to a volume through all of those I/O groups, an attempt to remove the host from just one of those I/O groups always fails, even when the **-force** parameter is specified.

User response

Before you run the command again, complete one of the following steps:

- Modify the I/O group list so that the list includes all of the I/O groups that provide access to the volumes for the host. The host mappings are removed when the command runs.
- Modify the I/O group list so that the list excludes I/O groups that contain host mappings.

CMMVC8263E **The command failed because the volume is associated with a file system and your requested action can not be completed under your current user role.**

Explanation:
You are attempting to complete an action on a volume that is associated with a file system. However, you do not possess the required role for file system actions.

User response

Execute command through volume command.

CMMVC8264E **The command failed because the volume associated with a file system and only the real capacity of a compressed file system volume can be changed.**

Explanation:
You are attempting to resize a volume that is associated with a file system. However, you can only resize the real capacity of a file system volume if it is compressed.

User response

The command cannot be completed on this volume. It will only succeed with a volume that is not associated with a file system or a with a compressed file system volume.

CMMVC8265E **The action failed because the specified port is for management only.**

Explanation:
The action failed because the specified port is for management only.

User response

Try another port which is not marked as management_only in the output of lsportip.

CMMVC8266E **The action failed because the specified port is not installed.**

Explanation:
The action failed because the specified port is not installed.

User response

Use a port that is shown in the output of lsportip but which is not marked as management_only.

CMMVC8267E **The attempt to add the node to the system failed because the node is being manually updated. The system must be prepared for update by using the same package that is currently installed on the node that is being updated.**

Explanation

The manual update requires that you prepare the system for the update first. Manual update mode (the **-pacedccu** parameter) is in use and an attempt is being made to add a node to the system at a newer code level. However, one of the following situations occurred:

- The code level of the node that is being added requires that you prepare the system for update before you attempt to use the **addnode** command
- The system was prepared at a different code level than the version of software on the node that is being added

User response

Prepare the system for update with the same software version as the manual update node that you are

adding. To do so, complete one of the following actions:

- Cancel the update and prepare the system with the correct package.
- Install a version of the code on the new node that matches the system version.

CMMVC8268E **The attempt to prepare the cluster for update has failed because the cluster has already been prepared with a different package level. The update must first be cancelled before reattempting the update.**

Explanation:
The user has prepared the cluster for update with one level package and then attempted to prepare the update with a different, higher level package or attempted to automate the update to a higher level after the cluster was prepared with a lower level package.

User response

The user must cancel the current update that is in progress and re-prepare with the desired update package.

CMMVC8269E **The attempt to prepare the cluster for update has failed because the previous update is in prepare_failed state. The previous update must first be aborted before reattempting the update.**

Explanation:
The current status of lsupdate reports the update as prepare_failed. This is an indication the user has already attempted to prepare an update, or started an update and in either scenario, the prepare has failed due to offline volumes. The cache flush has failed.

User response

The user needs to correct the error that caused the prepare to fail. Offline volumes are the most likely cause, also node resets might cause a prepare failure. Stop updates with prepare failures with the applysoftware -abort command, then reattempt the update.

CMMVC8270E **The applysoftware prepare timed out because an attempt to make the volume cache empty took too long. The command will be completed asynchronously. Use lsupdate to monitor the progress.**

Explanation:
The applysoftware prepare timed out because an attempt to make the volume cache empty took too long. The command will be completed asynchronously. Use lsupdate to monitor the progress. The state will be reported as prepared when successfully completed.

User response

Wait until the prepare completes and lsupdate reports prepared.

CMMVC8272E **Access iogrp parameter not valid when creating a file system volume.**

Explanation:
Access iogrp parameter not valid when creating a file system volume.

User response

Rerun mkvdisk without the -accessiogrp parameter or without the -filesystem parameter.

CMMVC8273E **Host cannot be removed as there is a volume that is accessible from multiple iogrps including one of the iogrps specified.**

Explanation:
If a volume is mapped to a host, it must be mapped in all of the iogrps in which it is accessible. The rmhostiogrp command fails if it would leave the system in this state.

User response

Use lshostvdiskmap to find the list of volumes that are mapped to the host in multiple iogrps. Then for each one either a) remove the host/volume mapping or b) remove the iorgp that the host is being removed from the volume's access iogrp set.

CMMVC8274E **The entry in the event log cannot be fixed because the given sequence number is out of range.**

Explanation:
The event log entry sequence number must be in the range 100 to 9,999,999 inclusive.

User response

Supply a valid event log entry sequence number in the range 100 to 9,999,999 inclusive.

CMMVC8275E **An entry with the given sequence number cannot be found in the event log.**

Explanation:

The fix request has failed because an entry with the given sequence number cannot be found in the event log.

User response

Supply a sequence number of an entry that exists in the event log.

CMMVC8276E **The entry in the event log cannot be fixed because it has expired or is in the monitoring state.**

Explanation:
The entry in the event log cannot be fixed because it has expired or is in the monitoring state.

User response

Expired and monitoring entries in the event log cannot be fixed.

CMMVC8277E **The MTM format must be XXXX-YYY where X is a numeric value, and Y is numeric or upper case character.**

Explanation:
The user has attempted to change the MTM but provided an incorrect format.

User response

Reissue the command with the MTM with the correct format. The format must be XXXX-YYY where XXXX are numeric values and YYY are alphanumeric characters. Any alphabetic characters must be upper case.

CMMVC8278E **The volume is accessible through more than one I/O group, and the host being mapped to the volume does not support volumes being mapped from multiple I/O groups.**

Explanation:
The volume is accessible through more than one I/O group, and the host being mapped to the volume does not support volumes being mapped from multiple I/O groups.

User response

Choose a different host or volume to map.

CMMVC8279E **The volume is accessible through more than one I/O group, and the host being mapped to the volume has an iSCSI name. iSCSI hosts do not support volumes being mapped from multiple I/O groups.**

Explanation:
The volume is accessible through more than one I/O group, and the host being mapped to the volume has an iSCSI name. iSCSI hosts do not support volumes being mapped from multiple I/O groups.

User response

Choose a different host or volume to map.

CMMVC8280E **The host has at least one volume mapped which is accessible through more than one I/O group, and the port being added is from a host system which does not support volumes being mapped from multiple I/O groups.**

Explanation

The host has at least one volume mapped which is accessible through more than one I/O group, and the port being added is from a host system which does not support volumes being mapped from multiple I/O groups.

Note: This error does not apply to release 7.6.0 and up.

User response

Choose a different port to add to the host.

CMMVC8281E **The host has at least one volume mapped which is accessible through more than one I/O group, and the port being added is from a host with an iscsi name. Iscsi hosts do not support volumes being mapped from multiple I/O groups.**

Explanation

The host has at least one volume mapped which is accessible through more than one I/O group, and the port being added is from a host with an iscsi name. Iscsi hosts do not support volumes being mapped from multiple I/O groups.

Note: This error does not apply to release 7.6.0 and up.

User response

Choose a different port to add to the host.

CMMVC8282E **At least one host mapped to the volume does not support volumes being mapped from multiple I/O groups.**

Explanation:
At least one host mapped to the volume does not support volumes being mapped from multiple I/O groups.

User response

Unmap the host which does not support access from multiple I/O groups.

CMMVC8283E **At least one host mapped to the volume has an iscsi name. Iscsi hosts do not support volumes being mapped from multiple I/O groups.**

Explanation:
At least one host mapped to the volume has an iscsi name. Iscsi hosts do not support volumes being mapped from multiple I/O groups.

User response

Unmap the host which does not support access from multiple I/O groups.

CMMVC8284E **The enclosure does not support setting a machine part number.**

Explanation

The user has attempted to set the machine part number on an enclosure that does not have a machine part number as part of its VPD, for example a 2076-112.

User response

Check the node panel name and try again.

CMMVC8285E **The machine type and model is not valid for this enclosure.**

Explanation

The user has attempted to change an enclosure machine type and model (MTM) to one that is not valid for the enclosure hardware, or one that is valid for the enclosure hardware but is not valid for the enclosure machine part number.

User response

Check the MTM and try again.

CMMVC8286E **The machine part number is not valid for this enclosure.**

Explanation

The user has attempted to change an enclosure's part number to one that is not valid for the enclosure hardware, or one that is valid for the enclosure hardware but is not valid for the enclosure's machine part number.

User response

Check the machine part number and try again.

CMMVC8287E **The machine part number and machine type and model are not valid for this enclosure.**

Explanation

The user has attempted to change an enclosure machine part number and machine type and model (MTM) to values that are not valid for the enclosure hardware.

User response

Check the MTM and the machine part number and try again.

CMMVC8289E **A new enclosure VPD field does not match the node copy.**

Explanation

The user has attempted to modify the enclosure serial number, machine part number, or machine type and model (MTM) for a replacement enclosure and the new value is not the value expected by the system.

User response

Check the service status view for the value the system is expecting and retry the command with the right value.

CMMVC8290E **The action failed because the requested combination of notification settings is not permitted.**

Explanation

The valid combinations of notification settings are info+warning+error and warning+error.

User response

Reissue the command using a valid combination of notification settings.

CMMVC8291E The command failed because it is not supported.

Explanation

The command entered is not supported on this platform. It might be supported on one of the other platforms and in the CLI help for those platforms.

User response

Check the command in the CLI guide; check the system is the intended one.

CMMVC8292E The command failed because a parameter is not supported.

Explanation

The user typed a command that is supported, but used a parameter that is not supported on this platform -- the parameter is supported on other platforms. For example the -mdisk parameter on the mkmdiskgrp command is not supported on this platform, but is supported on other platforms.

User response

Check the syntax in the CLI guide; check the system is the intended one.

CMMVC8293E The command failed because it is not supported for image-mode volumes.

Explanation

This error is returned by a migration command when issued against an image-mode volume on a platform that does not support migration for image-mode volumes.

User response

Perform these steps:

1. Use Volume Mirroring to change the volume's storage pool.
2. Add a volume copy in the desired storage pool.
3. Sync the volume.
4. Delete the first copy.

CMMVC8294E The command failed because the FlashCopy feature is not active and the maximum number of FlashCopy target volumes already exist.

Explanation

The user has already configured the maximum number of FlashCopy targets that are allowed without activating the FlashCopy licensed feature.

User response

Delete some FlashCopy targets or activate the FlashCopy licensed feature. Try the command again.

CMMVC8295E The command failed because a licensed feature is not activated.

Explanation

This command requires a licensed feature to be activated before it is used.

User response

Activate the licensed feature and try again.

CMMVC8296E The command failed because it is not supported for image-mode MDisks.

Explanation

This error is returned by the "remove mdisk" rmmdisk command when issued against an MDisk that is backing an image-mode volume on a platform that does not support migration for image-mode volumes.

User response

If the image-mode volume is not required then use rmvdisk to delete the volume. This deletes the MDisk as well. If the user wants to migrate the image-mode volume's data to internal storage, do that using Volume Mirroring and then delete the image-mode volume copy.

CMMVC8297E The drive cannot be managed because it has become unreachable.

Explanation

This message occurs if the user tries to change the use of the drive too soon after the drive has been inserted or the enclosure has been connected for the first time. It can also occur because of a hardware fault or if the user tries to change the use of a drive that has been removed from the system but is already managed.

User response

Wait and try again. If this doesn't work after 10 minutes, replace the drive. A drive that has been removed can still have its use changed to unused.

CMMVC8298E The system cannot open the file.

Explanation:
The file specified after the -file option cannot be opened.

User response

Refer to the documentation for the update file to ensure that the update file is correct.

Obtain a new copy of the correct package file, copy it to the system, then run the command again.

CMMVC8299E The system ran out of temporary resource while opening the file.

Explanation:
The temporary directory used to unpack the file cannot be opened.

User response

In the unlikely event of this error occurring, schedule either a node reboot or a node rescue maintenance task.

Once the maintenance task completes and the node is online, try the command again.

CMMVC8300E The specified file cannot be used because it contains too many drive software images.

Explanation:
The drive package contains more files than can be unpacked.

User response

Use a new package file and try the command again.

CMMVC8301E The system cannot read drive software from the specified file because the file contains an invalid drive firmware version string.

Explanation:
The drive package has probably become corrupt or has been made incorrectly.

User response

Refer to the documentation for the update file to ensure that the update file is correct.

Obtain a new copy of the correct package file, copy it to the system, then run the command again.

CMMVC8302E The system cannot read drive software from the specified file.

Explanation:
An internal error has occurred extracting the drive firmware from the drive package.

User response

Refer to the documentation for the update file to ensure that the update file is correct.

Obtain a new copy of the correct package file, copy it to the system, then run the command again.

CMMVC8303E The system can only program the FPGA of one drive at a time.

Explanation:
More than one drive has been specified and the -type option has been set to fpga.

User response

Change the command to only specify a single drive ID because svctask applydrivesoftware -type fpga only supports one drive by one command.

CMMVC8304E The system cannot apply the task because an earlier drive update task is still in progress.

Explanation:
The **svctask applydrivesoftware** command is already running, and only one instance of the command is allowed at a time.

User response

Try the command again when there is no drive update task in progress. Either use the **lsdriveupgradeprogress** command to determine when the updates complete, or use the **svctask applydrivesoftware -cancel** command to cancel the current drive update task.

CMMVC8305E The system cannot apply drive software to the specified drives because some of them are not online.

Explanation:
An svctask applydrivesoftware command has been issued, but some drives that have been specified are not online.

User response

Carry out problem determination to bring the drives online, then repeat the command. Alternatively, repeat the command but do not specify the drive ID of any drive that is not online.

CMMVC8306E **Some of the specified drives are offline. The system cannot apply drive software to an offline drive, even if the -force option is specified.**

Explanation:
An svctask applydrivesoftware command has been issued, but some drives that have been specified are offline.

User response

Ensure the drives specified are in the online or degraded state.

CMMVC8307E **None of the drives can be scheduled for drive software upgrade.**

Explanation:
None of the specified drives are in a suitable state to apply drive firmware updates.

User response

Ensure that the drive states meet all the requirements to permit drive software download before repeating the command.

CMMVC8309E **The task cannot be initiated because some of the specified drives have an unsupported drive technology.**

Explanation

The drive technology is the value of the **tech_type** field that is returned by the **svcinfo lsdrive** command. The following values are supported:

- `tier0_flash`
- `tier1_flash`
- `tier_enterprise`
- `tier_nearline`

Specifying a drive with any other technology type results in this error.

User response

Use the **lsdrive** command to determine which drives have an unsupported drive technology.

Repeat the command but do not include the drive ID of any drive with an unsupported drive type.

CMMVC8310E **The task cannot be applied to unused drives when multiple drives are specified.**

Explanation

Some of the specified drives cannot be upgraded since that are currently unused by the system.

User response

Repeat the command, but do not include the drive ID of any drive that is currently unused.

CMMVC8311E **The system cannot open the file.**

Explanation:
The file specified after the -file option cannot be opened.

User response

Refer to the documentation for the update file to ensure that the update file is correct.

Obtain a new copy of the correct package file, copy it to the system, then run the command again.

CMMVC8313E **The specified file cannot be used because it contains too many drive software images.**

Explanation:
The drive package contains more files than can be unpacked.

User response

Use a new package file and try the command again.

CMMVC8314E **The system cannot read drive software from the specified file because the file contains an invalid drive firmware version string.**

Explanation:
The drive package has probably become corrupt or has been made incorrectly.

User response

Refer to the documentation for the update file to ensure that the update file is correct.

Obtain a new copy of the correct package file, copy it to the system, then run the command again.

CMMVC8315E **The system cannot read drive software from the specified file.**

Explanation:
An internal error has occurred extracting the drive firmware from the drive package.

User response

Refer to the documentation for the update file to ensure that the update file is correct.

Obtain a new copy of the correct package file, copy it to the system, then run the command again.

CMMVC8316E **The system can only program the FPGA of one drive at a time.**

Explanation:
More than one drive has been specified and the -type option has been set to fpga.

User response

Change the command to only specify a single drive ID because svctask applydrivesoftware -type fpga only supports one drive by one command.

CMMVC8317E **The system cannot apply the task because an earlier drive update task is still ongoing.**

Explanation:
There is an existing svctask applydrivesoftware already running, and only one is allowed at a time.

User response

Try the command again when there is no drive update task in progress. Either use the lsdriveupgradeprogress command to determine when the updates will be completed, or use svctask applydrivesoftware -cancel to cancel the current drive update task.

CMMVC8318E **The system cannot apply drive software to the specified drives because some of them are not online.**

Explanation:
An svctask applydrivesoftware command has been issued, but some drives that have been specified are not online.

User response

Carry out problem determination to bring the drives online, then repeat the command. Alternatively, repeat the command but do not specify the drive ID of any drive that is not online.

CMMVC8319E **Some of the specified drives are offline. The system cannot apply drive software to an offline drive, even if the -force option is specified.**

Explanation:
The svctask applydrivesoftware command was issued, but some specified drives are offline.

User response

Ensure that the specified drives are in either online or degraded state, then reissue the command.

CMMVC8320E **None of the drives can be scheduled for drive software upgrade.**

Explanation:
None of the specified drives are in a suitable state to apply drive firmware updates.

User response

Ensure that the drive states meet all the requirements to permit drive software download before repeating the command.

CMMVC8321E **The '-all' or '-drive' option is required to specify the drive(s) you want to upgrade.**

Explanation

An attempt was made to apply software to one or more drives without specifying a drive.

User response

Retry the **applydrivesoftware** command and this time specify either the **-all** option or the **-drive** option followed by one or more colon-separated drive IDs.

CMMVC8323E **The task cannot be applied to unused drives when multiple drives are specified.**

Explanation

Some of the specified drives cannot be upgraded because the drives are currently unused by the system.

User response

Repeat the command, but do not include the drive ID of any drive that is currently unused.

CMMVC8324E **There are no drive software upgrades scheduled.**

Explanation

The command is not ongoing while user input `-cancel` option.

User response

No action is required.

CMMVC8325E **None of the specified drives needed to be upgraded or downgraded.**

Explanation

Every requested drive firmware is up-to-date in default. If the package is old, or contains no newer images for the drives on the package, the command cannot pass the level check.

User response

Ensure the package is correct or use `-allowreinstall` or `-allowdowngrade` option, then repeat the command.

CMMVC8326E **The task cannot be initiated because the drive use changed.**

Explanation

The command is stopped if user change the use of the drive, because some drive may be changed to "unused" while the command is ongoing.

User response

Check the use of the drive that you specified on the command line. If it is still appropriate to upload new firmware to the drive, repeat the command.

CMMVC8327E **The -allowreinstall and -allowdowngrade options cannot be used with option -type fpga.**

Explanation

When we applydrivesoftware fpga type drives, we don't permit to reinstall or downgrade drives.

User response

To download drive FGPA software, repeat the command but ensure the -allowreinstall and -allowdowngrade options are omitted.

CMMVC8328E **No package file is specified or invalid package file name is used.**

Explanation

User input invalid package file name.

User response

Input right package file name, and repeat the command.

CMMVC8329E **The task cannot be initiated because downloading to one or more drives could cause volumes to go offline.**

Explanation:
With any drive software update, there is a risk that the drive might become unusable. For systems that support RAID0 arrays, if the drive is a member of a RAID0 array consider whether to introduce more redundancy to protect the data on that drive.

User response:
If the drive is not a member of a RAID0 array, fix any errors in the event log that is related to the array. When the drive is a member of an array with sufficient redundancy, repeat the command. The **`-force`** option can be used to bypass the dependency checking but is not recommended. If a drive fails during an upgrade and the RAID array has no redundancy, the data on the array might be lost. To avoid this risk, restore array redundancy before attempting a drive software update.

CMMVC8330E **The %1 of %2 %3 in backup is %4; should be %5**

Explanation

The specified object has the specified property of the specified type with the specified incorrect value. The property most likely reflects the state of the object.

User response

Change the state to the required value, and resubmit the command.

CMMVC8331E **The command failed because at least one of the specified MDisks can only be used in image mode.**

Explanation

Image mode only MDisk is forbidden to be added into the storage pool.

User response

Do not include MDisks that can only be used in image mode in a storage pool.

CMMVC8332E **No MDisks were removed from the storage pool because at least one of the specified MDisks can only be used in image mode.**

Explanation

Cannot remove image mode only MDisk with this command.

User response

Use rmvdisk or rmvdiskcopy to remove the MDisks which could only be used in image mode.

CMMVC8333E **The task failed because it is not supported for image mode only MDisks.**

Explanation

Cannot migrate data away or to image mode only MDisk with this command.

User response

Use Volume Mirroring to migrate data from or to an image mode only MDisk.

CMMVC8335E **Cannot change attribute for IP address in an active partnership.**

Explanation

An attempt was made to change the attributes or the VLAN ID associated with an IP address where an active IP partnership was running on that address. You must stop the partnership before you change an attribute or a VLAN ID for this IP address.

User response

Stop the partnership by entering the **chpartnership -stop** command, then retry the original **cfgportip** command.

CMMVC8336E **Site was not specified. Site must be specified because topology is stretched.**

Explanation:
Site was not specified. Site must be specified because topology is stretched.

User response

Identify the site of the new node and resubmit the command with -site flag.

Or, change system topology.

Warning: Changing the topology will disable the DR feature.

CMMVC8337E **Site specified not valid. System topology is stretched and the other member has been configured to the same site.**

Explanation:
Site specified not valid. System topology is stretched and the other member has been configured to the same site.

User response

Identify a node in a different site to the existing node and resubmit.

Or, change system topology.

Warning: Changing the topology will disable the DR feature.

CMMVC8338E **Site parameter is not supported until the current update completes.**

Explanation:
Site parameter is not supported until the current update completes.

User response:
A node cannot be assigned to a site until the current update complete. Add the node without the -site parameter and configure sites after the update has completed.

CMMVC8339E **Not supported on this system.**

Explanation:
Not supported on this system.

User response:
The feature is not supported on this system. Wait for a later release.

CMMVC8340E **Cannot modify site because system topology is stretched.**

Explanation:
Cannot modify site because system topology is stretched.

User response

Set topology to standard then manipulate the site.

Warning: DR feature is not available when topology is standard.

CMMVC8341E **Site value is not valid. Can only specify site 1 or site 2.**

Explanation:
Site value is not valid. Can only specify site 1 or site 2.

User response:
Specify either site 1 or site 2.

CMMVC8342E **Cannot set stretched topology because some nodes do not have a configured site.**

Explanation:
Cannot set stretched topology because some nodes do not have a configured site.

User response:
Configure site for each node and then set topology.

CMMVC8343E **Cannot set stretched topology because some I/O groups have 2 nodes in the same site.**

Explanation:
Cannot set stretched topology because some I/O groups have 2 nodes in the same site.

User response:
Assign each node of the I/O group to a different site then set topology.

CMMVC8344E **Cannot change site because controller has one or more managed MDisks and the system topology is stretched or Hyperswap.**

Explanation

In stretched or HyperSwap system mode, it is not possible to change the site while the controller has one or more managed MDisks.

User response

Migrate or delete MDisks on the controller to make all MDisks unmanaged. Or, set the system topology to standard and then change site.

Note: The disaster recovery feature is not available with standard topology.

CMMVC8345E **Cannot change site because the controller is a SAS RAID Controller.**

Explanation

SAS RAID Controllers are not supported at this time.

User response

Use a FC connectivity controller.

CMMVC8346E **Cannot change quorum disk because there is already a quorum disk defined at that site, and the cluster topology is stretched or hyperswap.**

Explanation

Only one quorum disk is supported per site.

User response

Change either the identified current quorum disk ID, or a new quorum MDisk, so that they are not both in the same site.

Or, set the topology to standard.



Warning: Setting the topology to standard will disable the DR feature.

CMMVC8347E **Cannot modify rcauthmethod due to the presence of unstopped IP partnerships.**

Explanation

All IP partnerships must be stopped before rcauthmethod can be changed.

User response

Use chpartnership -stop to stop partnerships, then run the command again.

CMMVC8348E **Cannot modify chapsecret due to rcauthmethod set to chap and the presence of unstopped IP partnerships.**

Explanation

All IP partnerships must be stopped before rcauthmethod can be changed.

User response

Use chpartnership -stop to stop partnerships, then run the command again.

CMMVC8349E **The rc authentication method specified is not valid.**

Explanation

Value of rcauthmethod supplied is not 'chap' or 'none'.

User response

Correct the value of the parameter.

CMMVC8350E The command failed because the specified battery is offline.

Explanation

A command to adjust battery state cannot be completed because the specified battery is offline or has been removed.

User response

If the battery is offline but present, it is safe to remove without issuing a chnodebattery -remove command. To turn on the LED, the battery must be replaced and brought online.

CMMVC8351E The command has failed because the specified battery is not redundant.

Explanation

A request to prepare a battery for removal cannot be completed because the specified battery is not redundant.

User response

Remove the condition that is causing the lack of redundancy. Such conditions may include the partner battery being offline or not fully charged, or one of the boot drives being offline.

CMMVC8352E The task cannot be initiated because the download type is not valid.

Explanation

If -type is specified, only firmware or fpga is supported now.

User response

Check the input download type and repeat the command with a supported download type.

CMMVC8353E CHAP authentication failure

Explanation

The Partner Discovery has refused a discovery request because the CHAP secret that was specified is not correct.

User response

Correct CHAP secret must be provided.

CMMVC8354E Unreachable cluster IP address

Explanation

Incorrect partner cluster IP address parameter specified.

User response

Correct IP address of the cluster must be provided. IP address must be pingable.

CMMVC8355E Remote Copy port groups not configured or incorrectly configured.

Explanation

This error occurs when the administrator has not set up ethernet ports on any node of the local system with either remote copy port group 1 or 2. In addition, they may have been failed over or are offline.

User response

The administrator must execute the cfgportip CLI to configure one or more IP addresses on any one of the local nodes with remote copy port groups 1 or 2.

CMMVC8356E Remote Copy port groups not configured or incorrectly configured.

Explanation

This error occurs when the administrator attempts to set up more than one partnership of type IPv4 or IPv6.

User response

No action possible. The only option is to remove the existing partnership and create a new partnership.

CMMVC8357E Maximum number of allowed partnerships exceeded.

Explanation

This error occurs when the administrator attempts to set up more than 3 partnerships. A maximum of 3 FC partnerships, or 2 FC and 1 IP partnerships, can exist.

User response

No action possible. The only option is to remove one of the existing partnerships and create a new partnership.

CMMVC8358E **No local cluster IPs for the partnership type configured.**

Explanation

This error occurs when the administrator tries to create a partnership of type IPv4, but has not configured any IPv4 type cluster IPs on the local cluster. The same error is seen if the administrator tries to create a partnership of type IPv6, but has not configured any IPv6 type cluster IPs on the local cluster.

User response

The administrator must execute the cfgportip CLI to configure the local IP address depending on the type of IP partnership which is to be created.

CMMVC8359E **Partner already exists in candidate list. Cannot create partnership.**

Explanation

This error is seen when an attempt is made to create an IP partnership with a remote cluster that already appears in the candidate list displayed by the lspartnershipcandidate CLI.

User response

In this scenario, an FC link exists between the clusters. The adminstrator must run CLI mkfcpartnership to create partnership.

CMMVC8360E **Specified partner cluster IP address is used on local cluster. Cannot create partnership.**

Explanation

This error is seen when the administrator specifies a local IP address as the remote cluster IP address.

User response

A local IP address cannot be used to create an IP partnership. A remote IP address must be specified.

CMMVC8361E **All IP addresses of the partnership type are either down or not configured.**

Explanation

This error occurs when partner discovery is not reporting any matching remote ports. Example: The partnership type is IPv4, but all ethernet ports on the partner cluster are configured with IPv6 types (or vice versa). This can also happen when the partnership type is IPv4, but all IPv4 addresses configured on the partner cluster are offline.

User response

The adminstrator must specify the correct partnership type. This depends on the type of IP addresses configured on the remote cluster.

CMMVC8362E **The action failed because the cluster ID is not valid.**

Explanation

Partnership creation failed because the specified remote cluster ID is not valid.

User response

In case of CLI mkfcpartnership, the adminstrator should run CLI lspartnershipcandidate to check correct cluster id/name. In case of CLI mkippartnership, the adminstrator only specifies the IP address and clusterid is discovered. If this error occurs for mkippartnership, contact support.

CMMVC8363E **The remote cluster partnership was not created because it already exists.**

Explanation

This error occurs when an attempt is made to create a partnership with a cluster that is already in a partnership.

User response

No action. A partnership cannot be created with a cluster that is already in a partnership.

CMMVC8364E **Unsupported partnerships type specified.**

Explanation

This error occurs when the administrator specifies an unsupported type of partnership type for CLI chpartnership.

User response

The administrator must look for a possible option in CLI help, and specify the partnership type.

CMMVC8365E **Cannot change parameter if partnership is not in stopped state.**

Explanation

This error occurs when the administrator tries to modify partnership parameters without stopping the partnership.

User response

Administrator must run command chpartnership -stop *clusterid/name*

CMMVC8366E **Incorrect remote cluster IP specified.**

Explanation

This error occurs when the adminstrator running chpartnership specifies the type as IPv4, but specifies an IPv6 address value (or vice versa).

User response

Specify an IP address that is valid for the partnership type.

CMMVC8367E **Incorrect operation of FC partnership.**

Explanation

While running chpartnership, the CLI administrator species option -clusterip, -chapsecret, or -nochapsecret for FC partnership.

User response

No action. Options -clusterip, -chapsecret, and -nochapsecret are not valid for FC partnership

CMMVC8368E **Cannot set or reset attributes for unconfigured ports.**

Explanation

An attempt was made to complete one of the following actions when the corresponding IP address is not assigned:

- Change the attributes of a port.
- Add or remove the VLAN ID of a port.

User response

Configure the port by using the **cfgportip** command and specifying the missing IP address. You can change the port attributes, or add or remove the VLAN ID, as part of this same command, or make the changes in a separate **cfgportip** command.

CMMVC8369E **The action failed because the cluster ID is not valid.**

Explanation

The cluster ID used isn't valid for the command attempted. This could be because the cluster ID used doesn't correspond to any system, or it could be that the cluster ID used is that of the local system and the command attempted is only valid for remote systems.

User response

Use a correct cluster ID corresponding to a system that is valid for the command being attempted.

CMMVC8370E **The remote cluster partnership was not created because it already exists.**

Explanation

An attempt was made to create a partnership with a cluster that is already in a partnership.

User response

No action. A partnership cannot be created with a cluster that is already in a partnership.

CMMVC8371E **Unable to setup partnership due to partnership type mismatch.**

Explanation

This error occurs when the partnership types of local and remote systems don't match. Example: The partnership type specified on the remote system mkippartnership CLI invocation is done with IPv6, but the corresponding invocation on the local system is done with IPv4.

User response

Specify the same partnership type on both clusters.

CMMVC8372E **Unable to set up Remote Copy data paths with partner due to unavailability of matching Remote Copy port groups.**

Explanation

This error occurs when the partner returns Remote Copy configuration information that doesn't have matching Remote Copy port group IDs. This error could also occur in cases where matching local or remote ports are offline (link state is inactive), failed over to the partner node. Example, The local system has Remote Copy port group ID set to 1 while the partner has Remote Copy port group ID set to 2.

User response

Both clusters should have the same Remote Copy group port ID.

CMMVC8373E Easy Tier is active without a license for each enclosure.

Explanation

Easy Tier must not be active without a license for each enclosure.

User response

Ensure that a license exists for each enclosure.

CMMVC8374E The command failed because the source volume has pinned data.

Explanation

A FlashCopy mapping can only be started if there is not any pinned data on the source volume.

User response

Investigate why the source volume is offline. Fix the error and bring the volume back online. Try the command again.

CMMVC8375E The command failed because the volume has pinned data.

Explanation

The command cannot be completed because of pinned data on the volume.

User response

Investigate why the volume is offline. Fix the error and bring the volume back online. Try the command again.

CMMVC8376E Cannot change quorum disk because there is already a quorum disk defined at that site, and the cluster topology is stretched or hyperswap.

Explanation

In the stretched or Hyperswap system, every quorum disk should have a valid site to ensure there is only one quorum per site.

User response

Find an MDisk with a valid site without any other quorum disks existing in this site.

CMMVC8377E Cannot change the site for a quorum drive.

Explanation

In stretched system mode, drives as a quorum disk without a valid site are not supported.

User response

Find an MDisk with a valid site without any other quorum disks existing in this site.

CMMVC8378E Cannot modify cluster IP due to the presence of unstopped IP partnerships.

Explanation

This error will display when the administrator is trying to change a clusterip and there is an active IP partnership on the cluster.

User response

Once the Administrator stops the IP partnership, the cluster IP can be changed.

CMMVC8379E Partner state is stopped.

Explanation

This error occurs when creating an IP partnership with a remote cluster and the remote cluster partnership is partially_configured_local_stopped.

User response

Run chpartnership -start <Cluster ID> on the remote cluster.

CMMVC8380E Partner software version mismatch.

Explanation

This error occurs when an attempt is made to create a partnership with nodes that have incompatible versions of system software.

User response

None. A compatible software version is required.

CMMVC8381E The task cannot be cancelled as it has completed all drive downloads.

Explanation

After the multiple drive download task has completed, there is a 270 second delay. It is not necessary to cancel the task if it was performed within this period of time.

User response

Wait 270 seconds before trying the command again.

CMMVC8382E The system cannot apply the task because a delay has been imposed between commands. (About 270 seconds)

Explanation

When a previous applydrivesoftware task is completed, a 270 second delay is imposed between each command. This error message will appear if a new task is attempted during this waiting period.

User response

Wait 270 seconds before trying the command again.

CMMVC8383E System layer mismatch.

Explanation

This error occurs when creating an IP partnership between two clusters when both clusters in a different layer.

User response

The administrator should change the cluster layer so that they are the same. To do this, use command -svctask chsystem -layer<storage/replication>.

CMMVC8384E The command has failed because the specified fan module is offline.

Explanation

A command to adjust the fan module state can not be completed because the fan module specified is offline or has been removed.

User response

The offline fan module triggers a fan module replacement dmp. This procedure should be followed to replace the broken or missing fan module with a functional unit.

Note: Will be returned by chenclosurefanmodule if used on a fan module that has a current state of offline.

CMMVC8386E Sync operation not possible.

Explanation

Node boot drive sync operation is not possible.

User response

Check the can_sync field on the lsnodebootdrive view.

can_sync will be false if the drives are already in sync or when there are certain boot drive errors active (unsupported drive, wrong node, invalid drive contents).

In the event that drives are already in sync, do nothing.

If there are boot drive errors active, resolve the drive error before attempting to sync.

CMMVC8387E Partner cluster ID mismatch.

Explanation

The discovery of the partner cluster returned a cluster ID that was different than the ID that the partnership was working with.

User response

This error occurs in two scenarios:

- An attempt was made to run the **chpartnership** command where a remote cluster IP address was specified that is not part of the current partnership. In this case, retry the command with a valid cluster IP address.
- A T3 or T4 recovery on the remote cluster has changed the remote cluster ID. In this case, you must remove and re-create the IP partnerships and relationships:
  1. Stop the remote copy relationship by using the **stoprcrelationship** command.
  2. Delete the remote copy relationship by using the **rmrcrelationship** command.
  3. Delete the partnership by using the **rmpartnership** command.

4. Recreate the partnership by using the **mkippartnership** or the **mkfcpartnership** command.

CMMVC8389E [-size] is not a supported parameter for the specified mdiskgrp.

Explanation

Cannot change the size of a parent pool.

User response

Only the chmdiskgrp -size command can be used to change the size of a child pool.

CMMVC8390E Cannot set *easy_tier_option* for child storage pool, the value must be same as parent storage pool's easy tier setting.

Explanation

You cannot change the Easy Tier setting of a child storage pool.

User response

If you specified the wrong storage pool, retry the command using the correct storage pool.

If you want to change the Easy Tier setting of a parent pool, use the **chmdiskgrp -easytier** *easy_tier_option* command.

CMMVC8412E Cannot add mdisks to a child pool.

Explanation

The mdisk_group_id/_name specified in addmisk command must be a parent pool.

This error is reported from the CLI addmisk command.

User response

Specify a parent pool in the addmisk command.

CMMVC8415E The command cannot be initiated because the maximum supported number of drives already exists.

Explanation

Check the Configuration Limits page to understand the maximum drives supported on your system.

User response

Remove any unused drives and try again.

CMMVC8423E The command failed because the storage pool size is not a multiple of extent size.

Explanation

The child pool size must be a multiple of its extent size. For example, if the extent size is 256 MB, legal values for the storage pool size include 256 MB, 512 MB, 768 MB, and so on.

User response

Retry the command and specify a legal value for the storage pool size.

CMMVC8424E The command failed because the source or target storage pool is a child storage pool, and source and target are in different parent storage pools.

Explanation

An attempt was made to run a **migratevdisk** command that specified an invalid target **mdiskgrp** (child pool) parameter. Source and target storage pools must have the same parent storage pool.

User response

Retry the command and specify a valid target.

CMMVC8425E The command failed because the storage pool has child storage pools.

Explanation

An attempt was made to delete a parent storage pool that has child storage pools. This action is not permitted. To see how many child pools are associated with the parent, run the **lsmdiskgrp** command and look at the **child_mdisk_grp_count** field.

User response

To delete the parent storage pool, first remove all child storage pools by using the **rmmdiskgrp** command, then try again to remove the parent.

CMMVC8427E The command failed because the child storage pool can not be shrunk below its used capacity.

Explanation

The command failed because the child storage pool cannot shrink below its used capacity.

This error is reported from chmdiskgrp -size *new_size*.

User response

Specify a new size that is larger than the used capacity of the storage pool.

CMMVC8452E **The command failed because the specified parent mdiskgrp is a child pool.**

Explanation

The command failed because the child pool can not be shrunk below its used capacity.

This error is reported from chmdiskgrp -size **<new size>**.

User response

Specify a new size larger than its used capacity.

CMMVC8455E **The command cannot be initiated because the maximum supported number of drives already exists.**

Explanation

This error occurs when an attempt is made to manage more than the maximum number of supported drives. Check the Configuration Limits page to understand the maximum drives supported on your system.

User response

Remove any unused drives, and try again.

CMMVC8456E **At least one drive on the enclosure cannot find drive update package.**

Explanation

While updating the drives on the enclosure, the CLI returns this error if the drive update package filename is formatted incorrectly. Correct filename format: `mdisksw.`*`product_id`*`.firmware.`*`fw_level`*`.img`

User response

Check the .img file name in the .gpf file and retry the command.

CMMVC8457E **The mapping cannot be removed because the volume being unmapped has received I/O within the defined volume protection period.**

Explanation:
If volume protection is enabled, and the volume being unmapped has received I/O within the defined volume protection time period, then the command will fail.

User response:
Ensure no host I/O is being sent to the volume and ensure you have waited for the defined volume protection time period or disable volume protection.

CMMVC8458E **The volume cannot be removed because the volume has received I/O within the defined volume protection period.**

Explanation

If volume protection is enabled, and the volume has received I/O within the defined volume protection time period, then the command will fail.

This behaviour is not changed by the volume being mapped/unmapped.

The force flag does not affect this policing behaviour. For example, the force flag does not override the policing.

User response:
Ensure no host I/O is being sent to the volume and ensure you have waited for the defined volume protection time period or disable volume protection.

CMMVC8459E **The volume cannot be removed because the volume has received I/O within the defined volume protection period.**

Explanation

If the last volume copy is being deleted, volume protection is enabled, and the volume being deleted has received I/O within the defined volume protection time period, then the command will fail.

This behaviour is not changed by the volume being mapped/unmapped.

Deleting one (of two) volume copies is not affected.

The force flag does not affect this policing behaviour. For example, the force flag does not override the policing.

User response:
Ensure no host I/O is being sent to the volume and ensure you have waited for the defined volume protection time period or disable volume protection.

CMMVC8460E **The storage pool cannot be deleted because at least one volume in the pool has received I/O within the defined volume protection period.**

Explanation

If volume protection is enabled, and any volume in the pool being deleted has received I/O within the defined volume protection time period, then the command will be failed.

This behaviour is not affected by the force flag (the command would already be failed if there are any mdisks in the pool being deleted and the force flag was not used). For example, The force flag does not override the policing.

User response:

Ensure no host I/O is being sent to volumes in the storage pool and ensure you have waited for the defined volume protection time period or disable volume protection.

CMMVC8461E **The host cannot be removed because the host being deleted is mapped to volumes which have received I/O within the defined volume protection period.**

Explanation

If volume protection is enabled, and the host being deleted is mapped to any volume which has received I/O within the defined volume protection time period, then the command will fail.

This behaviour is not affected by the force flag (the force flag deletes the host even if it has any vdisk mappings). For example, the force flag does not override the policing.

If multiple hosts are mapped to the same volume, then the command will be allowed if the host that is being removed is already 'offline'. The protection will still prevent the last host being removed regardless of whether it is online or not (This is an effort to improve the policing behaviour for clustered hosts).

User response:

Ensure no host I/O is being sent to the volumes mapped to the host and ensure you have waited for the defined volume protection time period or disable volume protection.

CMMVC8462E **The host I/O group cannot be removed because at least one volume in the I/O group(s) being removed from the host has received I/O within the defined volume protection period.**

Explanation

If volume protection is enabled, and any volumes mapped to the host in the I/O group being removed from the host have received I/O within the defined volume protection time period, then the command will fail.

This behaviour is not affected by the force flag (the force flag deletes the host from I/O groups, even if it has any vdisk mappings in those I/O groups). For example, the force flag does not override the policing.

The same behaviour as rmhost, the policing will allow the I/O group if the host is already offline and at least one other host is mapped to the volume. It does not allow the io group to be removed if the host is the only host mapped to the volume and the volume is busy.

User response:

Ensure no host I/O is being sent to the affected volumes and ensure they have waited for the defined volume protection time period or disable volume protection.

CMMVC8463E **The port cannot be removed because the host is mapped to at least one volume which have received I/O within the defined volume protection period and the port is the last port associated with the host.**

Explanation

If volume protection is enabled and the host port being deleted is the last port for a host mapped to any volume which has received I/O within the defined volume protection time period, the command will fail.

The same behavior applies for **rmhost**: the policing will allow the last host port to be removed if the host is already offline and there is another host mapped to the volume.

User response:

Ensure no host I/O is being sent to the volumes and ensure you have waited for the defined volume protection time period or disable volume protection.

CMMVC8469E **The node cannot be added because it does not meet minimum hardware requirements.**

Explanation:

An attempt was made to add a node to the system that does not have enough cache RAM (memory) installed to run the current level of code.

User response

Either choose a different node to add, or upgrade the cache RAM (memory) on this node.

CMMVC8470E **The system update cannot be completed because it is not required.**

Explanation:
The user is attempting to complete a system update. The action is not permitted because the system is not in a state where this is required.

User response

None - the command is not appropriate for the system.

CMMVC8471E **The system update cannot be resumed because it is not stalled.**

Explanation:
The user is trying to resume a system update. The system update is not in the 'stalled' state so cannot be resumed.

User response

No further action is required.

CMMVC8472E **The action cannot be performed because paced update is not in progress.**

Explanation:
The user is requesting that the next node in the system be updated as part of a paced system update. The system is not performing a paced update.

User response

No further action is required.

CMMVC8473E **The next node cannot be updated because it is not ready for update.**

Explanation:
The user is requesting that the next node in the system be updated as part of a paced system update. The system has not finished updating the current node yet and is not ready to start updating another node.

User response

Wait until the current node update has completed and the next node is ready.

CMMVC8474E **The node cannot be updated because it would cause VDisks to go offline.**

Explanation:
You are requesting that the next node in the system be updated as part of a paced system update. If this node is updated then some volumes will temporarily go offline.

User response

Either address the issue causing the volumes not to be redundant, or resubmit the command using the -force option. If the -force option is used then some volumes will temporarily go offline.

CMMVC8475E **The node cannot be updated because it is offline.**

Explanation:
You are requesting that the next node in the system be updated as part of a paced system update. The node cannot be updated because it is offline.

User response

Either delete the node from the system or perform service actions to bring it back online.

CMMVC8476E **The node cannot be added because it would cause a paced update but the system has not completed its current update.**

Explanation:
You are trying to add a node that has been configured to perform a paced update. A system update is currently in progress and it is not a suitable time to begin a paced update.

User response

Wait until the system update has completed, then add the node again.

CMMVC8477E **The command failed because firmware applicable to a drive was not found in the package.**

Explanation

While updating drives on the enclosure, the CLI returns this error if an incorrect format was used for the file name of the drive update package.

The correct file name format is:

```
mdisksw.product_id.firmware.fw_level.img
```

User response

Check the .img file name in the .gpf file and retry the command.

CMMVC8478E **The mapping cannot be removed because the volume being unmapped has received I/O within the defined volume protection period.**

Explanation

Volume protection is enabled, but the volume being unmapped has received I/O recently, within the defined volume protection time period, so the unmapping command fails. Receiving I/O generally indicates that a volume is in use.

User response

1. Ensure that this volume is intended to be unmapped. If you selected the wrong volume, repeat the unmapping command with the correct volume.
2. To unmap this volume, ensure that no host I/O is sent to the volume, wait the time specified in the vdisk_protection_time field in the **lssystem** command since the last I/O was received, then retry the unmapping command.
3. To disable volume protection and its warnings, disable the vdisk_protection-enabled field in the **chsystem** command.

CMMVC8479E **The volume cannot be removed because the volume has received I/O within the defined volume protection period.**

Explanation

If volume protection is enabled and an attempt to delete the last volume copy occurs within the defined volume protection time period after the volume received I/O, the volume removal command fails. This policing behavior protects I/O integrity.

This policing behavior:

- Is not affected by the volume being mapped or unmapped
- Does not occur when deleting one of two volume copies
- Is not overridden by using the **force** flag

User response

1. Ensure that you intend to unmap the volume. If you selected the wrong volume, repeat the command with the correct volume.
2. To unmap the volume, ensure that no host I/O is sent to the volume, wait the time that is specified in the vdisk_protection_time field in the **lssystem** command, then retry the volume removal command.
3. Alternatively, to disable volume protection warnings and behavior, disable the vdisk_protection-enabled field using the **chsystem** command, then retry the volume removal command.

CMMVC8481E **The storage pool cannot be deleted because at least one volume in the pool has received I/O within the defined volume protection period.**

Explanation

If volume protection is enabled, and any volume in the storage pool being deleted received I/O within the defined volume protection time period, the storage pool deletion command fails. This is a policing behavior to protect I/O integrity.

This policing behavior is not overridden by the **force** flag.

User response

1. Ensure that no volumes are in use in the pool. If you selected the wrong pool, repeat the command with the correct pool.
2. To unmap all volumes in the pool, ensure that no host I/O is sent to any volume in the pool, wait the time that is specified in the vdisk_protection_time field in the **lssystem** command, then retry the storage pool deletion command.
3. Alternatively, to disable volume protection warnings and behavior, disable the vdisk_protection-enabled field using the **chsystem** command, then retry the storage pool deletion command.

CMMVC8482E **The host cannot be removed because the host being deleted is mapped to volumes which have received I/O within the defined volume protection period.**

Explanation

If volume protection is enabled, and the host being deleted is mapped to any volume that has received I/O within the defined volume protection time period, the host deletion command fails. This policing behavior protects I/O integrity.

This policing behavior:

- Does not occur when multiple hosts are mapped to the same volume, and the host being removed is already `offline`
- Is not overridden by using the **force** flag
- Prevents the last host being removed regardless of whether the host is online or not, which improves the policing behaviour for clustered hosts

User response

1. Ensure that you intend to remove the host. If you selected the wrong host, repeat the command with the correct host.
2. To remove the host, ensure that no host I/O is sent to any volume mapped to the host, wait the time that is specified in the vdisk_protection_time field in the **lssystem** command, then retry the host removal command.
3. Alternatively, to disable volume protection warnings and behavior, disable the vdisk_protection-enabled field using the **chsystem** command, then retry the host removal command.

CMMVC8483E **The host I/O group cannot be removed because at least one volume in the I/O group(s) being removed from the host has received I/O within the defined volume protection period.**

Explanation

If volume protection is enabled, but any volume mapped to the host in the removal candidate I/O group has received I/O recently (within the defined volume protection time period), the host I/O group removal command, **rmhostiogrp**, fails.

This behavior is not affected by the **force** parameter of the command. The force parameter deletes from a host the I/O groups that have volume mappings. The force parameter does not overrride the volume protection policing.

Removing a host I/O group behaves the same way as does removing a host in that the policing behavior allows the removal of an I/O group if the host is offline, and if one other host is mapped to the volume. The policing policy does not allow removing the I/O group if the host is the only host mapped to the volume and the volume is busy.

User response

1. Ensure that the host I/O group is intended to be unmapped. If you selected the wrong host I/O group , repeat the removal command with the correct host I/O group.
2. To unmap the volumes in the host I/O group and remove the I/O group, ensure that no host I/O is sent to the volumes, wait the time specified in the vdisk_protection_time field in the **lssystem** command since the last I/O was received, then retry the host I/O group removal command.
3. To disable volume protection and its warnings, disable the vdisk_protection-enabled field in the **chsystem** command.

CMMVC8484E **The port cannot be removed because the host is mapped to at least one volume which have received I/O within the defined volume protection period and the port is the last port associated with the host.**

Explanation

If volume protection is enabled, but the host port is the last port for a host mapped to any volume that has received I/O recently (within the defined volume protection time period), the **rmhostport** command fails.

The **rmhostport** command shares a similar behavior with the **rmhost** and the **rmhostiogrp**commands: volume protection policing allows removal of the last host port if the host is already offline and there is another host mapped to the volume. The policing policy does not allow removing the host port if the host is the only host mapped to the volume and the volume is busy.

This behavior is not affected by the force parameter of the command.

User response

1. Ensure that the host port is intended to be removed. If you selected the wrong host port, repeat the removal command with the correct host port.
2. To unmap the volumes in the host I/O group and remove the host port, ensure that no host I/O is sent to the volumes, wait the time specified in the vdisk_protection_time field in the **lssystem** command since the last I/O was received, then retry the host port removal command.
3. To disable volume protection and its warnings, disable the vdisk_protection-enabled field in the **chsystem** command.

CMMVC8485E **The command cannot be initiated because the drive that you have specified has not been formatted with the correct protection information scheme.**

Explanation:
The drive has not been formatted to the correct protection information scheme (Type 2), so initialization cannot be performed. This usually

implies a process step has been missed during the manufacturing of the drive.

User response:
If you are confident that you have no data on the drive, use chdrive -task format *<drive_id>* to reformat the drive to the correct protection information scheme. Then, retry the command.

CMMVC8518E The command failed because image-mode volumes are not supported in child storage pools.

Explanation

You cannot create an image mode volume from the child pools.

This error is reported by the **mkvdisk -mode image** command.

User response

Change the volume mode to **striped**.

CMMVC8519E The command failed because sequential-mode volumes are not supported in child storage pools.

Explanation

You cannot create a sequential mode volume from the child pools.

This error is reported by the **mkvdisk -mode seq** command.

User response

Change the volume mode to **striped**.

CMMVC8523E The command failed because the specified storage pool is a child storage pool.

Explanation

This error is reported by the **mkmdiskgrp -parentmdiskgrp** command. The **parentmdiskgrp** parameter must identify a parent storage pool. You cannot create a child pool from another child pool.

User response

Specify a parent storage pool for the **-parentmdiskgrp** parameter.

CMMVC8525E The command failed because mdisks cannot be added to a child storage pool.

Explanation

The **mdisk_group_id** or **mdisk_group_name** that was specified in the **addmdisk** command must be a parent pool.

User response

Specify a parent pool in the **addmdisk** command.

CMMVC8526E The command failed because mdisks cannot be removed from a child storage pool.

Explanation

The **mdisk_group_id** or **mdisk_group_name** that was specified in the **rmmdisk** command must be a parent pool.

User response

Specify a parent pool in the **rmmdisk** command.

CMMVC8528E The system does not support encryption.

Explanation

The hardware on this system does not support encryption.

User response

Use a system hardware that supports encryption.

CMMVC8529E The encryption function has not been activated.

Explanation

The system supports encryption but licenses have not been activated.

User response

Activate the encryption function on all I/O groups by installing license keys.

CMMVC8530E A rekey operation is in progress.

Explanation

Command can't be carried out whilst rekey operation is in progress.

User response

Run chencryption to commit or cancel current rekey operation.

CMMVC8531E The command failed because the last encryption method cannot be disabled while encrypted storage exists.

Explanation:
The user has attempted to disable an encryption method while there are encrypted arrays, encrypted pools, or encrypted cloud accounts on the system. Disabling the last encryption method while encrypted objects exists is not allowed, as it would mean permanently losing access to that data.

User response

Delete all encrypted objects and then run the command again.

CMMVC8532E The encryption function is not enabled.

Explanation

The encryption key cannot be set up until encryption is enabled.

User response

If your system supports encryption, enable encryption and try again. You can find more information by searching for "Enabling encryption".

CMMVC8533E No rekey operation in progress.

Explanation

No rekey operation to commit or cancel.

User response

Prepare key first and retry commit task.

No action required to cancel.

CMMVC8534E node [X] has insufficient entropy to generate key material.

Explanation

An internal error occurred during an attempt to generate encryption key material.

User response

Identify the node that is used for key generation, and then retry the command. If the command fails again, complete the following steps, retrying the command after each step until the command runs successfully.

1. Restart the node.
2. If the Trusted Platform Module (TPM) is replaceable, replace the TPM.
3. Replace the node.

CMMVC8535E The system does not have access to current encryption keys.

Explanation

The system cannot complete the action because it does not have access to the encryption keys.

User response

If the encryption key is stored on USB flash drives, then locate the USB flash drives and plug them into the system.

If the encryption key is stored on key servers, then ensure the key servers are online and can be accessed by the system.

If the encryption key is the recovery key, then run the **chencryption -recoverykey** command to provide the recovery key.

CMMVC8536E The system does not have sufficient USB flash drives attached.

Explanation

The system does not have sufficient USB flash drives attached. The number needed is displayed.

User response

Insert the specified number of USB flash drives needed into the system and retry.

CMMVC8537E Not enough USB copies made, needs (%1).

Explanation

The keys have not been written to the minimum number of USB devices.

User response

- Wait for the copies count to reach the minimum required
- Check for USB errors
- Replace flash drive and cancel. Then, retry prepare.

CMMVC8538E System is busy commit in progress.

Explanation

The system is busy. Commit in progress.

User response

Wait for I/O to complete and commit to complete. Then, use lsencryption to check state.

CMMVC8539E **Encryption is not supported on drives IO group**

Explanation

The drive selected is in an I/O group that does not support encryption.

User response

Select a drive in an I/O group that supports encryption.

CMMVC8540E **Key not accepted by SAS adapter because of an internal error.**

Explanation

Key not accepted by SAS adapter.

User response

Follow DMP for error shown in lsencryption.

CMMVC8541E **The command cannot be initiated because the object is being used by an automatic system configuration task.**

Explanation

This message appears when:

- A failed drive has been reseated or replaced. The system has automatically configured the device.

  A drive was automatically managed in the configuration as a replacement for a failed drive. This might have been due to user action of replacing the failed hardware with a new drive or reseating the old hardware on advice from a DMP.
- An attempt to automatically configure a reseated or replaced drive failed.

  A failed drive has been reseated or replaced. The system attempted to automatically configure the reseated or replaced drive for use in the system. This attempt failed.

User response

Wait until the drive auto manage process is complete. An informational or error message appears in the event log.

CMMVC8542E **The system update test utility has not been run before updating the system.**

Explanation:
Before you can update your system, you need to run the latest version of the test utility to verify that there are no issues with the current system.

User response

Run the system update test utility before you attempt a system update.

CMMVC8544E **The command cannot be initiated because the object is being used by an automatic system configuration task.**

Explanation:
The chenclosureslot command will no longer work while this new drive auto manage code is running.

User response

Wait until drive automanage finishes before running the chenclosureslot command.

CMMVC8549E **The command failed because the source volume is from a child pool.**

Explanation

The volumes of a child pool can be migrated only to its parent pool or to a child pool that comes from the same parent.

User response

Choose a target pool that has the same parent as the source volume and retry the command.

CMMVC8550E **Cluster ID alias cannot be changed whilst encryption is enabled.**

Explanation

Encryption requires that the cluster ID alias does not change while encryption is enabled.

User response

To change the cluster ID alias, disable encryption.

Once completed, re-enable encryption.

CMMVC8567E **A host of type adminlun is required to be mapped to all IO groups.**

Explanation

An attempt was made to change the host type to adminlun, but the host is not in all I/O groups.

User response

Add the host to all I/O groups by using the following command:

```
addhostiogrp -iogrpall host_id
```

You can now change the host type to adminlun.

CMMVC8570E **A subsidiary volume may only be mapped to hosts of type adminlun.**

Explanation

An attempt was made to map a subsidiary volume to a host that wasn't an adminlun type. Subsidiary LUN mappings are typically created or removed by the IBM Spectrum® Connect application. These mappings apply only to adminlun host types, and so it is unlikely that a storage administrator would encounter this message.

User response

Contact IBM Support.

CMMVC8571E **The volume is already mapped as a subsidiary volume.**

Explanation

An attempt was made to map a volume as a subsidiary volume, but the volume is already mapped as one. Subsidiary LUN mappings are typically created or removed by the IBM Spectrum Connect application. These mappings apply only to adminlun host types, and so it is unlikely that a storage administrator would encounter this message.

User response

Contact IBM Support.

CMMVC8573E **The action failed because there is one special volume owned by metadata volume in this storage pool.**

Explanation

The **rmmdiskgrp** command failed even though the **-force** parameter was specified. One special volume in the specified storage pool is owned by the metadata volume.

User response

Complete the following steps:

1. Run the **lsmetadatavdisk** command to find the special volume that is owned by the metadata volume.
2. Either use the **migratevdisk** or **addvdiskcopy**/**rmvdiskcopy** commands to move the special volume from the specified storage pool to a different storage pool. (If vVols or metadata volume or both are no longer needed, run the **rmmetadatavdisk** command to remove.)
3. Retry the **rmmdiskgrp** command.

CMMVC8574E **The action failed because some volumes with owner type vVol depend on the metadatavdisk and -ignorevvolsexist was not specified.**

Explanation

An attempt was made to remove a metadata volume that some other volumes depend on.

User response

Complete one of the following actions:

- Remove the vVol type volumes by using the **rmvdisk** command, then retry the **rmmetadatavdisk** command.
- Alternatively, specify the **-ignorevvolsexist** parameter when you retry the command. Using this parameter means that the metadata volume is removed, but the dependent vVol volumes are retained.

CMMVC8575E **The action failed because it depends on NTP server but NTP server is not configured.**

Explanation

During the creation of the metadata volume, an incorrect timestamp was found. The error was caused by the lack of a network time protocol (NTP) server, which is required for the metadata volume.

User response

Create the NTP server by using the following command:

```
chsystem -ntpip
```

After the NTP server is created, retry the **mkmetadatavdisk** command.

CMMVC8587E The command failed because the volume is fast formatting.

Explanation

The command cannot be run while the volume is fast formatting.

User response

Wait for the fast formatting process to complete, and then try the command again.

CMMVC8590E The command failed because the Compare Volume ID or name is invalid.

Explanation:
The specified Compare Volume ID or name is not valid.

User response:
Specify a valid Compare Volume ID or name and resubmit the command.

CMMVC8591E The command failed because the Base Volume ID or name is invalid.

Explanation:
The command failed because the Base Volume ID or name is invalid.

User response:
Specify a Base Volume ID or name and resubmit the command.

CMMVC8592E The command failed because the Compare Volume does not exist.

Explanation:
The command failed because the Compare Volume does not exist.

User response:
Specify the ID or name of an existing Compare Volume and resubmit the command.

CMMVC8593E The command failed because the Base Volume does not exist.

Explanation:
The command failed because the Base Volume does not exist.

User response:
Specify the ID or name of an existing Base Volume and resubmit the command.

CMMVC8594E The command failed because the Base and Compare Volumes are the same.

Explanation:
The command failed because the Base and Compare Volumes are the same.

User response:
Specify Base and Compare Volumes that are not identical and resubmit the command.

CMMVC8595E The command failed because the startlba of the scan segment is invalid.

Explanation:
The command failed because the startlba of the scan segment is invalid.

User response:
Specify a valid startlba and resubmit the command.

CMMVC8596E The command failed because the lbacount of the scan segment is invalid.

Explanation:
The command failed because the lbacount of the scan segment is invalid.

User response:
Specify a valid lbacount and resubmit the command.

CMMVC8597E The command failed because the chunk size is invalid.

Explanation:
The command failed because the specified chunk size is not valid.

User response:
Specify a valid chunk size and resubmit the command.

CMMVC8598E The command failed because the volume and the base volume are not in the same dependency chain.

Explanation:
The command failed because the volume and the base volume are not in the same dependency chain.

User response:
Specify base and compare volumes that are in the same dependency chain and resubmit the command.

CMMVC8599E The command failed because a node is offline.

Explanation:
A node is offline which caused the command to fail.

User response:
Bring the node back online and resubmit the command.

CMMVC8600E **The command failed because the write cache could not be flushed in time.**

Explanation:
The command failed because the write cache could not be flushed in time.

User response:
Retry the command.

CMMVC8601E **The command failed because the start of the scan segment is not chunk aligned.**

Explanation:
The command failed because the start of the scan segment is not chunk-aligned.

User response:
Specify a startlba that is chunk-aligned and resubmit the command.

CMMVC8620E **The active-active relationship could not be created because the master and auxiliary volumes do not have different well-defined sites.**

Explanation:
The active-active relationship could not be created because the master and auxiliary volumes do not have different well-defined sites.

User response

Contact IBM Support.

CMMVC8621E **The active-active relationship could not be created because the system topology is not HyperSwap.**

Explanation:
The active-active relationship could not be created because the system topology is not HyperSwap.

User response

Contact IBM Support.

CMMVC8622E **The active-active relationship could not be created because the master and auxiliary volumes are not in the same system.**

Explanation:
The active-active relationship could not be created because the master and auxiliary volumes are not in the same system.

User response

Contact IBM Support.

CMMVC8623E **I/O group was not specified. I/O group must be specified because topology is HyperSwap.**

Explanation:
You cannot set HyperSwap topology without specifying the I/O group.

User response

Specify the I/O group when using HyperSwap technology.

CMMVC8624E **Cannot set Hyperswap topology because some nodes do not have a configured site.**

Explanation:
Cannot set HyperSwap topology because some nodes do not have a configured site.

User response

Configure the site for each node and then set the topology.

CMMVC8625E **Cannot set HyperSwap topology because an I/O group has nodes assigned to different sites.**

Explanation:
You cannot set HyperSwap topology when the I/O group has nodes assigned to different sites.

User response

Assign both nodes in the I/O group to the same site or do no attempt to set HyperSwap.

CMMVC8626E **Cannot modify site because topology is HyperSwap.**

Explanation:
You cannot change the site of a node because the topology is set to HyperSwap.

User response

Change the system topology to `standard` to change the site of the node.

Warning: Changing the topology disables the disaster recovery (DR) feature.

CMMVC8627E **Site was not specified. Site must be specified because topology is Hyperswap.**

Explanation:

Site was not specified. Site must be specified because topology is HyperSwap.

User response

Identify the site and resubmit command with -site flag. Alternatively change system topology.

Warning: Changing the topology disables the disaster recovery feature.

CMMVC8628E **The host mapping was not created because the volume is not the host accessible volume in an active-active relationship.**

Explanation:
A volume cannot be mapped to a host if it is an auxiliary volume in an active-active relationship.

User response

Create the host mapping to the master volume in the relationship.

CMMVC8629E **Site specified is not valid. System topology is Hyperswap and the other member of the IO group has been configured to a different site.**

Explanation:
Site specified is not valid. System topology is HyperSwap and the other member of the I/O group has been configured to a different site.

User response

Identify the site of the other member of the I/O group. Assign the same site and resubmit. Alternatively, change the system topology.

Warning: Changing the topology will disable the DR feature.

CMMVC8630E **Site specified is not valid. System topology is Hyperswap and the node or control enclosure was previously configured with a different site.**

Explanation:
Site specified is not valid. System topology is HyperSwap and the node or control enclosure was previously configured with a different site.

User response

Identify the site specified originally for this node or control enclosure and resubmit the command. Or, change the system topology.

Warning: Changing the topology disables the disaster recovery feature.

CMMVC8631E **The change VDisk could not be disassociated because it is configured in an active-active relationship.**

Explanation:
A change volume cannot be disassociated from an active-active relationship.

User response

Contact IBM Support.

CMMVC8632E **The copy type of an active-active relationship or group cannot be changed.**

Explanation:
The copy type of an active-active relationship or group cannot be changed.

User response

Contact IBM Support.

CMMVC8633E **The relationship could not be added to the consistency group because it does not have an up-to-date copy in the same site as a set of up-to-date copies in the group.**

Explanation:
When adding an active-active relationship to a consistency group, the relationship must have an up-to-date copy in the same site as a set of up-to-date copies of each relationship in the consistency group.

User response

Before attempting to add the relationship to the consistency group, verify that the copy in the candidate relationship is as up-to-date as the set of up-to-date copies of any existing relationship in the consistency group.

CMMVC8634E **The host mapping was not created because the host does not have a site defined and the volume is in an active-active relationship.**

Explanation:
A volume in an active-active relationship can only be mapped to a host if the host has a site defined.

User response

Set the site for the host using the management GUI or the lshost and chhost CLI commands.

CMMVC8635E **Cannot un-set the host site as the host is mapped to a volume in an active-active relationship.**

Explanation:
A host mapped to a volume in an active-active relationship must have a site defined.

User response

Keep the site defined or remove the host mapping if un-setting the host site is required.

CMMVC8636E **The active-active relationship was not created because one or more hosts mapped to the master volume do not have a site defined.**

Explanation:
A host mapped to a volume in an active-active relationship must have a site defined.

User response

Set the site of all hosts that are mapped to the master volume before resubmitting the command.

CMMVC8637E **The active-active relationship could not be created because there are one or more host mappings to the auxiliary volume.**

Explanation:
The active-active relationship could not be created because there are one or more host mappings to the auxiliary volume.

User response

Contact IBM Support.

CMMVC8638E **An active-active relationship or group may not be manually started unless it has a state of idling and access has previously been enabled to an out-of-sync consistent copy.**

Explanation:
An active-active relationship or group may not be manually started unless it has a state of idling and access has previously been enabled to an out-of-sync consistent copy.

User response

Contact IBM Support.

CMMVC8639E **An active-active relationship or group may not be manually stopped unless it has a state of consistent_copying and access is required to the out-of-sync consistent copy.**

Explanation:
An active-active relationship or group may not be manually stopped unless it has a state of consistent_copying and access is required to the out-of-sync consistent copy.

User response

Contact IBM Support.

CMMVC8640E **Active-active relationships and groups alter their direction automatically and cannot have their direction switched manually.**

Explanation:
Active-active relationships and groups alter their direction automatically and cannot have their direction switched manually.

User response

Contact IBM Support.

CMMVC8642E **The command failed because the copy specified is not synchronized. The copy must be synchronized to create an active-active relationship.**

Explanation:
Only volumes where both copies are synchronized can be split when creating an active-active relationship. The command failed because the specified copy is not synchronized.

User response

Contact IBM Support.

CMMVC8643E **The active-active relationship could not be created because the existing and new volumes do not have different well-defined sites.**

Explanation:
The active-active relationship could not be created because the existing and new volumes do not have different well-defined sites.

User response

Contact IBM Support.

CMMVC8644E **The command failed because the auxiliary volume does not contain**

an up-to-date copy or have host access enabled.

Explanation:
The master volume in an active-active relationship (keeping host access to the auxiliary volume) can only be removed when the auxiliary volume is already processing host I/O. Either it is an up-to-date copy of an active-active relationship or the copy has host access enabled with the "stoprcrelationship -access" command.

User response

Contact IBM Support.

CMMVC8646E **The active-active relationship could not be created, either because the master or auxiliary volume is the target of a FlashCopy mapping or because the master or auxiliary volume is the source of a FlashCopy mapping with a target volume in a different site.**

Explanation

The active-active relationship could not be created because:

- The master or auxiliary volume is the target of a FlashCopy mapping.
- The master auxiliary volume is the source of a FlashCopy mapping with a target volume in a different site.

User response

Contact Support.

CMMVC8649E **Cannot change topology because active-active relationships are defined, and only supported in Hyperswap topology.**

Explanation:
It is not possible to change the topology because active-active relationships are defined and only supported in a HyperSwap topology.

User response

Remove all active-active relationships before changing the topology.

CMMVC8650E **System topology is Hyperswap or stretched, and the site of the MDisk being added does not match that of the storage pool.**

Explanation:
The system topology is either HyperSwap or stretched and the site of the MDisk you are adding does not match the storage pool.

User response

Do one of the following:

- Add the MDisk to an empty storage pool or to one already containing MDisks of a matching site.
- Remove and re-add the controller of the MDisk in the desired site.

Then, retry the command.

CMMVC8652E **The command failed because the volume is owned and has restricted use.**

Explanation

You cannot specify this command for this volume because the volume is owned and has restricted use.

User response

Choose a volume that is not owned with restricted use.

CMMVC8653E **The command failed as an MDisk group is owned and has restricted use.**

Explanation:
An MDisk group specified in the command is owned and either the action is not allowed for this type of owner or the user requires a specific role.

User response:
Check if the MDisk group is owned and then make sure the command is permitted for this type of owner or the required role is being used.

CMMVC8654E **The storage pool specified is not valid. The volume is a participant in an active-active relationship, and the storage pool for the new copy is in a different site to the volume's current site.**

Explanation

Participants in an active-active relationship must be located in the same site.

User response

Move the volume to a storage pool in the same site, or move the other copy in the active-active relationship to the target storage pool.

CMMVC8655E **The volume being moved is a participant in an active-active relationship, and the target storage pool is in a different site to the volume's current site.**

Explanation

The members of an active-active relationship must be located on the same site.

User response

Move the volume to a storage pool in the same site, or move the other copy in the active-active relationship to the target storage pool.

CMMVC8657E **The command failed as the specified storage pool is offline.**

Explanation

The storage pool that was specified in the command is offline.

User response

1. Determine why the storage pool is offline. You can use the **lseventlog** command to determine the event that caused the problem.
2. Correct the problem that is shown in the output of the **lseventlog** command.
3. Retry the **mkmetadatavdisk** command.

CMMVC8658E **The create failed because there are no iogroups supporting compression.**

Explanation

An attempt was made to create a compressed volume where the **-iogrp** parameter was not specified and none of the candidate I/O groups for the command support compression.

User response

Retry the command and use the **-iogrp** parameter to specify an I/O group that supports compression. Alternatively, retry the command without the **-compressed** parameter.

CMMVC8659E **The Create FlashCopy mapping task cannot be initiated because the target volume is in an active-active relationship.**

Explanation

An attempt was made to create a FlashCopy mapping with a HyperSwap volume as a target. HyperSwap volumes cannot be used as targets for FlashCopy mappings. This rule applies to both the master and auxiliary volumes of the active-active relationship.

User response

If you must use FlashCopy to create a point-in-time copy of a volume onto a HyperSwap volume, that target HyperSwap volume must be reconfigured as a regular volume until the FlashCopy process completes. Use the management GUI or the **rmvolumecopy** command to delete one copy of the target HyperSwap volume before you create the FlashCopy mapping.

CMMVC8660E **The Create FlashCopy mapping task cannot be initiated because the source volume is in an active-active relationship and the target volume and map are not in the same site as the source volume.**

Explanation

HyperSwap volumes contain a copy on each site, and copying data from one site to the other introduces significant unnecessary data movement. When you create a FlashCopy mapping to copy data on a HyperSwap volume, the source volume of the FlashCopy mapping must be the HyperSwap volume copy that is on the same site as the target volume.

User response

Retry the command, specifying as a source volume the HyperSwap volume copy that is on the same site as the target volume.

CMMVC8661E **The command has failed because the specified volume is a target of a FlashCopy mapping with a source volume in an active-active relationship and the new I/O group is in a different site to the source volume.**

Explanation

An attempt was made to move the preferred node of a volume to an I/O group on a different site when the volume is a target of a FlashCopy mapping with a source volume in an active-active relationship. For FlashCopy mappings that copy data from a HyperSwap volume, the source volume of the FlashCopy mapping must be on the same site as the target volume.

User response

Retry the command and move the volume to another I/O group on the same site. Alternatively, to move the volume to an I/O group on the other site, remove the FlashCopy mapping or convert the HyperSwap source volume to a regular volume by using the management GUI or the **rmvolumecopy** command first.

CMMVC8662E **The command has failed because the specified volume is a target of a FlashCopy mapping with a source volume in an active-active relationship and the target storage pool is in a different site to the source volume.**

Explanation

HyperSwap volumes contain a copy on each site, and copying data from one site to the other introduces significant unnecessary data movement. For a FlashCopy mapping that copies data from a HyperSwap volume, the source volume of the FlashCopy mapping must be on the same site as the target volume.

User response

Retry the command and move the volume to a storage pool on the same site. Alternatively, to move the volume to a storage pool on the other site, convert the HyperSwap source volume to a regular volume by using the management GUI or the **rmvolumecopy** command first.

CMMVC8663E **The command has failed because the specified volume is a target of a FlashCopy mapping with a source volume in an active-active relationship and the storage pool for the new copy is in a different site to the source volume.**

Explanation

HyperSwap volumes contain a copy on each site, and copying data from one site to the other introduces significant unnecessary data movement. For a FlashCopy mapping that copies data from a HyperSwap volume, the source volume of the FlashCopy mapping must be entirely on the same site as the target volume. This command would spread the volume across both sites.

User response

Complete one of the following actions:

- Delete the FlashCopy mapping and retry the command.
- Move the FlashCopy target volume to the same site as the source volume:
  1. Use the **movevdisk** command to change the I/O group for the target volume to one on the same site as the source volume.
  2. Use the **migratevdisk** command to change the storage pool for the target volume to a storage pool on the same site as the source volume.

  Then retry the command.

CMMVC8664E **The change volume could not be associated because it does not have the same well-defined site as the associated volume in the relationship.**

Explanation

Change volumes support HyperSwap volumes by capturing consistent data during synchronization. This process allows access to the consistent data if the remote site is lost. The change volume must be on the correct site for the volume in the active-active relationship it is being associated with.

User response

Configure the change volume in an I/O group and storage pool on the same site as the volume in the active-active relationship you are associating it with.

CMMVC8665E **The FlashCopy mapping was not started because the source volume is in an active-active relationship and does not contain an up-to-date copy or have host access enabled.**

Explanation

HyperSwap volumes contain a copy on each site, and copying data from one site to the other introduces significant unnecessary data movement. For a FlashCopy mapping that copies data from a HyperSwap volume, the HyperSwap volume must have an up-to-date copy on the same site as the target volume of the FlashCopy mapping.

User response

Allow the HyperSwap volume to synchronize fully before you start the FlashCopy mapping. Alternatively, create and use a FlashCopy mapping on the same site as the up-to-date copy of the HyperSwap volume.

CMMVC8666E **The FlashCopy consistency group was not started because one or more of the source volumes are in an active-active relationship and do not contain an up-to-date copy or have host access enabled.**

Explanation

HyperSwap volumes contain a copy on each site, and copying data from one site to the other introduces significant unnecessary data movement. For a FlashCopy consistency group that has one or more maps that copy data from a HyperSwap volume, the HyperSwap volume must have an up-to-date copy on the same site as the target volume of the FlashCopy mapping.

User response

Allow the HyperSwap volume to synchronize fully before you start the FlashCopy mapping. Alternatively, create and use a FlashCopy mapping on the same site as the up-to-date copy of the HyperSwap volume.

CMMVC8667E **The command failed because a specified host port has more than four logins on a node.**

Explanation

At least one host port (WWPN or IQN) has more than four logins to the same node. The network or SAN might not be zoned correctly. The system supports up to four logins per node from the same host port.

User response

Complete the following steps. If at any point you need more assistance, contact your service support representative.

1. Create a list of the problem hosts, WWPNs, and nodes.
   a. Run the **`svcinfo lsfabric -host`** command and parse the output into a human-readable format.
   b. Sort by WWPN, then by node.
   c. For any WWPN and node combination that shows more than four logins, complete the following steps:
      i) Get the host port mask from the mask field of the **`lshost`** detailed view.
      ii) Ignore any row where the local_port field does not match the appropriate bit in the host port mask.
      iii) Make a note of any hosts that still show more than four logins after the host port mask is applied.
2. Fix the issue either by changing the zoning or by putting the host and host ports in a portset such that there cannot be more the four logins from host WWPN to SVC node WWPN.

CMMVC8668E **The host mapping was not created because the SCSI logical unit number (LUN) ID is not valid for this particular host type.**

Explanation

An attempt was made to create a mapping between a volume and an adminlun host, but the SCSI LUN ID is not within the range for adminlun hosts. The highest SCSI LUN ID that is permitted for an adminlun host is 512.

User response

Retry the command with a valid SCSI LUN ID.

CMMVC8670E **Cannot modify host type since it has subsidiary mappings.**

Explanation

An attempt was made to change the host type from adminlun to another host type, but the adminlun host has subsidiary mappings. Only adminlun hosts can have subsidiary mappings.

User response

Contact a VMware administrator to ensure that all Virtual Volume datastores are unmounted from the specified host because this operation removes all remaining subsidiary mappings. After the process is complete, retry the command.

CMMVC8671E **An existing volume is mapped with a SCSI ID which exceeds the maximum for this host type.**

Explanation

An attempt was made to change the host type to adminlun, but an existing volume is mapped with a SCSI ID that exceeds the maximum SCSI ID for an adminlun host. The SCSI ID for an adminlun host type must be below 512.

User response

Ensure that all volumes that are mapped to the host have a SCSI ID below 512, and then change the host

type to adminlun. If not all volumes that are mapped to the host have a SCSI ID below 512, then remove mappings with a SCSI ID of 512 or above. If a mapping is still needed, re-create the mapping with a SCSI ID that is below 512.

CMMVC8672E The command failed because the access IO group set of a subsidiary volume may not be modified.

Explanation

An attempt was made to modify the access I/O group set for a subsidiary volume, which is required to be available through all I/O groups.

User response

You cannot modify access of a subsidiary volume. Retry the command for a different volume.

CMMVC8673E The volume is already mapped as a SCSI LUN which must first be unmapped/removed in order to create a subsidiary volume mapping.

Explanation

An attempt was made to create a subsidiary volume mapping, but the volume is already mapped as a SCSI LUN. Subsidiary LUN mappings are typically created or removed by the IBM Spectrum Connect application. These mappings apply only to adminlun host types, and so it is unlikely that a storage administrator would encounter this message.

User response

Contact IBM Support.

CMMVC8674E The volume is not mapped as a subsidiary LUN.

Explanation

An attempt was made to remove a subsidiary mapping between an adminlun host and a volume, but the volume was not a subsidiary volume.

User response

Specify a valid volume ID and retry the command.

CMMVC8676E The command failed because the volume is not a master volume in an active-active relationship.

Explanation:

This command retains access to a HyperSwap volume when deleting the master volume of its active-active relationship. It is not necessary or permitted when deleting the auxiliary volume of its active-active relationship, or any non-HyperSwap volume.

User response

Retry the command on the correct object type.

CMMVC8677E Retaining access to the auxiliary VDisk failed because either the master or the auxiliary VDisk is in mirrored mode.

Explanation:
The master volume in an active-active relationship can only be removed, keeping host access to the auxiliary volume, when both the master and auxiliary volumes are in non-mirrored volume.

User response

Remove additional mirrors of the master and auxiliary volumes of the active-active relationship so each has only one mirror, then retry the command.

CMMVC8696E The command failed because the volume or a volume copy is in the deleting state.

Explanation

An action was requested on a volume or a volume copy that was in the process of being deleted.

User response

Make sure that you specified the correct volume. If so, the requested action is not supported on volumes or volume copies that are deleting.

CMMVC8697E The command failed because the source or target volume is in the deleting state.

Explanation

An attempt was made to delete a volume that was already in the process of being deleted.

User response

Complete the following actions:

1. Make sure that you specified the correct volume.
2. If so, wait until the previous **rmvdisk** command completes.

3. If you believe that the previous command failed, and you are sure that you want to proceed, you can retry the command with the **-force** parameter. Using the **-force** parameter bypasses all built-in safety checks.

CMMVC8698E The command failed because the volume has host mappings, use -removehostmappings.

Explanation

An attempt was made to delete a volume that has host mappings.

User response

Make sure that you specified the correct volume. If so, and you want to delete the host mappings along with the volume, retry the command and include the **-removehostmappings** parameter.

CMMVC8699E The command failed because the volume is image mode and has FlashCopy maps, either remove the FlashCopy maps first or use -force.

Explanation

An attempt was made to delete an image mode volume that has FlashCopy maps.

User response

Complete the following steps:

1. Make sure that you specified the correct volume.
2. If so, remove the FlashCopy maps by using the **rmfcmaps** command and retry the **rmvdisk** command.
3. Alternatively, you can retry the **rmvdisk** command with the **-force** parameter. Using the **-force** parameter bypasses all built-in safety checks.

CMMVC8706E The command failed because the primary copy's autodelete flag is on, meaning type conversion is in progress.

Explanation

Changing the primary copy (using the **chvdisk -primary** command) is disabled when the primary copy's autodelete flag is on.

User response

Change the primary copy when the primary copy's **autodelete** is **no**.

Either wait for the type conversion to finish or cancel the process, if required, by using the **rmvdiskcopy** command.

CMMVC8707E The command failed because the caching IO groups for the volume are not in two independent sites.

Explanation

When you create HyperSwap volumes, the caching I/O groups must be located in both site 1 and site 2.

User response

Ensure that you have caching I/O groups in both site 1 and site 2, then retry the command.

CMMVC8708E The command failed because the system does not have a hyperswap topology.

Explanation

An attempt was made to create a HyperSwap volume on a system that does not have a HyperSwap topology. The **mkvolume** command attempts to create a HyperSwap volume in either of the following circumstances:

- Two I/O groups are specified in the command
- Two storage pools are specified within two independent sites that have a HyperSwap topology. In this case the **iogrp** parameter can be omitted and the system will auto-select an I/O group.

User response

Complete one of the following actions:

- If you entered the command incorrectly, correct the command and retry.
- Enter the **mkvolume** command with a single **-iogrp** parameter to create a basic volume.
- Configure the system to use a HyperSwap topology and retry the command.

CMMVC8709E The command failed because the caching IO groups are not in the same sites as the storage pools.

Explanation

When you create volumes under a HyperSwap topology, the caching I/O groups must be in the same sites as the storage pools.

User response

Ensure that the caching I/O groups are in the same sites as the storage pools, then retry the command.

CMMVC8710E The command failed because there were not enough extents in storage pool *storage_pool*.

Explanation

Volumes are created from the extents that are available in a storage pool. An attempt was made to create a volume when not enough extents were available.

User response

If you specified the correct storage pool, complete one of the following tasks:

- Reduce the size of the volume you are creating.
- Add more objects to the pool.

Then, retry the command.

CMMVC8711E The command failed because too many IO groups were specified.

Explanation

The maximum number of I/O groups is 2.

User response

Retry the command, specifying 2 or fewer I/O groups.

CMMVC8712E The command failed because too many storage pools were specified.

Explanation

An attempt was made to create a volume or volume copy that referenced more than two storage pools. The following limits are enforced:

- For the **mkvdisk** command: two mirrored copies.
- For the **mkvolume** command with standard or stretched topology: two mirrored copies.
- For the **mkvolume** command with HyperSwap topology: two volume copies.

User response

Retry the command with no more than two storage pools specified.

CMMVC8713E The command failed because the storage pool must be in site 1 or 2.

Explanation

When you create a volume under a stretched or HyperSwap topology, the storage pool must be located in either site 1 or site 2.

User response

Ensure that your storage pool is in either site 1 or site 2, then retry the command.

CMMVC8714E The command failed because the number of IO groups specified was different to the number of storage pools specified.

Explanation

When you create HyperSwap volumes, the number of I/O groups that you specify must equal the number of storage pools that you specify.

User response

Retry the command, specifying an equal number of I/O groups and storage pools.

CMMVC8715E The command failed because the storage pools for the volume are not in two independent sites. Volumes can only be created using storage in site 1 and site 2.

Explanation

When you create stretched or HyperSwap volumes with two storage pools, the pools must be located in independent sites, that is, in both site 1 and site 2.

User response

Ensure that you have storage pools in both site 1 and site 2, then retry the command.

CMMVC8716E The create failed because there are no iogroups in site *site_number*.

Explanation

When no I/O group is specified, the system automatically selects an I/O group for the volume in

the same site as the MDisk or storage pool that was specified in the command. However, an attempt was made to create a volume in a site where no non-empty I/O groups existed.

User response

Add at least one non-empty I/O group to the site and then retry the command.

CMMVC8717E The create failed because there are no iogroups supporting compression in site *site_number*.

Explanation

When no I/O group is specified, the system selects the I/O group for the volume automatically. No non-empty I/O groups that support compression were found in the site of the specified MDisk or storage pool.

User response

Complete one of the following tasks:

- Specify an existing I/O group that supports compression.
- Specify an MDisk or storage pool in a site with a non-empty I/O group that supports compression.
- Reconfigure your system so that the specified MDisk or storage pool is in a site that contains a non-empty I/O group that supports compression.

CMMVC8718E The command failed because the managed disk must be in site 1 or 2.

Explanation

Under a stretched or a HyperSwap topology, the managed disk must be in site 1 or site 2.

User response

Ensure that the managed disk is in site 1 or site 2 and retry the command.

CMMVC8719E The command failed because one or more service IP addresses aren't set.

Explanation:
The quorum application requires a connection to all nodes in the cluster by using the service IP addresses.

User response

1. Determine which IP version needs to be used by the quorum application to connect to the cluster: IPv4 (default) or IPv6. To change the default IP version, use the **mkquorumapp -ip_v6** command.
2. The service IPs of the IP version you are not using might be blank and are not used.
3. Check that the IPv4 or IPv6 service IPs of all nodes in the cluster are set.
4. Set any blank IPv4 or IPv6 service IPs. To set a service IP, use the **satask chserviceip** command.
5. Reissue the command.

CMMVC8720E The command failed because the IP address provided is already configured as a DNS server.

Explanation

Each configured DNS server must have a different IP address.

User response

Retry the command with a different IP address.

CMMVC8721E The command failed because the IP address provided is already in use by this system.

Explanation

The IP address that is specified in the command is an IP address on the current system and the system does not provide a DNS service.

User response

Provide the correct IP address of a DNS server.

CMMVC8722E The command has failed because DNS server cannot be removed while DNS clients are configured.

Explanation

The last configured DNS server can be deleted only if no DNS clients are configured.

User response

Stop all DNS clients for the specified server and then retry the command. Because different kinds of DNS clients can trigger this message, specific instructions for stopping the clients cannot be provided here. Refer to the product documentation or contact your service support representative.

CMMVC8723E **The command failed because the system does not support the cloud gateway function.**

Explanation

An attempt was made to enable the cloud gateway function when at least one node in the system does not support the function.

User response

Check the hardware compatibility matrix in the documentation. Remove the nodes that do not support the function from the system and try again.

CMMVC8724E **The command failed because the candidate does not support cloud snapshots and there is a configured cloud account.**

Explanation

An attempt was made to add a node or enclosure that does not support cloud snapshots to a system on which the cloud snapshot function is enabled.

User response

Add a different node, or disable cloud snapshots by using the **rmcloudaccount** command and try again.

CMMVC8725E **The command failed because DNS servers are not configured.**

Explanation

An attempt was made to configure a cloud account, but the system does not have any configured DNS servers.

User response

Configure at least one DNS server by using the **mkdnsserver** command, and then retry the **mkcloudaccountawss3** command.

CMMVC8726E **The command failed because the encryption capability is not currently available on this system.**

Explanation

An attempt was made to configure an encrypted cloud account, but the encryption capability is not currently available on this cluster.

User response

Enable encryption on the system by using the **chencryption** command, then retry the **mkcloudaccount** command.

CMMVC8727E **The command failed because the cloud gateway service is offline.**

Explanation

An attempt was made to test or modify the properties of a cloud account. The system cloud gateway service has reset too often and is being held offline. Cloud accounts cannot be manipulated when they are in this state.

User response

Check the log for an alert event that is related to this error. Mark the event as fixed, and the system will try to bring the account online. If the attempt fails, a new error message is displayed. Refer to the documentation for the displayed error code for further instructions, then resubmit the command.

CMMVC8728E **The command failed because the cloud account is with a different provider.**

Explanation

An attempt was made to modify an existing cloud account, but the wrong command was used. For example, if a cloud account is created that uses AWS S3 as the cloud provider, only the **chcloudaccountawss3** command can be used to modify it.

User response

Use the command that is appropriate to the account type to modify an account.

CMMVC8729E **The command failed because the maximum number of systems are already using the cloud account.**

Explanation

An attempt was made to configure a cloud account with credentials that identify cloud storage that is already being used by the maximum number of systems.

User response

Specify different cloud storage or stop some systems from using the specified cloud storage.

CMMVC8730E **The command failed because of a problem accessing metadata objects in the cloud storage.**

Explanation

An attempt was made to access cloud storage that was written by a different system, or by the current system at a different time. The current system cannot correctly read the metadata that was written to the cloud storage by the originating system.

User response

Verify with your cloud service provider that the cloud storage is working correctly. If you are trying to access cloud storage that was written by a different system, verify that the system is working correctly. Retry the command.

CMMVC8731E **The command failed because the metadata objects were created by a more recent version of code that is incompatible with this system.**

Explanation

An attempt was made to access cloud storage that was written by a different system, or by the current system at a different time. The metadata that was created by the other system is not compatible because the other system is running a more recent level of code than the current system.

User response

Update the code level of the current system to match that of the other system, then retry the command.

CMMVC8732E **The command failed because the new credentials identify different cloud storage.**

Explanation

An attempt was made to change the login details of a cloud account. The new details work, but the cloud storage resources that are accessible by using the new details are different from the storage resources that the system is using.

User response

If you want to use different cloud storage resources, set up a new cloud account object. If you do not want to use different resources, check the details that you are trying to change and try the command again.

CMMVC8733E **The command failed because there is at least one cloud account configured.**

Explanation

Ordinarily, this error is seen only by service support representatives. An attempt was made to modify the cluster ID alias. At least one cloud account is configured that is using the current cluster ID alias to identify data in the cloud. The cluster ID alias cannot be changed in this situation because data exists in cloud storage that refers to it.

User response

If you specified the correct cluster and must change the cluster ID alias, you must first delete the associated cloud account by using the **rmcloudaccount** command. You can then retry the **chsystem** command.

CMMVC8734E **The command failed because the system cannot resolve the hostname of the cloud provider.**

Explanation

The system cannot resolve the host name that is associated with the cloud provider by using the configured DNS servers for the system. For some cloud providers, the user enters the host name. For others, such as Amazon S3 or Microsoft Azure, the system determines the host name.

User response

If an endpoint URL that contains a host name is associated with the cloud account, verify that the host name is correct. Verify that the system DNS servers are configured correctly and are working correctly. After you correct any problems that you find, retry the command.

CMMVC8735E **The command failed because the system cannot contact the cloud provider server on the management network.**

Explanation

The system cannot contact the cloud provider over the management Ethernet network.

User response

If an endpoint URL that contains an IP address is associated with the cloud account, verify that the IP address is correct. Determine why the system

management IP address cannot contact the cloud provider; a firewall might be in the way, or the cloud account provider might be experiencing problems. After you correct any problems that you find, retry the command.

CMMVC8736E The command failed because the system cannot establish a connection with the cloud provider software.

Explanation

The system cannot communicate with the cloud provider server software. The cloud provider address on the system might be configured incorrectly or the cloud provider might be having a problem.

User response

If an endpoint URL is associated with the cloud account, verify that the URL is correct. Check that the cloud provider is functioning correctly and is compatible with the system. After you correct any problems that you find, retry the command.

CMMVC8737E The command failed because the system does not have a CA SSL certificate for the cloud provider server.

Explanation

An attempt was made to configure a cloud account that uses SSL. The cloud provider server presented a certificate, but the system cannot verify its authenticity because the system does not have a corresponding certificate authority (CA) certificate.

User response

Contact the cloud provider and obtain the appropriate CA certificate. Or, if SSL is not a requirement, see whether the cloud provider will support a connection where SSL is disabled.

CMMVC8738E The command failed because the SSL certificate has expired.

Explanation

An attempt was made to configure a cloud account that uses SSL, and an SSL certificate was provided, but the current system time is not within the validity period for the certificate.

User response

Verify that the system time is correct. If so, contact the cloud provider to obtain a valid SSL certificate.

CMMVC8739E The command failed because the SSL certificate is not valid.

Explanation

An attempt was made to configure a cloud account that uses SSL. The SSL certificate that was supplied does not work because it is not valid or because it uses features that the system does not support.

User response

Verify that the certificate that you supplied is a valid X509 SSL CA certificate in PEM format. If not, replace the certificate. Verify that the system security setting is compatible with the cloud provider. After you correct any problems that you find, retry the command.

CMMVC8740E The command failed because the account credentials were rejected by the cloud provider.

Explanation

The cloud provider did not accept the login credentials that were configured on the local system.

User response

Verify that the time is set correctly on the system. Verify that the credentials that were entered on the system match valid credentials on the cloud provider. After you correct any problems that you find, retry the command.

CMMVC8741E The command failed because the account credentials do not give permission to access the cloud storage containers.

Explanation

The cloud provider accepted the login credentials that were configured on the local system, but they do not give the system enough permissions to use cloud storage.

User response

Verify that the container prefix that is associated with the cloud account is available and not in use by other cloud clients. Either change the credentials so that they specify a user that has more permissions, or access the cloud provider and give the current user

the necessary permissions, as described in the system documentation.

CMMVC8742E The command failed because of an error communicating with the cloud provider.

Explanation

The system encountered an unexpected error when it attempted to communicate with the cloud provider.

User response

Verify that the cloud provider is working normally. Check the system event log for alerts. If you cannot determine the source of the problem, contact your service support representative.

CMMVC8743E The command failed because the region specified is not valid.

Explanation

An attempt was made to configure an Amazon Web Services (AWS) S3 cloud account, but the specified AWS region was not recognized.

User response

Retry the command with a valid region. To find a valid region, check the list of regions that are supported by AWS S3 on the AWS website:

https://aws.amazon.com/about-aws/global-infrastructure/regional-product-services/

CMMVC8744E The command failed because the bucket prefix or container prefix is already in use.

Explanation

An attempt was made to configure a cloud account that uses Amazon S3 or Microsoft Azure, but the bucket prefix or container prefix was already in use.

User response

Retry the command with different bucket or container prefix. For more information about bucket or container naming, see Amazon S3 or Microsoft Azure documentation.

CMMVC8746E The command failed because an account with an SSL certificate must have an https endpoint.

Explanation

If you provide an SSL certificate when you configure a cloud account, you must provide an endpoint URL that requires SSL. If you provide an endpoint that does not use https, you must not provide an SSL certificate.

User response

Retry the command, either with an https endpoint URL or without the SSL certificate.

CMMVC8748E The command failed because the cloud data is encrypted with the wrong key.

Explanation

The master key associated with the cloud data does not match the master key that was used when the system was created. Cloud snapshot services remain unavailable until this problem is fixed.

User response

Locate the correct master key and make it available to the cluster, either on a USB drive or on a network key server. Then, retry the command.

CMMVC8749E The command failed because the specified volume group is not empty.

Explanation

An attempt was made to delete a volume group that contains volumes. Volume groups cannot be deleted while they contain volumes.

User response

Remove the volumes from the group by using the **chvdisk -novolumegroup** command and try again.

CMMVC8750E The command failed because the specified volume group already contains the maximum number of volumes.

Explanation

The volume cannot be added to the specified volume group after its maximum volume member limit is reached.

User response

Try one of the following solutions:

- Use the specified volume as a stand-alone volume.

- Remove an existing volume from the group by using the **chvdisk -novolumegroup** command, then try again to add the new volume.
- Retry the command and specify a different volume group.

CMMVC8751E A new cloud snapshot could not be created because the volume group does not contain any volumes.

Explanation

An attempt was made to back up an empty volume group.

User response

Add at least one volume to the group and try again. Alternatively, specify a different volume group.

CMMVC8752E A cloud snapshot could not be created because the cloud FlashCopy function is not enabled for all the volume group members.

Explanation

An attempt was made to back up a volume group where at least one volume in the group does not have the cloud snapshot function enabled. A new cloud snapshot can be created only when the cloud backup function is enabled for all the volume group members.

User response

Enable the cloud snapshot function for all the volumes in the group and try again. Use the following command to enable cloud snapshots for a volume:

```
chvdisk -backup cloud -enable -account
account volume
```

CMMVC8753E A new cloud snapshot could not be created because a cloud snapshot operation is already in progress.

Explanation

A new cloud snapshot can be started only when the volume **backup_status** is **ready**.

User response

Wait for the existing cloud snapshot operation to complete, or cancel the cloud snapshot. Use the management GUI or the **lsvolumebackupprogress** command to monitor the progress of the cloud snapshot, or the **rmvolumebackupgeneration** command to cancel the cloud snapshot.

CMMVC8754E A cloud snapshot could not be created because a restore operation is already in progress.

Explanation

A new cloud snapshot cannot be started while the volume is being restored.

User response

Make sure that you specified the correct volume. If so, wait for the existing restore operation to complete, or cancel the restore. Use the management GUI or the **lsvolumerestoreprogress** command to monitor the progress of the restore, or the **restorevolume** command to cancel the restore.

CMMVC8755E The command failed because the system does not support the volume group function.

Explanation

An attempt was made to enable the volume group function, but at least one node in the system does not support the function.

User response

Check the hardware compatibility table in the documentation. Remove the nodes that do not support the function from the system and try again.

CMMVC8756E A new cloud snapshot could not be created because the volume is part of a volume group.

Explanation

A new cloud snapshot can be created only if the volume is not part of a volume group.

User response

Remove the volume from the group or use the **backupvolumegroup** command to create a cloud snapshot of the entire group.

CMMVC8757E Cannot change facility because CADF reporting is set.

Explanation

An attempt was made to change the facility code for the syslog server. The facility code is automatically set to 8 when Cloud Auditing Data Federation (CADF) reporting is enabled.

User response

To change the facility code, you must first disable CADF reporting by using the **chsyslogserver -cadf off** command.

CMMVC8758E **Cloud snapshots cannot be enabled because there are insufficient free volume copy IDs in the system.**

Explanation

The cloud snapshot function requires IDs to create two internal volumes in the same caching I/O group as the volume.

User response

Remove unwanted volumes or volume copies from the system by using the **rmvolume** or **rmvolumecopy** command, and then retry the command.

CMMVC8759E **Cloud snapshots cannot be enabled because there are insufficient free volume copy IDs in the caching I/O group of the volume.**

Explanation

The cloud snapshot function requires IDs to create two internal volumes in the same caching I/O group as the volume.

User response

Complete one of the following actions:

- Remove unwanted volumes from the caching I/O group.
- Move the volume to a different I/O group.

Then, retry the command.

CMMVC8760E **Cloud snapshots cannot be enabled because there is insufficient free capacity in the storage pool.**

Explanation

The cloud snapshot function requires added storage capacity in the storage pool of the volume.

User response

Complete one of the following actions:

- Add capacity to the storage pool by using the **addmdisk** command, or use the **mkarray** command if using internal drives.
- Remove unwanted volumes from the storage pool by using the **rmvdisk** command.
- Move the volume to a different storage pool by using the **migratevdisk** command.

CMMVC8761E **Cloud snapshots cannot be enabled because the maximum number of FlashCopy mappings in the system has been reached.**

Explanation

The cloud snapshot function configures two internal FlashCopy mappings per volume.

User response

Remove unwanted FlashCopy mappings from the system and retry the command.

CMMVC8762E **Cloud snapshots cannot be enabled because there is insufficient memory available for FlashCopy.**

Explanation

The restore operation requires an internal FlashCopy mapping to be created, but the system limit was reached.

User response

Remove unwanted FlashCopy mappings from the system by using the **rmfcmap** command and retry the command.

CMMVC8763E **Cloud snapshots cannot be enabled because this volume is part of a FlashCopy mapping.**

Explanation

The cloud snapshot function cannot be used with volumes that are the source or target of a FlashCopy mapping.

User response

Make sure that you specified the correct volume. If so, remove all FlashCopy mappings that include this volume and retry the command.

CMMVC8764E **Cloud snapshots cannot be enabled because this volume is**

part of a Metro Mirror or Global Mirror relationship.

Explanation

The cloud snapshot function cannot be used with volumes that are in a remote copy relationship.

User response

Make sure that you specified the correct volume. If so, remove the remote copy relationship that includes this volume and retry the command.

CMMVC8765E **Cloud snapshots cannot be enabled because this is a HyperSwap volume.**

Explanation

The cloud snapshot function cannot be used with HyperSwap volumes.

User response

Make sure that you specified the correct volume. If so, convert the volume to a basic volume by removing a copy at one site, then retry the command.

CMMVC8766E **Cloud snapshots cannot be enabled because the volume is mirrored between two different pools.**

Explanation

The cloud snapshot function cannot be used with volumes that have copies in different pools.

User response

Make sure that you specified the correct volume. If so, convert the volume to a basic volume by removing a copy, then retry the command.

CMMVC8767E **Cloud snapshots cannot be enabled because the volume is being migrated between pools.**

Explanation

An attempt was made to enable the cloud snapshot function while the volume was being migrated between pools.

User response

Wait for the migration operation to finish before you enable cloud snapshots for this volume. Use the management GUI or the **lsmigrate** command to monitor the progress of the migration.

CMMVC8768E **Cloud snapshots cannot be enabled because the volume capacity is zero bytes.**

Explanation

The cloud snapshot function cannot be used with volumes that have a capacity of zero bytes.

User response

Make sure that you specified the correct volume. If so, expand the size of the volume and retry the command.

CMMVC8769E **Cloud snapshots cannot be disabled because there is a cloud snapshot of the volume in progress.**

Explanation

An attempt was made to disable the cloud snapshot function while a cloud snapshot was in progress.

User response

Wait for the cloud snapshot operation to complete or cancel the cloud snapshot. Use the management GUI or the **lsvolumebackupprogress** command to monitor the progress of the cloud snapshot, or the **rmvolumebackupgeneration** command to cancel the cloud snapshot. When the snapshot is no longer in progress, retry the command.

CMMVC8771E **The specified parameter lists must be of equal length.**

Explanation

The **driveclass** parameter must contain the same number of colon-separated list elements as the **drivecount** parameter.

User response

Validate the command line and retry the command.

CMMVC8772E **The command cannot be completed because the I/O group of the drive class has the maximum number of arrays configured.**

Explanation

The **mkdistributedarray** command fails because the maximum number of arrays already exists, as configured in the I/O group of the drive class.

Note: Each distributed array occupies 16 slots, starting at an MDisk ID divisible by 16, in the MDisk table. For more information, refer to the documentation for the **lsmdisk** command.

User response

Remove existing arrays in the I/O group before trying the command again.

CMMVC8773E Unable to create the array due to insufficient available drives.

Explanation

The command fails because no drives of the specified drive class in the appropriate drive state are available to create the full array.

Some drives might be available, but not all of the required drives.

Perhaps some drives are not in the candidate state.

Perhaps an enclosure might be offline.

User response

After correcting the drive availability problem for the proper class, use the **lsdriveclass** command or the lsdriveclass view in the GUI to verify the number of available drives. Then try the command again.

CMMVC8774E An array cannot be created with the drive count, stripe width and rebuild areas entered.

Explanation

The **mkdistributedarray** command fails because the values for the drive count, stripe width, and rebuild areas do not work together.

The drive count value must equal or be greater than the combined value of the stripe width and rebuild areas count.

User response

After determining what the proper values should be, try the command again to create the array.

CMMVC8775E Command cannot be initiated because the distributed array does not have an available rebuild space. Replace failed member drives or swap members occupying rebuild spaces before retrying.

Explanation

The **charraymember** or the **chdrive** command fails because not enough rebuild space is available.

A distributed array has a limited number of spaces where it rebuilds. The command requires the use of a rebuild space, but there are no free spaces to use.

Replacing failed drives creates space in the array to begin a copyback that copies data back from rebuild space. Replacing failed drives thus frees rebuild space for the command to use.

User response

Verify that the array member ID is correct. Replace failed hardware to free rebuild space. If failing a drive, allow the command to degrade redundancy. Then try the command again.

CMMVC8776E A parameter is not supported for this array type.

Explanation

The **charraymember** or the **charray** command fails because you used a parameter that is not supported for the array type (distributed or traditional) to attempt a change that is not supported.

Changes to the `rebuildareasgoal` and `initnewextents` parameters are supported for distributed arrays only.

Changes to the `sparegoal` and `balanced` parameters are supported on traditional arrays only.

User response

Try the command again using parameters that are appropriate for the type of array.

CMMVC8777E The command failed because the maximum number of encryption keys has been reached.

Explanation:
The resources that are needed to create a new key have all been used.

User response:
If your object has an individually selectable encryption attribute, consider selecting the **-encrypt no** option for that object. Otherwise, review all encrypted objects in the system. If possible, free up some existing encrypted objects and then retry the command.

CMMVC8778E **The command failed because the object is the last active copy of its volume.**

Explanation

The volume must contain at least one copy.

User response

If you must delete the copy, delete the entire volume by using the **rmvolume** command.

CMMVC8779E **The provided information was not sufficient to identify a specific volume copy.**

Explanation

The command uses one or more of the following criteria to delete a volume copy:

- Site number
- Storage pool
- Copy ID

The command did not specify enough information to uniquely identify the volume copy to be deleted.

User response

Retry the command, providing more information.

CMMVC8780E **The volume copy was not deleted because it is part of a FlashCopy mapping.**

Explanation

You cannot remove a volume copy that is part of a FlashCopy mapping.

User response

First, verify that you specified the correct volume copy. If not, retry the command with a volume copy that is not part of a FlashCopy mapping.

If you specified the correct volume copy, complete one of the following tasks:

- Remove the FlashCopy mapping and then retry the command.
- Specify the **-removefcmaps** parameter to force-delete the volume copy. This parameter stops any dependent FlashCopy maps and must be used with caution.

CMMVC8782E **The volume copy was not deleted because the copy has an image mode mdisk and a change volume is providing the consistent image.**

Explanation

A volume copy cannot be deleted when a change volume is in use.

User response

First, verify that you specified the correct volume copy. If not, retry the command with a volume copy that does not have an image mode MDisk.

If you specified the correct volume copy, complete one of the following tasks:

- Wait for the use of the change volume to complete and then retry the command.
- Specify the **-discardimage** parameter to force-delete the volume copy. A forced deletion can result in data loss on the image mode copy.

CMMVC8783E **The volume copy was not deleted because the volume is part of a consistency group.**

Explanation

A volume copy cannot be deleted when it is part of an active-active relationship.

User response

First, verify that you specified the correct volume copy. If not, retry the command with a volume copy that is not part of a consistency group.

If you specified the correct volume copy, remove the active-active relationship for the volume from the consistency group, then retry the command.

CMMVC8784E **The command failed because an additional feature activation is required to use remote copy**

Explanation:
An additional license is required to use remote copy before the action can be completed.

User response

Activate the feature to use remote copy and then retry the command.

CMMVC8785E **The command failed because the node does not support encryption, and encryption is required in the IO group.**

Explanation:

The existing I/O node is capable of encryption. This error occurs when adding an I/O group partner node that does not support encryption.

User response:
Acquire an encryption capable node to add to the I/O group.

CMMVC8786E **The node cannot be added because it does not support encryption, and encryption is in use for some SAN MDisks.**

Explanation:
Encryption is enabled. An attempt was made to add a node that is not encryption-capable while storage pools were being encrypted that include SAN MDisks that are not self-encrypting.

User response:
Acquire an encryption capable node, or remove all non-self-encrypting SAN MDisks from all encrypting storage pools.

CMMVC8787E **The volume was not deleted because it is mapped to a host.**

Explanation

You cannot delete a volume that is mapped to a host.

User response

If you specified the wrong volume, retry the command. If you are sure that you want to remove the specified volume, remove the host mapping and retry the command. Alternatively, you can include the **-removehostmappings** parameter with the **rmvolume** command to force the removal of the host mappings.

CAUTION: When you use a "force" parameter such as **-removehostmappings**, you risk loss of data.

CMMVC8788E **The volume was not deleted because it is part of Metro Mirror or Global Mirror relationship.**

Explanation

You cannot delete a volume that is part of a remote copy relationship.

User response

If you specified the wrong volume, retry the command. If you are sure that you want to remove the specified volume, remove the remote copy relationship and retry the command. Alternatively, you can include the **-removercrelationships** parameter with the **rmvolume** command to force the removal of the remote copy relationship.

CAUTION: When you use a "force" parameter such as **-removercrelationships**, you risk loss of data.

CMMVC8789E **The volume was not deleted because it is part of a FlashCopy mapping.**

Explanation

You cannot delete a volume that is part of a FlashCopy mapping.

User response

If you specified the wrong volume, retry the command. If you are sure that you want to remove the specified volume, remove the FlashCopy mapping and retry the command. Alternatively, you can include the **-removefcmaps** parameter with the **rmvolume** command to force the removal of the FlashCopy mappings.

CAUTION: When you use a "force" parameter such as **-removefcmaps**, you risk loss of data.

CMMVC8790E **The volume was not deleted because it would result in inconsistent data on an image mode mdisk.**

Explanation

You cannot delete an image mode volume that has inconsistent data.

User response

If you specified the wrong volume, retry the command. If you are sure that you want to remove the specified volume, wait until the data on the image mode volume is consistent before you retry the command. Alternatively, you can include the **-discardimage** parameter with the **rmvolume** command to force the removal of the volume.

CAUTION: When you use a "force" parameter such as **-discardimage**, you risk loss of data.

CMMVC8791E **There is already an outstanding certificate request. Use the -cancel flag to discard the request and then generate a new certificate request.**

Explanation:

You issued **chsystemcert -mkrequest** twice without successfully installing the certificate from the first request that uses **chsystemcert -install**.

User response

Do one of the following:

- Get the outstanding request signed and install it by using the **chsystemcert -install** command.
- Abandon the outstanding request and start a new one by using the 'chsystemcert -cancel' command.

Note: If a new request is started, any certificate that is created from a previous request can no longer be used.

CMMVC8792E **Cannot install certificate because there is no outstanding certificate request.**

Explanation:
You are attempting to install a certificate file using **chsystemcert -install** without first creating a certificate request. Only a signed certificate that is produced from a certificate request that is created by the system can be installed.

User response

Only a signed certificate request can be installed. Use the **chsystemcert -mkrequest** command to create a certificate request and then get this request signed by a certificate authority (CA). The resulting certificate can then be installed.

CMMVC8793E **Cannot install certificate because it has the wrong key.**

Explanation

You are attempting to install a certificate file by using **chsystemcert -install**, which is not derived from the outstanding certificate request. Either:

- You tried to install the wrong certificate
- 'mkrequest -force' was used between the certificate request being generated and the installation of the signed certificate.

  Note: Get the latest request signed and install this certificate.

User response

Find the correct certificate to install or use the **chsystemcert -mkrequest -force** command to start a new certificate request.

CMMVC8794E **Invalid certificate file.**

Explanation:
The certificate file that you are attempting to install cannot be parsed. The certificate must be provided in base64 encoded PEM format.

User response

Check that the certificate to be installed is in the correct format and then copy the file to the system. Attempt to install it again.

CMMVC8795E **This command is no longer supported. Use the CLI command 'chsystemcert' to manage the SSL certificate for the system.**

Explanation:
The CLI command 'chsystem -regensslcert' is no longer supported. A new CLI command 'chsystemcert' that provides more options for managing the SSL certificate for the system, replaced it.

User response

Use the 'chsystemcert' command instead.

CMMVC8796E **Cannot make unencrypted MDisk Group as parent MDisk Group [%1] has an encryption key.**

Explanation:
You are not allowed to make an unencrypted child MDisk Group when the parent MDisk Group has an encryption key.

User response:
The use case is to make an unencrypted child MDisk Group on an encrypted system. You will need to make a new MDisk Group without an encryption key instead.

CMMVC8797E **Unable to change mdisk property while being part of an MDisk Group with an encryption key.**

Explanation:
The MDisk Group and/or one of its child pools has an encryption key that it is using for the MDisk. This can not be changed while the MDisk contains customer data.

User response:
You need to remove the MDisk from the MDisk Group, apply the change in property, and then add back to the MDisk Group.

CMMVC8798E **Cannot add unencrypted MDisk to encrypted MDisk Group**

Explanation:
The user is not allowed to add an unencrypted MDisk to an encrypted storage pool.

User response:

Only MDisks in encrypting I/O groups can be added to an encrypted storage pool.

CMMVC8799E It is not possible to create an encrypted image disk.

Explanation:
You are trying to create an image mode disk that would be encrypted.

User response:
Image mode disks can be created in an unencrypted storage pool only.

CMMVC8800E Cannot migrate as pool *storage_pool* is encrypted.

Explanation:
Encrypted data cannot be moved unless the source and destination have the same encryption key.

User response:
Create or add a volume copy and then (auto) delete the source once copied.

CMMVC8801E Insufficient drive count specified.

Explanation

The **lspotentialarraysize** command fails when its `drivecount` parameter is less than its `stripewidth` parameter.

The number of drives in the array must be an integer from 3 - 128. The stripe width varies per RAID type:

RAID
Stripe width

R1
2-16

R5
3-16

R6
5-16

R10
Even integers from 2-16

User response

Retry the command with a drive count that is equal to or greater than the stripe width.

CMMVC8802E The command failed because the parent pool contains an unencrypted external MDisk.

Explanation:
An attempt was made to create an encrypted child pool where the parent pool has an unencrypted external MDisk and at least one I/O group contains hardware that is not reported as encryption capable. It is not possible for non-encryption-capable nodes to perform encrypted operations.

User response:
Either remove the non-encryption-capable I/O group members from the cluster or create the pool as not encrypted.

CMMVC8803E Stripe width incorrect for RAID type.

Explanation

The **lspotentailarraysize** command fails because the stripe width is not good. Valid values are:

RAID
Stripe width

R0
1-8

R1
2

R5
3-16

R6
5-16

R10
Even integers from 2-16

User response

Retry the command with a valid stripe. Use the **lsarrayrecommendation** command to recommend an optimal value.

CMMVC8804E The number of rebuild areas specified must be greater than or equal to the rebuild area goal parameter.

Explanation

The command fails because it creates a distributed array with a goal of more rebuild areas than are available in this array.

User response

If you use the `rebuildareas` parameter of the **mkdistributedarray** command, ensure that its value is greater than or equal to the `rebuildareasgoal` parameter value.

If you do not use the `rebuildareas` parameter, but do use the `rebuildareasgoal` parameter, the goal value is greater than the default areas value of 1 for RAID5 or 2 for RAID6. Either reduce the goal or specify a greater areas value than the default.

CMMVC8805E **An array of this RAID level cannot be created using this drive capacity.**

Explanation

The command fails because it cannot create a distributed array with drives that have a capacity exceeding the limit for the RAID level. For example, you cannot create a RAID5 array if the drive capacity is greater than or equal to 8 TB.

User response

Change the RAID level or use drives with a capacity less than the limit for the RAID level, such as 8 TB for RAID5.

CMMVC8806E **The specified drive count exceeds available drives in the drive class.**

Explanation

The **lsarrayrecommendation** command fails because the colon-delimited list of drive counts has a drive count that is greater than the number of available drives in the respective drive class.

User response

Try the command again with each drive count equal to or less than the number of available drives in its respective drive class.

CMMVC8807E **A duplicate drive class ID was entered in the list.**

Explanation

The **lsarrayrecommendation** command fails when the colon-delimited list of drive classes has duplicate IDs.

User response

Try the command again after editing the list of drive classes.

CMMVC8808E **Cannot create the encrypted child pool because the parent pool contains an array that is provided by a non-encryption capable IO group.**

Explanation:
You attempted to create an encrypted child pool where the parent pool contains an array that is provided by an enclosure that does not support encryption.

User response:
Either remove the array from the parent group, create the child pool with a parent pool that doesn't contain such arrays or create the child pool without encryption.

CMMVC8809E **Cannot create the encrypted child pool because the parent pool contains an unencrypted DRAID array.**

Explanation:
You attempted to create an encrypted child pool where the parent pool contains an unencrypted distributed array.

User response:
Remove the array from the parent group or create the child pool without encryption.

CMMVC8810E **The volume was not deleted because it is configured as a change volume for a Remote Copy relationship. Either disassociate the volume from the relationship or use -removefcmaps to force delete the volume, which might result in the loss of data.**

Explanation

An attempt was made to delete a volume that is configured as the master or auxiliary change volume in a remote copy relationship.

User response

Complete one of the following actions:

- If you entered the wrong volume, correct the command and retry it.
- Remove the volume from the remote copy relationship and retry the command.
- Retry the command with the **-removefcmaps** parameter to force the deletion of the volume.



CAUTION: A forced deletion can result in data loss.

CMMVC8811E **The command is not supported for this array type.**

Explanation

Any of the following actions can cause this error message to be displayed:

- Trying to change an array to be encrypted by using the **charray -encrypt** command.
- Trying to change the balanced or spare goal parameters for a distributed array.

- Trying to change the rebuild areas goal or init new extents parameters for a nondistributed array.
- Trying to change the balance parameter for a distributed array member.
- Trying to change a distributed array to use latency mode.
- Trying to create a distributed array in latency mode.
- Running the **lspotentialarraysize** command for RAID0
- Running the **lspotentialarraysize** command for RAID1 if rebuild areas set or for RAID10 if rebuild areas set.

User response

Specify an array of the correct type and retry the command.

CMMVC8812E **The command failed because the existing volume may have only one mirrored copy.**

Explanation

In a HyperSwap topology, the existing volume can have only a single mirrored copy.

User response

Specify a different volume or remove an existing mirrored copy, and then retry the command.

CMMVC8813E **The command failed because the existing volume has mirrored copies in different sites.**

Explanation

In a HyperSwap topology, all of the mirrored copies for the existing volume must have a matching site of 1 or 2 (not 3 or none). This configuration ensures that the existing volume itself has a consistent and well-defined site of 1 or 2.

User response

Ensure that all mirrored copies of the existing volume have a matching site of 1 or 2 and then retry the command.

CMMVC8814E **The command failed because the existing volume has mirrored copies in storage pools whose site is not 1 or 2.**

Explanation

In a HyperSwap topology, all of the mirrored copies for the existing volume must have a matching site of 1 or 2 (not 3 or none). This configuration ensures that the existing volume itself has a consistent and well-defined site of 1 or 2.

User response

Ensure that all mirrored copies of the existing volume have a matching site of 1 or 2 and then retry the command.

CMMVC8815E **The command failed because the IO group is not in the same site as the storage pool.**

Explanation

The **-iogrp** parameter applies only when you create a HyperSwap volume copy and requires that the system be configured with a HyperSwap topology. The caching I/O group must be in the same site as the storage pool that you specify for the volume copy that you are creating.

User response

Correct the command and retry.

CMMVC8816E **The command failed because the specified FlashCopy mapping is controlled by FlashCopy.**

Explanation

This command is not valid for an internally controlled FlashCopy map.

User response

Specify a FlashCopy map that is not internally owned and retry the command.

CMMVC8818E **The command failed because the volume already has a copy in the specified site. Use the addvdiskcopy command to add a second copy in the same site.**

Explanation

An attempt was made to add a volume copy to a site by using the **addvolumecopy** command when the volume already has a copy in the specified site.

User response

Use the **addvdiskcopy** command to create an additional copy in the same site.

CMMVC8819E **The volume copy could not be created because the existing source copy is not compressed.**

Explanation:
When creating new volume copies, use the `ignoresyncerrors` parameter only when the existing source copy is a compressed copy.

User response

Use the `ignoresyncerrors` parameter for compressed volume copies.

Note: The `ignoresyncerrors` parameter can also help create valid copies of compressed volume copies that were corrupted due to a bug in the compression engine.

CMMVC8820E **The command failed because the specified volume group did not meet all the requirements to be associated with a Safeguarded policy.**

Explanation

For a volume group to be associated with a Safeguarded policy, all the volumes must meet the criteria to become Safeguarded sources. The error occurred when one or more volumes in the volume group do not meet the criteria.

User response

Review the volumes to determine why do the volumes not meet the criteria to be a Safeguarded source. Rectify the problem or modify the set of volumes in the volume group.

CMMVC8821E **The command failed because a volume group cannot be associated with a Safeguarded policy and have an ownership group.**

Explanation:
The **-safeguardedpolicy** parameter is mutually exclusive with the **-ownershipgroup** parameter.

User response

Remove anyone of the two parameters and resubmit the command. If the error still persists, then specify a different volume group.

CMMVC8822E **The command failed because the -retentiondays parameter must be specified when a target volume is a Safeguarded backup.**

Explanation

The **-retentiondays** parameter is only allowed if a FlashCopy consistency group contains a target volume that is a Safeguarded backup. The parameter must not be used if the target volume is not a Safeguarded backup.

User response

Specify the number of days to retain the Safeguarded backup if the target volume is a Safeguarded backup in the FlashCopy mapping.

CMMVC8823E **The command failed because all the FlashCopy mappings in a consistency group must have a Safeguarded backup as their target volume to copy the data.**

Explanation

A FlashCopy mapping defines the relationship that copies data between a source volume and a target volume. The error occurs when all the FlashCopy mappings in a consistency group do not have a safeguarded backup as the target volume.

User response

Add the FlashCopy mappings to a different consistency group that has a Safeguarded backup and resubmit the command.

CMMVC8824E **The command failed because the user's role is not authorized to submit the command.**

Explanation:
This error occurs when an attempt is made by a user who is not authorized to run the command or to run the command with certain parameters, based on the user's role. For example, only the administrator and the superuser role users have permission to shrink the size of a Safeguarded source volume.

User response

Select the one of the following steps to resolve the issue:

1. Log in as a user that has a role that is authorized to run the command. Then, resubmit command.

2. Contact someone with the administrator or superuser role and have them run the command.
3. Select a different source volume.

CMMVC8825E **The command failed because the -backup parameter cannot be specified when the specified volume is a Safeguarded source.**

Explanation:
The error occurs when **-backup** parameter value specifies when the volume specified is a Safeguarded source.

User response

Submit the command without specifying **-backup** parameter or select a different volume.

CMMVC8826E **The command failed because the selected volume in the volume group does not meet all the requirement to be a Safeguarded source.**

Explanation:
The **-volumegroup** parameter that specifies a new volume group to a volume, fails if the volume group is associated with a Safeguarded policy and the volume does not meet the criteria to be a Safeguarded source.

User response

Check the volume to meet all the requirements to be a Safeguarded source or assign the volume to a different volume group.

CMMVC8827E **The command failed because with the current user role, the requested action cannot be completed on Safeguarded backups that depend on FlashCopy mapping.**

Explanation:
This error occurs when an attempt is made to run a command by a user who is not authorized to run the command or to run the command with certain parameters, based on the user's role. For example, when the Safeguarded backups depend on a FlashCopy mapping, only the Storage Virtualize administrator and the superuser role users have permission to use **-force** parameter on **stopfcmap** command.

User response

Select the one of the following steps to resolve the issue:

1. Log in as a user that has a role that is authorized to run the command. Then, resubmit command.
2. Contact someone with the administrator or superuser role and have them run the command.
3. Select a different FlashCopy mapping.

CMMVC8828E **The command failed because an ownership group cannot be assigned to Safeguarded backup location.**

Explanation:
For the Safeguarded Copy function, a source volume cannot be in an ownership group. If the specified pool is a Safeguarded backup location, **-ownershipgroup** parameter cannot be used.

User response

Select the one of the following steps to resolve the issue:

1. Select another storage pool that is not a Safeguarded backup location.
2. Run the command without specifying **-ownershipgroup** parameter.

CMMVC8829E **The command failed because the -grainsize parameter cannot be specified when the target volume is a Safeguarded backup.**

Explanation:
When the target volume is a Safeguarded backup, the **-grainsize** parameter cannot be specified.

User response

Resubmit **mkfcmap** command without specifying the **-grainsize** parameter.

CMMVC8830E **The command failed because with the current user role the requested action cannot be completed on a Safeguarded FlashCopy mapping.**

Explanation:
When a FlashCopy mapping is a Safeguarded FlashCopy mapping, only the Security Administrator role and the superuser role users have permission to use **stopfcmap** command.

User response

Select the one of the following steps to resolve the issue:

1. Log in as a user that has a role that is authorized to run the command. Then, resubmit command.

2. Contact someone with the administrator or superuser role and have them run the command.
3. Select a different FlashCopy mapping.

CMMVC8831E **The command failed because the current user role is not authorized to submit the command when FlashCopy consistency group contains Safeguarded FlashCopy mappings.**

Explanation:
When the FlashCopy consistency group contains Safeguarded FlashCopy mappings, only the Security administrator role and the superuser role users have permission to use **stopfcconsistgrp** command.

User response

Select the one of the following steps to resolve the issue:

1. Log in as a user that has a role that is authorized to run the command. Then, resubmit command.
2. Contact someone with the administrator or superuser role and have them run the command.
3. Select different FlashCopy consistency group.

CMMVC8832E **The command failed because the safeguarded backup start time format is invalid.**

Explanation:
The invalid time format fails the validation tests. The valid time format is YYMMDDHHMM.

User response

Enter the valid time format and resubmit the command.

Note: Do not enter the seconds value.

CMMVC8833E **The command failed because the predefined Safeguarded policies cannot be modified or removed.**

Explanation:
The management GUI supports several predefined Safeguarded policies to simplify policy management. The management GUI does not support creating customized Safeguarded policies. The command modifies the existing safeguarded policy. However, it cannot be used on a predefined Safeguarded policy.

User response:
Specify the predefined Safeguarded policies.

CMMVC8834E **The -copyrate and the -autodelete cannot be modified if the target volume is a Safeguarded FlashCopy mapping.**

Explanation

When the mapping is in the copying state, the **-copyrate** determines the priority that is given to the background copy process. When **-autodelete** is set to on, the mapping is automatically deleted when the mapping reaches the idle or copied state and the progress is 100%. You cannot modify these parameters when the target volume is a Safeguarded backup.

User response:
Select a different FlashCopy mapping or do not specify the **-copyrate** and the **-autodelete**.

CMMVC8835E **The command failed because with the current user role, the requested action cannot be completed.**

Explanation:
This error occurs when an attempt is made to run a command with the **-size** parameter to shrink the size of a Safeguarded backup location based on the user's role.

User response

Select the one of the following steps to resolve the issue:

1. Log in as a user that has a role that is authorized to run the command with the parameter. Then, resubmit the command.
2. Contact someone with the Security Administrator or superuser role and have them run the command.
3. Run the command without specifying **-size** parameter.

CMMVC8836E **The command failed because the Safeguarded Copy operations are suspended.**

Explanation

When the Safeguarded Copy operations are suspended, new volumes in the Safeguarded backup location cannot be created.

User response

Resume the Safeguarded Copy operations by using the **chsystem** command, and resubmit the command.

CMMVC8837E **The command failed because the Safeguarded backup location that is associated with the Safeguarded source volume cannot be removed.**

Explanation

The **rmmdiskgrp** command deletes the specified storage pool. The **-force** parameter is required to delete the volumes that have been created from this storage pool. If Safeguarded sources use the Safeguarded backup locations for backup, then the Safeguarded backup location cannot be deleted.

User response

Login as a Security Administrator to reconfigure your Safeguarded backup location before removing the **mdiskgrp**.

CMMVC8838E The command failed because the system does not support Safeguarded Copy.

Explanation:
This error occurs when the command runs on the system that does not support Safeguarded Copy operations.

User response:
Try to run the command on a system that supports Safeguarded copy policy.

CMMVC8839E A target volume in a Safeguarded FlashCopy mapping cannot be mirrored.

Explanation

A Safeguarded backup must not be mirrored when a Safeguarded FlashCopy mapping is created between a Safeguarded source as the source volume and a Safeguarded backup as the target volume of the mapping.

User response

Either remove one of the copies of the target volume or create a new target volume in the Safeguarded backup location, and resubmit the **mkfcmap** command.

CMMVC8840E The command failed because the policy start time format was invalid. The valid time format is YYMMDDHHMM.

Explanation:
The time format specified is invalid. The valid time format is YYMMDDHHMM.

User response

Enter the valid time format and resubmit the command.

Note: Do not enter the seconds value.

CMMVC8841E The command failed because the predefined snapshot policies cannot be modified or removed.

Explanation:
A predefined snapshot policy cannot be modified or removed.

User response:
Specify a user defined snapshot policy or do not attempt the command.

CMMVC8844E The snapshot policy cannot be removed because of one or more associated volume groups.

Explanation:
There are volume groups that have been associated to this snapshot policy. The snapshot policy cannot be removed unless -removefromvolumegroups is specified.

User response:
Remove the volume group associations via volume group commands before reattempting the removal of the snapshot policy. Or if the user has the proper role, use the -removefromvolumegroups option to remove the volume group associations along with the snapshot policy.

CMMVC8846E The command failed because volume groups do not support both Safeguarded and snapshot policies.

Explanation

The user entered -snapshotpolicy and the volume group already has a Safeguarded policy or -safeguardedpolicy and the volume group already has a snapshot policy.

User response

Remove the existing Safeguarded or snapshot policy before adding a policy of the other type.

CMMVC8847E The command failed because [-backupinterval] cannot be smaller than 60 when [-backupunit minute] is specified.

Explanation

The user attempted to a snapshot backup frequency less than 60 minutes.

User response

Enter a minute interval 60 minutes or greater.

CMMVC8848E **The command failed because the system does not support this function.**

Explanation

The command fails because it is not supported by the product.

User response

Review the documentation for the product to select an appropriate command.

CMMVC8849E **The command failed as -policystarttime cannot be set for a Safeguarded policy that is externally controlled.**

Explanation

You cannot set or modify the start time with an external Safeguarded policy.

User response

Do not specify external safeguarded policy.

CMMVC8850E **The command failed because the snapshot policy has no start time and -policystarttime is then required.**

Explanation

A start time was not specified when snapshot policy was created or assigned to a volume group.

User response

Enter a valid start time i.e. YYMMDDHHMM time format when assigning snapshot policy to volume group or specify a snapshot policy that has its own start time already defined.

CMMVC8851E **The command failed because the snapshot is Safeguarded and your requested action cannot be completed under your current user role.**

Explanation:
Only Security Administrator users or superuser can remove Safeguarded snapshots.

User response:
Use a Security Administrator user or superuser, and retry the command.

CMMVC8852E **The command failed because snapshot policies are suspended at system level.**

Explanation

Cannot resume a volume group's snapshot policy if snapshot policies are suspended at a system level.

User response

Resume snapshot policies at system level with chsystem -snapshotpolicysuspended no, then retry the command.

CMMVC8853E **The command failed because the snapshot policy is Safeguarded and your requested action cannot be completed under your current user role.**

Explanation

This error occurs when the user is not authorized to remove the assignment of a Safeguarded snapshot policy, only users with SecurityAdmin role or superuser are authorized to remove the assignment of a Safeguarded snapshot policy.

User response

Retry the command with superuser or user with SecurityAdmin role.

CMMVC8854E **The current user cannot create a two person integrity request.**

Explanation:
This error might occur if you currently have an elevated role request that is pending approval, is already approved, or you are a superuser. When two person integrity is enabled, you can't request a role elevation when any of these conditions are true.

User response:
Cancel the approved request and then request an elevated role again. You can also wait until the duration of the approved request expires and then request an elevated role again. If you are a superuser, you must log in with an account that is not a superuser to request an elevated role.

CMMVC8855E **The command failed because there are Safeguarded snapshots that depend on objects affected by this command and the requested action cannot be completed under your current user role.**

Explanation:

Only Security Administrator users or superuser can use **-force** if Safeguarded snapshots are present.

User response:
Use a Security Administrator user or superuser, and retry the command.

CMMVC8856E **The command failed because the specified volume group is associated with a snapshot policy and volume does not meet criteria to create snapshot.**

Explanation:
The specified volume group is associated with a snapshot policy and might contain volumes that have properties that prohibit automatic snapshots.

User response:
Ensure that the specified volume group does not contain volumes that are restricted. For example, mirrored volumes and volume with a zero capacity are restricted.

CMMVC8857E **The command failed because two person integrity requires a minimum of two local SecurityAdmin users or at least one remote SecurityAdmin user group and remote authentication service enabled.**

Explanation

This error occurs when an attempt is made to enable two person integrity under the following conditions:

- Less than two Security Administrator users are available on the system. A superuser account does not contribute to this requirement.
- The remote authentication service is not enabled.
- There is no remote user group of Security Administrator role.

User response:
Ensure that the system has at least two local Security Administrator users (not including the superuser account) or has a remote Security Administrator user group and the remote authentication service enabled. Retry the command.

CMMVC8858E **The command failed because two person integrity cannot be enabled while superuser password reset is disabled.**

Explanation:
The error occurred when an attempt was made to enable two person integrity while the superuser password reset is disabled.

User response:
Run the **svctask setpwdreset -enable** command to enable superuser password reset and then retry the command.

CMMVC8859E **The command failed because two person integrity is only supported on platforms with a dedicated technician port.**

Explanation:
A dedicated technician port is mandatory to support two person integrity.

User response:
No action possible.

CMMVC8860E **The command failed because superuser locking cannot be disabled if two person integrity is enabled.**

Explanation:
The user locking policy on the system must be enabled for the superuser. You can't use the **superuser_locking** command to disable superuser locking while two person integrity is enabled.

User response:
To disable superuser locking, disable two person integrity and then run the **superuser_locking disable** command on the system.

CMMVC8861E **The command failed because the two person integrity request specified is already approved.**

Explanation:
This error occurs when two person integrityis enabled and you try to approve a request for role elevation role that is already approved.

User response:
You can't modify the duration of an approved request. However, you can cancel the approved request and then request role elevation again. You can also wait until the duration of the approved request expires and then create another request.

CMMVC8862E **The command failed because the two person integrity request specified cannot be approved by the same user that created it.**

Explanation:
This error occurred when you tried to approve your own request for a role elevation. When two person integrity is enabled, requests must be approved by another user with the Restricted Security Administrator or Security Administrator roles.

User response:
Contact a user with the Restricted Security Administrator or Security Administrator roles to

approve the request for role elevation. When the request is approved, retry the command.

CMMVC8863E **The command failed because this action compromises two person integrity which is enabled on the system.**

Explanation

This error can occur when two person integrity is enabled on the system and you attempt to run one of the following security-related commands:

- **chuser** or **rmuser** affecting a Security Administrator user.
- **chusergrp** or **rmusergrp** affecting a user group that has Security Administrator role.
- **chauthservice** disabling remote authentication.

User response:
Make sure the changes that need to be made to the user, user group, or remote authentication service do not compromise two person integrity. Check that if the changes complete successfully, there are still at least two local users of Security Administrator role (excluding superuser) or a remote user group with Security Administrator role and remote authentication enabled.

CMMVC8864E **The current user is not allowed to create new users in SecurityAdmin user groups when Two Person Integrity is enabled.**

Explanation:
This error occurs when you don't have the required authorization to create users in Security Administrator user groups when two person integrity is enabled. To create users, you must elevate your security role.

User response

User must follow the steps to gain temporary access:

1. Request a role elevation.
2. Check with a user having Restricted Security Administrator or Security Administrator roles and provide them with the user name and ID of the request created by step 1, and get that request approved.
3. After the request is approved, retry before the temporary access expires.

CMMVC8865E **The command failed because the user name specified is currently in use by an existing two person integrity request.**

Explanation:
This error occurs when you try to remove a user who has a role elevation request for two person integrity that is pending or approved. Users cannot be removed when their requests are pending or are active and not yet expired.

User response

To remove a user, wait for their role elevation request to expire and then try again. If you don't want to wait, complete the following steps to cancel the request and remove the user:

1. Locate the two person integrity request.
2. Cancel the request.
3. After the role elevation is canceled or expired, retry the command.

CMMVC8866E **The current user is not allowed to remove user members of SecurityAdmin user groups when two person pntegrity is enabled.**

Explanation:
This error occurs when you try to remove a user from a SecurityAdmin user group and you don't have an approved role elevation request.

User response

To resolve this error, complete these steps:

1. Create a two person integrity request, note down the user name and ID assigned to the request.
2. Locate a user with Security Administrator or Restricted Security Administrator access, provide the user name and ID assigned for two person integrity request and get it approved.
3. After it is approved, retry before the temporary access expires.

Related reference
“mktwopersonintegrityrequest” on page 635
Use the **mktwopersonintegrityrequest** command to request a role elevation for yourself. This command requests on behalf of the current user invoking this command.

CMMVC8867E **The command failed because the user to be removed has an outstanding two person integrity request.**

Explanation:
This error occurs when the user that you want to remove has an approved or pending role elevation request.

User response

Follow these steps to remove the user and retry the command or wait until the request expires:

1. Locate the two person integrity request.
2. Cancel the request.
3. After the role elevation is canceled or expired, retry the command.

Related reference

"chtwopersonintegrityrequest" on page 633
Use the **chtwopersonintegrityrequest** command to approve or cancel an existing role elevation request. You must have the Restricted Security Administrator role or the Security Administrator role to use this command.

CMMVC8868E The current user is not allowed to create new user groups of SecurityAdmin user role while two person integrity is enabled.

Explanation:
This error occurs because you don't have the required elevated role to create user groups with the Security Administrator user role.

User response

To resolve this error, complete these steps:

1. Create a two person integrity request, note down the user name and ID assigned to the request.
2. Locate a user with Security Administrator or Restricted Security Administrator access, provide the user name and ID assigned for two person integrity request and get it approved.
3. After it is approved, retry before the temporary access expires.

Related reference

"mktwopersonintegrityrequest" on page 635
Use the **mktwopersonintegrityrequest** command to request a role elevation for yourself. This command requests on behalf of the current user invoking this command.

CMMVC8869E The command failed because two person integrity is not enabled.

Explanation:
This error code appears when two person integrity is disabled on the system.

User response:
Security Administrator must enable two person integrity by using **chsecurity -twopersonintegrity yes** command on the system, and then retry the command.

Related information

Configuring two person integrity (TPI)

CMMVC8870E The volume was not created because a capacity of zero bytes is not allowed for Hyperswap volumes.

Explanation

A HyperSwap volume can only be created with a provisioned capacity greater than zero bytes.

User response

Either specify a non-zero size for the volume or create a different type of volume.

CMMVC8879E I/O group specified is not present.

Explanation

If you must perform discovery by using a specified I/O group, then you must specify a valid I/O group number. An invalid or non-existent I/O group was specified.

User response

Obtain a valid I/O group number by using the management GUI or the **lsiogroup** command.

CMMVC8880E Some of the nodes in the cluster/ specified IO group are not in online state.

Explanation

One or more of the nodes that are part of the specified system or I/O group are not in an active state. You must initiate discovery from all of the nodes in the system or in the specified I/O group.

User response

Check the node status by using the management GUI or the **lsnode** or **lsnodecanister** command, and ensure that all nodes are online before you restart discovery.

CMMVC8881E Discovery cannot be run as the source port ID is invalid for one or more nodes in the system or specified iogroup.

Explanation

Different models or generations of nodes might have different numbers of Ethernet ports. If the system is created with multiple nodes that have different

supported node models, the source port ID that you specify (either by using the management GUI or by using the **detectiscsistorageportcandidate** command) cannot be greater than the highest port ID on any node that is part of the system or specified I/O group.

User response

Check the number of Ethernet node ports present on each node by using the management GUI or the **lsportip** command. Run discovery again, and this time specify a port ID within the available range.

CMMVC8882E **Some of the specified source Ethernet ports are not configured for iSCSI use.**

Explanation

The Ethernet ports on all nodes in the system or specified I/O group must be in an online state before a backend controller discovery can be initiated through the port. The port can be in one of the following states:

unconfigured
No iSCSI address is assigned to the port.

configured
An IP address is assigned to the port.

User response

Use the management GUI or the **lsportip** command to identify the Ethernet port on the unconfigured node whose source port ID you specified for discovery. Use the management GUI or the **cfgportip** command to assign an IP address of the required type (IPv4 or IPv6). Ensure that all required node Ethernet ports are in an online state before you restart discovery.

CMMVC8883E **Link state is inactive for one or more source Ethernet ports in the system or specified IO group.**

Explanation

The link must be active for all node Ethernet ports that are specified with a source port ID. If any of the links are down, you cannot initiate discovery.

User response

Use the management GUI or the **lsportip** command to identify the node Ethernet port with source port ID whose **link_state** is inactive. Verify that the Ethernet cable is connected to the port for all affected nodes, then restart discovery.

CMMVC8884E **IP type mismatch between IP of source ports and target IP specified.**

Explanation

To establish an iSCSI session, the initiator and target connection endpoints must both have either IPv4 or IPv6 addresses. If you are performing cluster-wide discovery and the specified target IP type is IPv6, then an IPv6 address must be assigned to all of the initiator source ports with the specified port ID. If you are performing I/O-group-specific discovery and the specified target IP type is IPv6, then an IPv6 address must be assigned to all of the initiator source ports in the I/O group with the specified port ID.

User response

Use the management GUI or the **cfgportip** command to assign IP addresses with an IP type that matches the target IP. Alternatively, you can configure an IP address that matches the IP type of the source ports on the target controller port.

CMMVC8885E **New controller cannot be discovered as limit of managed iSCSI controllers has been reached.**

Explanation

The **detectiscsistorageportcandidate** command found a new controller with no existing established connections after the maximum of 64 iSCSI controllers was reached.

User response

Use the management GUI or the **lscontroller** command to view the list of managed controllers. Identify any controllers that you no longer need to manage. Remove the mdiskgroups that are composed of mdisks that were exported from the identified controller. Use the **rmiscsistorageport** command to remove the unused iSCSI port. After the **lscontroller** command shows fewer than the maximum of 64 controllers, retry the command to discover the new controller. Also, consider how you might consolidate iSCSI storage onto fewer controllers.

CMMVC8886E **The targetip/ targetip6 parameters do not adhere to the standard IPv4/ IPv6 formats.**

Explanation

An IPv4 or IPv6 address was used in a **targetip** or **targetip6** argument that does not adhere to standard IP address format.

User response

Reenter the command, this time using the correct format in the **targetip** or **targetip6** argument. See the following table for accepted formats.

Table 171. Accepted IP address formats

IP type	IP address list format
IPv4	1.2.3.4
Full IPv6	1234:1234:abcd:0123:0000:0000:7689:6576
Full IPv6, leading zeros suppressed	1234:1234:abcd:123:0:0:7689:6576
IPv6 with zero compression	1234:1234:abcd:123::7689:6576

CMMVC8887E Authentication with target controller using the credentials specified did not succeed

Explanation

The **username** and **chapsecret** arguments that the iSCSI initiator presented do not match the credentials expected by the iSCSI target controller.

User response

Make sure that you specify the correct credentials in the **username** and **chapsecret** arguments.

CMMVC8888E Discovery could not start as a connection could not be established to the backend controller.

Explanation

A problem in the IP network or the iSCSI target controller prevented the login or discovery request from reaching the target controller, or prevented the target controller from responding to the discovery request.

User response

Verify the following conditions:

- The IP network is properly configured
- The target ports are configured
- The target IPs are reachable

CMMVC8889E Discovery could not be completed due to problems at target or initiator.

Explanation

A problem occurred that prevented a successful login or prevented discovery output from the target controller. Possible problems include protocol errors; a transient target, initiator, or network problem; or other errors returned by the target in response to the discovery request.

User response

Correct any configuration problems at the initiator or target. If the problem persists, contact IBM Support.

CMMVC8890E Sessions cannot be established as there are no discovered controllers for the specified id.

Explanation

A row ID was not specified that indicates the target iSCSI name.

User response

Use the management GUI or the **lsiscsistorageportcandidate** command to list the discovered targets and select a valid row ID as an input to the **addiscsistorageport** command.

CMMVC8891E The specified target IP must not refer to the IP of a system management port

Explanation

An IP conflict was detected between IP addresses of source ports and targets. This error prevents discovery requests being sent from one source port to the same or different configured source ports on any node in the cluster.

User response

Use the management GUI or the **lsportip** command to locate any source port IP addresses that are the same as IP addresses configured on the target. Resolve any conflicts and retry discovery.

CMMVC8892E The specified target IP must not refer to an iSCSI port on the system

Explanation

An IP conflict was detected between assigned clustered system IP addresses and target IP addresses. This error prevents discovery requests being sent from one source port to a system management port. The clustered system IP addresses are assigned during cluster creation.

User response

Resolve conflicting IP addresses and retry the command. If necessary, you can use the management GUI or the **lssystemip** command to locate conflicting IP addresses, and use the **chsystemip** command to change system IP addresses.

CMMVC8893E Sessions cannot be established as discovery has not been done through the iogroup specified.

Explanation

Discovery must be done through an I/O group before you can use the **addiscsistorageport** command to establish sessions through that I/O group. You can run the **addiscsistorageport** command with a specific I/O group ID if you previously ran cluster-wide discovery.

User response

Use the management GUI or the **lsiscsistorageportcandidate** command to view the I/O groups through which discovery was done. Use this information to correct the command, and then resubmit it.

CMMVC8894E Limit of max IQN+IP tuples per external iSCSI storage system has been reached.

Explanation

A maximum of 128 unique IQN + IP combinations (tuples) are allowed per controller from the initiating cluster.

User response

Use the management GUI or the **lsiscsistorageport** command to view the list of established sessions. Remove sessions that are no longer required by using the management GUI or the **rmiscsistorageport** command. Retry the **addiscsistorageport** command.

CMMVC8895E Sessions cannot be established as maximum session count per node has been reached.

Explanation

The maximum of 256 initiator sessions per initiator node were established from one or more target controllers. Each invocation of the **addiscsistorageport** command creates an additional session through the nodes in the specified I/O group or cluster.

User response

Use the management GUI or the **lsiscsistorageport** command to check the total number of sessions per node. Use the management GUI or the **rmiscsistorageport** command to remove any sessions that are no longer in use. When the number of sessions is below the maximum, try again to establish the session.

CMMVC8896E One source port cannot have more than one session to the same target IQN through different target ports.

Explanation

One-to-many connectivity from source ports to backend controller target ports can create bottlenecks in the I/O path.

User response

Use the management GUI or the **lsiscsistorageport** command to view the sessions that were established between initiator ports and the target IQN+IP. Create a session to a target IQN that does not already have a connection to a source port.

CMMVC8897E One node cannot have more than one session with same target IQN+IP through different source ports.

Explanation

Many-to-one connectivity from source ports to backend controller target ports can create bottlenecks in the I/O path.

User response

Use the management GUI or the **lsiscsistorageport** command to view the sessions that were established between initiator ports

to target IQN+IP. Use a source port for the session that is not already connected to a target IQN.

CMMVC8898E Sessions cannot be torn down as the session does not exist.

Explanation

A valid existing session identifier was not specified for the tear-down.

User response

Use the management GUI or the **lsiscsistorageport** command to list the sessions that are candidates for removal.

CMMVC8901E Enter at least one throttle parameter.

Explanation

No throttle parameters (such as IOPS or bandwidth) were specified.

User response

Specify one or more throttle parameters.

CMMVC8902E Volume is invalid or does not exist.

Explanation

The ID or name of the volume that is to be throttled was not valid or was not found.

User response

Use a valid ID or name for the volume when you retry the command.

CMMVC8903E Throttle is already associated with this volume.

Explanation

The volume ID or name that you specified in a command line parameter is already associated with a throttle value.

User response

Either specify a different ID or name for the volume, or use the **chthrottle** command to change an existing throttle value for the current volume.

CMMVC8904E Throttle parameter missing or invalid.

Explanation

You entered an invalid value for one or more of the parameters, or one or more required parameters are missing, or both.

User response

Verify that all required parameters are specified and that all values are valid, and then retry the command.

CMMVC8905E Offload IO throttle already exists.

Explanation

A throttle exists for offloaded I/O. Only one offloaded I/O throttle can be created for each cluster.

User response

To change the offloaded I/O throttle parameters, use the **chthrottle** command.

CMMVC8906E Throttle object does not exist.

Explanation

The value that you specified for the **throttle_name** parameter was not found.

User response

Retry the command, this time specifying a valid throttle name.

CMMVC8910E The specified IP is a reserved or special IP and cannot be used.

Explanation

The following special or reserved IP addresses cannot be assigned to target controllers:

- Broadcast addresses

  Note: Because no subnet mask is specified, the system does not detect the incorrect use of subnet broadcast addresses.
- Multicast addresses
- Unspecified addresses ("0.0.0.0" for IPv4, "::" for IPv6)
- Loopback addresses
- Link-local addresses

Use APIPA addresses with caution because the system cannot verify them.

User response

Configure a supported IP address on the target controller and retry discovery.

CMMVC8911E **Limit of IQNs per target controller has been reached.**

Explanation

A maximum of 64 iSCSI Qualified Names (IQNs) per target iSCSI controller is permitted.

User response

Use the management GUI or the **lsiscsistorageport** command to view the sessions for a controller. Remove sessions to IQNs that are no longer in use by using the management GUI or the **rmiscsistorageport** command. Rerun the **addiscsistorageport** command.

CMMVC8912E **The specified Ethernet port is not configured for connectivity with external storage systems.**

Explanation

The Ethernet port must be configured to allow connectivity to a backend storage controller before you can begin discovery or establish a session through the port.

User response

Use the management GUI or the **cfgportip** command to enable storage system connectivity. Use the **storage** flag for IPv4 addresses and the **storage_6** flag for IPv6 addresses.

.

CMMVC8913E **The port IP address cannot be unconfigured as there are active sessions established to backend controller(s) through the port.**

Explanation

An attempt was made to remove an IP address from a port that was used to establish backend controller sessions. Removing an IP address through which sessions are established puts mdisks into a degraded state and is not permitted.

User response

Use the management GUI or the **rmiscsistorageport** command to remove sessions through the source port before you remove an IP address that is used for backend controller connectivity. Alternatively, if the intent is only to remove the IP and then configure a new IP, you can achieve the same goal by using the **cfgportip** command with a **-force** flag.

CMMVC8914E **Storage flag cannot be reset as the IP is already in use for backend controller(s) connectivity.**

Explanation

If you set the **storage** or **storage_6** flag to yes, the IPv4 or IPv6 address for the port can be used to establish connections to the backend iSCSI controller. After you use the **addiscsistorageport** command to establish connectivity from the source port to one or more backend controllers, then you cannot set the **storage** or **storage_6** flag to no unless all the sessions to backend controllers are removed.

User response

Verify the information, make a correction, and resubmit the command.

CMMVC8915E **Changing the IP address attributes may cause mdisks to go into a degraded state for some time.**

Explanation

If a source IP address is in use for backend controller connectivity, changing the IP address, mask, or gateway causes the existing sessions to be removed and new sessions to be established. During this phase, the managed disks that are visible though the source port that is being reconfigured go into a degraded state until the new sessions are established.

User response

If you did not intend to reconfigure, reenter the command with the correct parameters. If you do want to reconfigure, an administrator can use the management GUI or the **cfgportip** command with the **-force** flag to go ahead with the reconfiguration. Even if no IP address is already assigned to the port that you are trying to use for the **cfgportip** command, use the **-force** flag. No sessions are disconnected because no sessions exist.

CMMVC8919E **The command failed because the requested transition is invalid.**

Explanation

An invalid transition was attempted, such as from disabled to enabled.

User response

Use multiple steps to make the required transition. For example, move from disabled to transitional to enabled. Refer to the help for specific configuration instructions.

CMMVC8920E **The command failed because one or more host ports are currently failed over.**

Explanation

A state transition was attempted while a failover was in progress. In practice, this error is rarely seen, due to the requirement that both nodes be online. A small timing window for the error can occur after a node has just unpended and before a controlled node shutdown.

User response

Wait until the **lstargetportfc** command reports that all ports in this I/O group are online on their owner nodes, and then retry.

CMMVC8922E **The command failed because the volume copies would not be in two independent sites.**

Explanation

In a stretched system topology, you can add a volume copy only if the result is a volume copy in each of site 1 and site 2.

User response

Ensure the existing volume copy has a well defined **site** argument (site 1 or site 2), and retry by adding the new volume copy to the other site.

CMMVC8923E **The command failed because a volume in a Metro Mirror or Global Mirror relationship cannot have copies in two sites using HyperSwap.**

Explanation

In a HyperSwap system topology, you cannot add a volume copy to a different site if the existing volume is in a Metro Mirror or Global Mirror relationship.

User response

If you specified the wrong volume, correct the command and retry. If you do need to add the volume copy to a different site, first remove the Remote Copy relationship by using the **rmrcrelationship** command, then retry the command.

CMMVC8924E **The command failed because the volume already has a copy in the specified site.**

Explanation

In a stretched system topology, a volume copy can be added only if the result is a volume copy in each of site 1 and site 2

User response

Retry by specifying a storage pool with a different site for the new volume copy.

CMMVC8925E **The command failed because the cache state for an image mode volume or copy is corrupt.**

Explanation

The data on the image mode mdisk might be different from the data on the image mode volume.

User response

Use the **recovervdisk** command to acknowledge the volume data loss, then retry the command. Alternatively, specify the **-discardimage** parameter to force the removal of the image mode copy.

CMMVC8926E **The volume copy was not deleted because a change volume is providing the consistent image for a copy on the other site.**

Explanation

An attempt was made to remove an image mode copy of a HyperSwap volume while a change volume was still in use.

User response

Wait until the change volume is no longer in use, then retry the command. Alternatively, specify the **-discardimage** parameter to force the removal of the volume copy, which might result in data loss.

CMMVC8927E **The requested operation cannot be applied to sra_monitor and sra_privileged.**

Explanation

An attempt was made to delete or modify support assistance users by using the **rmuser** or **chuser** commands.

User response

Remove support assistance users by using the **chsra -disable** command.

Note: The **chsra -disable** command removes sra_monitor and sra_privileged users and also disables local and remote support assistance if configured.

CMMVC8928E **User names sra_monitor and sra_privileged are reserved, please try again with a different name.**

Explanation

An attempt was made to add a user with the reserved name sra_monitor or sra_privileged.

User response

Pick a different user name and retry the command.

CMMVC8929E **Cannot stop email when support assistance is enabled.**

Explanation

An attempt was made to stop the email service while support assistance is enabled.

User response

Disable support assistance by entering the **chsra -disable** command, and then try again to stop the email service.

Note: The **chsra -disable** command removes sra_monitor and sra_privileged users and also disables local and remote support assistance if configured.

CMMVC8930E **The command failed because the support assistance feature has not been enabled.**

Explanation

An attempt was made to enter support assistance commands such as **chsra -disable** or **chsra -updatetoken** before support assistance was enabled.

User response

Enable support assistance by entering the **chsra -enable** command, and then retry the command that failed.

CMMVC8931E **Cannot remove email user of usertype support when support assistance is enabled.**

Explanation

An attempt was made to remove a support-type email user by using a command such as **rmemailuser** while support assistance is enabled.

User response

Disable support assistance by using the **chsra -disable** command, and then try again to remove the email user.

Note: The **chsra -disable** command removes sra_monitor and sra_privileged users and also disables local and remote support assistance if configured.

CMMVC8932E **Cannot delete email server configurations when support assistance is enabled.**

Explanation

An attempt was made to remove an email server by entering a command such as **rmemailserver** while support assistance is enabled.

User response

Disable support assistance by running the **chsra -disable** command, and then try again to remove the email server.

Note: The **chsra -disable** command removes sra_monitor and sra_privileged users and also disables local and remote support assistance if configured.

CMMVC8933E **Limit of host clusters has been reached.**

Explanation

An attempt was made to create a host cluster when the maximum number of host clusters had already been created. The maximum number of host clusters is 512.

User response

Remove one or more existing host clusters before you add a new one.

CMMVC8934E Host already belongs to a host cluster.

Explanation

An attempt was made to add a host to a host cluster when the host already belongs to another host cluster.

User response

- If you specified the wrong host, retry the command and specify a different host.
- Otherwise, remove the host from the host cluster to which it currently belongs, then retry the command.

CMMVC8935E Host cluster does not exist.

Explanation

A host cluster was specified that does not exist.

User response

Retry the command with a different host cluster.

CMMVC8936E Host cluster has shared mappings. Use -keepmappings or -removemappings flag.

Explanation

An attempt was made to delete a host cluster that has one or more shared mappings without setting a flag to specify how those mappings are to be treated.

User response

Retry the command and specify one of the following flags:

-removemappings
To delete the shared mappings along with the host cluster.

-keepmappings
To keep the shared mappings as private mappings for the host.

CMMVC8937E Limit of hosts in the host cluster has been reached.

Explanation

An attempt was made to add a host to a host cluster that is full. A host cluster can include a maximum of 128 hosts.

User response

- Add the host to a different host cluster.
- Alternatively, remove one or more hosts from the host cluster before you add another.

CMMVC8938E Host is not part of the host cluster.

Explanation

An attempt was made to carry out one of the following actions, specifying a host that did not belong to the specified host cluster:

- Remove the host
- Convert a shared mapping to a private mapping

User response

Retry the command and specify a different host, a different host cluster, or both.

CMMVC8939E Some source Ethernet ports are temporarily unavailable due to failover.

Explanation

In the case of a node-pend event, the iSCSI IP addresses assigned to node Ethernet ports fail over to the partner node. While the IP addresses are in a failover state, discovery cannot be initiated from the node those addresses are assigned to, even though the node ports show as configured. Discovery can be initiated when the IP addresses fail back. The IP addresses fail back immediately if the node unpends 5 minutes or more after the pend event. If the node unpends less than 5 minutes after the pend event, the IP addresses fail back only after pend time + 5 minutes.

User response

Use the management GUI or the **lsportip** command to verify that the IP addresses failed back to the owning node before retrying discovery.

CMMVC8940E Discovery has not been initiated from nodes of specified site.

Explanation

The **addiscsistorageport** command was not run with the same site ID argument that was used to run discovery.

User response

Run the management GUI or the **addiscsistorageport** command with the same site ID argument that was used to run discovery with the **detectiscsistorageportcandidate** command.

CMMVC8941E **Sessions cannot be added to the target controller due to conflict with existing connectivity or discovery output.**

Explanation

One of the following situations occurred:

- Discovery was run with one value for the **iogroup** parameter and the **addiscsistorageport** command was run with a different value for the **iogroup** parameter.
- Discovery was run with a value for the **iogroup** parameter and the **addiscsistorageport** command was run cluster-wide (**iogroup** not specified).
- Discovery was run cluster-wide and the **addiscsistorageport** command was run with a value for the **iogroup** parameter.
- Sessions were already established through one I/O group, and the **addiscsistorageport** command was run to add connectivity through another I/O group or cluster-wide.
- Sessions were already established cluster-wide, and the **addiscsistorageport** command was run to add connectivity through an I/O group.
- Sessions were already established through a site, and the **addiscsistorageport** command was run to add connectivity through an I/O group or cluster-wide.
- Sessions were already established through an I/O group or were established cluster-wide, and the **addiscsistorageport** command was run to add connectivity through a site. (You cannot mix I/O groups of different topologies in the same system.)

User response

Run the management GUI or the **addiscsistorageport** command with the same site ID parameter that was used to run discovery with the **detectiscsistorageportcandidate** command.

CMMVC8942E **The volume copy was not deleted because the volume is part of a FlashCopy mapping. All FlashCopy mappings must be removed from the volume before this copy can be deleted.**

Explanation

An attempt was made to remove a volume copy when the volume had FlashCopy mappings.

User response

Remove all FlashCopy mappings from the volume, then try again to delete the copy.

CMMVC8943E **The command failed because the consistent image on an image mode copy is being provided by a FlashCopy map.**

Explanation

An attempt was made to remove an image mode volume or volume copy that is the target of an active (copying) FlashCopy mapping. While the FlashCopy mapping is copying, the data on the image mode mdisk might be different from the data on the image mode volume.

User response

Wait for the FlashCopy operation to complete, then retry the command. Alternatively, specify the **-discardimage** parameter to force the delete operation, which might result in data loss.

CMMVC8944E **The command failed because the volume already has copies in two sites using HyperSwap. Use the addvdiskcopy command to add a second copy in the same site.**

Explanation

In HyperSwap system topology, the **addvolumecopy** command cannot be used to add a mirrored copy in the same site as an existing volume copy.

User response

Use the **addvdiskcopy** command to add a second copy in the same site.

CMMVC8945E **The command failed because it would allow access to an image mode copy in a Metro Mirror or Global Mirror relationship that does not contain consistent data.**

Explanation

An attempt was made to remove an image mode volume or volume copy that is a secondary in a Metro

Mirror or Global Mirror relationship and the data on the volume is not consistent.

User response

Wait for the relationship to synchronize the volume from the primary volume, then retry the command. Alternatively, specify the **-discardimage** parameter to force the delete operation.

CMMVC8946E **The command failed because an image mode copy is not synchronized with the host-accessible copy.**

Explanation

An attempt was made to remove an image mode volume or volume copy, and the data on an image mode copy is not synchronized with the host-accessible copy.

User response

Wait for the volume copy to resynchronize, then retry the command. Alternatively, specify the **-discardimage** parameter to force the delete operation.

CMMVC8947E **The target controller is not in the storage layer.**

Explanation

The target controller is configured in the replication layer.

User response

Consider whether the target controller is configured correctly. If appropriate, use the management GUI or the **chsystem -layer storage** command on the target controller to configure the controller in the storage layer. If the target controller is in the correct layer, check that the IP address that you specified points to the correct target. Correct the layer for the target controller or specify the correct IP address and rerun discovery by using the **detectiscsistorageportcandidate** command.

CMMVC8952E **The action failed because there was a mismatch with authentication credentials previously specified for same target.**

Explanation

The **addiscsistorageport** command configures a session between an initiator port and target IP addresses for a target iSCSI Qualified Name (IQN). If you run the command again for the same target parameters, it returns successfully without initiating any action as the sessions was established after the first invocation. However, if you retry the command with a different **username** argument, **chapsecret** argument, or both, then the command fails because the initiator cannot verify which credentials are correct.

User response

To change the target credentials, complete the following steps by using the GUI or the specified command:

1. Use the **lsiscsistorageport** command to list the previous set of sessions that were established.
2. Use the **rmiscsistorageport** command to remove the sessions that match the source port, IP address, and IQN for the target.
3. Initiate a new discovery and establish sessions with new credentials.

CMMVC8953E **There are no nodes in the site and/or I/O group specified.**

Explanation

Nodes in the cluster were not configured to be part of the site before discovery of the iSCSI backend controller was initiated.

User response

If the system topology is a stretched cluster or HyperSwap configuration, use the management GUI or the **addnode** or **chnode** command to add the nodes to the specified site. Restart discovery.

CMMVC8954E **The system is not in the replication layer.**

Explanation

You can initiate discovery from a system in the replication layer only.

User response

Consider whether the source system is configured correctly. If appropriate, use the management GUI or the **chsystem -layer replication** command on the source system to configure the system in

the replication layer. Rerun discovery by using the **detectiscsistorageportcandidate** command.

CMMVC8956E The command failed because the volume is the auxiliary volume in an active-active relationship.

Explanation

The specified volume is an auxiliary volume in an active-active Remote Copy relationship. You cannot throttle an auxiliary volume.

User response

Specify a different volume or break the Remote Copy relationship before you create or modify the throttle on this volume.

CMMVC8957E The command failed because the volume specified does not exist.

Explanation

A volume ID or name was specified that does not exist.

User response

View the list of volumes by means of the GUI or the **lsvdisk** command to ensure that you specified the correct volume ID or name. For a HyperSwap volume involved in an active-active relationship, the volume ID and name are the same as those of the master volume.

CMMVC8958E The command failed because the volume has vVol ownership and may not have copies in two sites using HyperSwap.

Explanation

A volume created as a vVol may not be created as, or converted to, a HyperSwap volume that is involved in an active-active relationship.

User response

Use a different volume as a HyperSwap volume.

CMMVC8959E The command failed because the master volume has vVol ownership and may not participate in a Remote Copy relationship.

Explanation

A volume created as a vVol cannot participate in a Remote Copy relationship as the master or auxiliary volume.

User response

Use VMware vCenter to manage system objects like volumes and pools.

CMMVC8960E The command failed because the specified node is the last node of the cluster.

Explanation

Only one node is left in the system. If this node leaves the system, the system is removed and cannot process the command. This error never occurs with **-action replace**.

User response

To put the last node of the system into service state, run **satask startservice**. This command does not run at the system level. The last node of a system cannot be exchanged with a spare.

CMMVC8961E The action failed because the node hardware is incompatible with the previous node.

Explanation

A candidate node has the same World Wide Node Name as the node that you are trying to swap, but is a different node type, or has a different hardware configuration.

User response

Change your Fibre Channel configuration, and then retry the command.

CMMVC8962E The action failed because the specified node is not online.

Explanation

An attempt was made to run the **swapnode -action service** command on a node that is not online.

User response

First, make sure that you specified the correct node. If not, retry the action and specify an online node.

If you specified the correct node, then no further action is required because the node is already offline.

CMMVC8964E The command failed because the support assistance feature has already been enabled.

Explanation

The **chsra -enable** command was entered when support assistance was already enabled.

User response

Enter the **chsra -updatetoken** command to update the shared token or enter the **chsra -disable** command to disable support assistance.

CMMVC8965E **The command failed because volumes are dependent on the specified node and -permitofflinevolumes was not specified.**

Explanation

Performing this operation may cause a host I/O outage. This error never occurs with -action replace.

User response

Correct the redundancy issue, or, if an I/O outage is acceptable, retry the command with the **-permitofflinevolumes** parameter.

CMMVC8966E **This node cannot be added to the IO Group due to having lower configured RAM than required.**

Explanation

An attempt was made to add a node to an I/O group where the other nodes in the group have more memory than the added node. It is also possible that faulty RAM is preventing the amount of memory in the node from being correctly detected. All nodes in an I/O group must have the same amount of configured RAM.

User response

Add the node to a different I/O group or increase its memory to match that of the other nodes in the group and retry the command.

CMMVC8967E **The command failed because the auxiliary volume has vVol ownership and may not participate in a Remote Copy relationship.**

Explanation

A volume created as a vVol cannot participate in a Remote Copy relationship as the master or auxiliary volume.

User response

Use VMware vCenter to manage system objects like volumes and pools.

CMMVC8968E **This operation cannot be completed because iSCSI Initiator sessions are present.**

Explanation

The operation is not permitted when iSCSI initiator sessions are present and you try to make one of the following changes:

- Changing from the replication layer to the storage layer on a system that is acting as an iSCSI initiator
- Changing from the storage layer to the replication layer on a system that is acting as an iSCSI target

User response

Follow these steps to remove the initiator sessions:

1. Use the management GUI or the **lsiscsistorageport** command to identify the iSCSI initiator sessions on the system that is acting as an iSCSI initiator.
2. Use the management GUI or the **rmiscsistorageport** command to remove all iSCSI initiator sessions.
3. After you remove all initiator sessions, retry the command.

CMMVC8969E **This operation cannot be completed because migration to different parent storage pool is not supported.**

Explanation

You cannot migrate volumes from a child storage pool to a different parent storage pool or to a child storage pool that has a different parent than the source.

User response

First migrate the volume from the child pool to its own parent, then migrate it from that parent pool to that of a different parent.

CMMVC8970E **The storage pool was not deleted because there are volumes associated with the storage pool. A forced deletion is required.**

Explanation

You cannot delete a storage pool that has any volumes still associated with it.

User response

Remove the associated volumes and retry the command, or use the **-force** flag:

```
rmmdiskgrp -force storage_pool
```

where *storage_pool* is either the name or the ID of a storage pool.

CAUTION: Use of the **-force** flag can result in data loss. Contact IBM Support before using this flag.

CMMVC8971E The command failed because data in the cache has not been committed to disk.

Explanation

The command failed because data in the cache has not been committed to disk.

User response

Check your command to ensure you have specified the correct volume and target. Make the correction and resubmit the command. Otherwise, investigate why the data has not been committed and how the data must be committed.

CMMVC8972E The command failed because the MDisk is an array.

Explanation

An attempt was made to change the encryption settings of an MDisk that is an array. The **chmdisk** command applies only to external MDisks.

User response

You cannot directly convert the array from unencrypted to encrypted or vice versa. Instead, you must complete one of the following actions:

- Specify a different MDisk and retry the command.
- Delete and re-create the array with the new encryption setting.

CMMVC8973W The IO group *io_group_name* has been restored with ID *new_id* instead of *old_id* .

Explanation

This situation can occur when the configuration node is different from the node that was used to create the original cluster. This change affects the SCSI Inquiry value for the I/O group.

User response

This message is a warning only. No user response is required.

CMMVC8974E The action failed because of incompatible code.

Explanation

The code version on one or more nodes is incompatible with the new version.

User response

Refer to the compatibility requirements for the code version you are adding. Update the cluster to meet the compatibility requirements, and then perform the upgrade.

CMMVC8975E The node could not be added because of incompatible code. The status code is *status_code* .

Explanation

An attempt was made to add a node with code that has a different revision level than the code in the cluster.

User response

Update the code on the rejected node to the same level as that on the cluster to which it will be added, and then resubmit the command.

CMMVC8976E The cluster was not modified because the IP address is not valid.

Explanation

An attempt was made to change the IP address of a cluster to an address that is not valid.

User response

Correct the address and reissue the command.

CMMVC8977E The action failed because the directory that was specified was not one of the following directories: /dumps, /dumps/iostats, /dumps/iotrace, /

dumps/feature, /dumps/config, /dumps/elogs, /dumps/ec or /dumps/pl.

Explanation

An attempt was made to clear a file from, or copy a file to, a directory that is not valid.

User response

Ensure that the command accesses a valid directory and try again.

CMMVC8978E **The action failed as the resulting disk size would be less than, or equal to, zero.**

Explanation

An attempt was made to shrink a disk, however the resulting size would have been less than or equal to zero.

User response

Check your command to ensure that you have the correct disk size. You can also check the **shrinkvdisksize** command documentation for additional information. Make the correction and resubmit the command.

CMMVC8979E **The action failed as the resulting disk size would be less than, or equal to, zero.**

Explanation

An attempt was made to shrink a volume to a size that is below the permitted minimum.

User response

Check your command to ensure that you have the correct disk size. Make the correction and resubmit the command. For more information, see the documentation for the **shrinkvdisksize** command.

CMMVC8980E **Metadata recovery can not use the provided MDisk id - invalid or destroyed.**

Explanation

Metadata recovery cannot use the provided MDisk ID, which is not valid or refers to a disk that was destroyed.

User response

Correct the specified MDisk and resubmit the command.

CMMVC8981E **The update failed as a file containing the code for the specified MCP version was not found.**

Explanation

Two files are required to successfully complete a code update. One file contains the files that make up the base operating system, and the other file contains the code. This message appears if the OS version is incompatible with the code.

User response

Upload two compatible files and resubmit the command.

CMMVC8982E **The action failed because the volume is part of a Remote Copy relationship.**

Explanation

An action was performed against a volume that is part of a Remote Copy relationship.

User response

Remove the volume from the Remote Copy relationship and then resubmit the command.

CMMVC8983E **The action failed because the volume is part of a FlashCopy mapping.**

Explanation

An action was performed against a volume that is part of a FlashCopy mapping.

User response

Remove the volume from the FlashCopy mapping before reissuing the command.

CMMVC8984E **The FlashCopy mapping or consistency group could not be started in a reasonable time. The mapping or group is instead being prepared.**

Explanation

The FlashCopy mapping or consistency group could not be started in a reasonable time. The mapping or group is instead being prepared.

User response

Resubmit the command.

CMMVC8985E The command has failed because a virtual medium error exists on the image mode volume or copy.

Explanation

When you submit this command, you cannot specify an image mode volume that has a virtual medium error on the volume or on any copy of the volume because the medium errors cannot be maintained on the ejected MDisk image copy.

User response

If an exact image copy is required, ensure that there is no virtual medium error on the image mode volume that you specify or on any of its copies, and resubmit the command.

If an exact copy is not required, you can use the -force option of the command, but all of the virtual medium errors will be lost.

CMMVC8986E The command failed as a migrate to image was in progress.

Explanation

An attempt was made to execute a command against a volume that was involved in a migrate to image operation.

User response

Wait for the migration to complete and reissue the command.

CMMVC8987E You are trying to recover region data that was created by a code level different from the one you are currently running on the node.

Explanation

You are trying to recover region data that was created by a code level different from the one you are currently running on the node.

User response

Inform your administrator of this error. You might need to update the code level on your server. Wait until the server is updated before you resubmit the command.

CMMVC8988E Failed to recreate the cluster you are trying to rebuild.

Explanation

An attempt was made to rebuild a cluster, but the attempt failed.

User response

Check your command. The source and target names might not match. Make the correction and resubmit the command.

CMMVC8989E The FlashCopy mapping was not created or modified because the consistency group already contains the maximum number of mappings.

Explanation

An attempt was made to create a FlashCopy mapping in, or move a FlashCopy mapping to, a consistency group that has the maximum number of FlashCopy mappings that it can contain.

User response

Create or move the FlashCopy mapping to another consistency group or remove an existing FlashCopy mapping from the desired group and then reissue the command.

CMMVC8990E The Remote Copy relationship was not created because the master or auxiliary volume is a member of a Remote Copy relationship.

Explanation

An attempt was made to create an active-active relationship with a volume copy in another site. This relationship is not supported if the volume is already in a remote copy relationship.

User response

Make sure that you specified the correct volume. If so, delete the existing relationship and try again.

CMMVC8991E The maximum number of hosts for the cluster is already configured.

Explanation

You must remove at least one host definition before you can resubmit the command.

User response

Determine whether the action is required.

If the action is required, review the current configuration to determine whether any current host definitions are not required. Remove at least one host definition that is not required, and resubmit the command.

CMMVC8992E The maximum number of host/IO group pairs for the cluster is already configured.

Explanation

You must remove at least one host I/O group pair definition before you can resubmit the command.

User response

Determine whether the action is required.

If the action is required, review the current configuration to determine whether any current host I/O group pair definitions are not required. Remove at least one host I/O group pair definition that is not required, and resubmit the command.

CMMVC8993E The maximum number of WWPNs plus iSCSI qualified names (IQNs) for the cluster is already configured.

Explanation

The command cannot be initiated because the maximum number of WWPNs plus iSCSI qualified names for the cluster has been reached.

User response

Determine whether the action is required.

If the action is required, review the current configuration to determine whether any current WWPN or iSCSI qualified name definitions are not required. Remove at least one WWPN definition or iSCSI qualified name that is not required, and resubmit the command.

CMMVC8994E The maximum number of hosts for one or more IO groups is already configured.

Explanation

You must remove at least one host I/O group pair definition from the I/O group that you have specified before you can resubmit the command.

User response

Determine whether the action is required.

If the action is required, review the current configuration to determine whether any current host I/O group pair definitions for the I/O group that you have specified are not required. Remove at least one host I/O group pair definition that is not required from the I/O group that you have specified, and resubmit the command.

CMMVC8995E The maximum number of WWPNs for one or more IO groups is already configured.

Explanation

You must remove at least one WWPN definition from the I/O group that you have specified before you can resubmit the command.

User response

Determine whether the action is required.

If the action is required, review the current configuration to determine whether any current WWPN definitions for the I/O group that you have specified are not required. Remove at least one WWPN definition that is not required from the I/O group that you have specified, and resubmit the command.

CMMVC8996E The maximum number of WWPNs for the host is already configured.

Explanation

You must remove at least one WWPN definition for the host that you have specified before you can resubmit the command.

User response

Determine whether the action is required.

If the action is required, review the current configuration to determine whether any current 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.

CMMVC8997E **The host does not belong to one or more of the IO groups specified or inferred.**

Explanation

The host does not belong to one or more of the I/O groups specified or inferred.

User response

Specify a host and I/O group combination that is currently defined, and resubmit the command.

CMMVC8998E **The host already belongs to one or more of the IO groups specified.**

Explanation

The host already belongs to one or more of the I/O groups specified.

User response

Check your command. Change the host name if applicable and resubmit the command.

CMMVC8999E **An IO group cannot be removed from a host because of one or more associated volumes.**

Explanation

An I/O group cannot be removed from a host because of one or more associated volumes.

User response

Check your command and ensure that you have specified the correct I/O group. Make the correction if one is needed and resubmit.

CMMVC9000E **The action was not completed because the cluster has reached the maximum number of extents in storage pools.**

Explanation

The cluster reached the maximum number of extents in the storage pool; therefore, the action did not complete. An attempt was made to use additional extents, for example by creating or expanding a volume. The action cannot be initiated because it would cause the maximum number of extents for a cluster to be exceeded.

User response

Free up extents by deleting other volumes, and resubmit the command.

CMMVC9001I **The package installed successfully.**

Explanation

The package installed successfully.

User response

This message is for your information only. No response is required.

CMMVC9002E **The cluster was recovered and the CLI functionality is limited until the cause of the failure is determined and any corrective action taken. Contact technical support for assistance.**

Explanation

The cluster was recovered and the CLI function is limited.

User response

Contact IBM Support.

CMMVC9003E **The action failed as the SSH key has been revoked.**

Explanation

The action failed because the SSH key was revoked.

User response

Check your command and correct the number that you specified for the SSH key. Resubmit after you make the correction.

CMMVC9004E **The action failed as the SSH key index (SSH_LABEL_ID) is invalid.**

Explanation

The action failed because the SSH key index (SSH_LABEL_ID) is not valid.

User response

Correct the SSH key index and retry the command.

CMMVC9006E **This operation cannot be performed because the cluster is**

currently cancelling the previous update command.

Explanation

This operation cannot be performed at the same time that the previous update command is being canceled.

User response

Wait until the previous update command stops running and resubmit the command.

CMMVC9007E **This operation cannot be performed because, either an update has not been started, or an update is in progress but is not in a state where it can be aborted.**

Explanation

This operation cannot be performed because the update is progressing.

User response

Wait until the update finishes and resubmit your command.

CMMVC9008E **The update cannot be cancelled because at least one node has already committed to a new code level.**

Explanation

You cannot cancel the update because one or more nodes already committed to a new code level.

User response

Contact IBM Support.

CMMVC9009E **An invalid response has been entered. The command has not been executed. Input is case sensitive. Enter either yes or no.**

Explanation

A response that is not valid was entered. The command was not run.

User response

Enter either yes or no.

CMMVC9010E **The command has not completed. A limited availability parameter has been entered without the required environment setting being set.**

Explanation

The command did not finish. A limited availability parameter was entered without the required environment setting.

User response

Check your command. If the command is correct, contact your administrator to determine whether you need an environment setting so that the command can run. Resubmit the command when the problem is resolved.

CMMVC9011E **The command failed as the remote cluster does not support global mirror.**

Explanation

The command failed because the remote cluster does not support Global Mirror.

User response

Check your command and ensure that you specified the correct cluster. Make the correction and resubmit the command. If you specified the correct cluster, research why it does not support Global Mirror.

CMMVC9012E **The copy type differs from other copies already in the consistency group.**

Explanation

The copy type differs from other copies already in the consistency group.

User response

Ensure that the copy type of the mapping that you are attempting to add is the same copy type as the mappings in the consistency group to which you are attempting to add the mapping, and resubmit the command.

CMMVC9013E **The FlashCopy mapping was not prepared because the mapping or consistency group is in the stopping state. The mapping or consistency group must first complete the stop operation and then be prepared**

Explanation

You cannot prepare a FlashCopy mapping or consistency group when the FlashCopy mapping or consistency group is in the stopping state. If you want to prepare a FlashCopy mapping or consistency group, the FlashCopy mapping or consistency group must be in the Stopped or idle_or_copied state.

User response

Wait until the FlashCopy mapping or consistency group reaches the Stopped or idle_or_copied state and then resubmit the command.

CMMVC9014E **The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the stopping state.**

Explanation

You cannot modify the consistency group of a FlashCopy mapping when the FlashCopy mapping is in the stopping state. If you want to modify the consistency group of a FlashCopy mapping, the FlashCopy mapping must be in the Stopped or idle_or_copied state.

User response

Wait until the FlashCopy mapping reaches the Stopped or idle_or_copied state and then resubmit the command.

CMMVC9015E **The FlashCopy mapping was not deleted because the mapping or consistency group is in the stopping state. The mapping or consistency group must be stopped first.**

Explanation

You cannot delete a FlashCopy mapping or consistency group when the FlashCopy mapping or consistency group is in the stopping state. If you want to delete a FlashCopy mapping or consistency group, the FlashCopy mapping or consistency group must be in the Stopped or idle_or_copied state.

User response

Wait until the FlashCopy mapping or consistency group reaches the Stopped or idle_or_copied state and then resubmit the command.

CMMVC9016E **The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the stopping state. The mapping or consistency group must first complete the stop operation and then be prepared.**

Explanation

You cannot start a FlashCopy mapping or consistency group when the FlashCopy mapping or consistency group is in the stopping state. If you want to start a FlashCopy mapping or consistency group, the FlashCopy mapping or consistency group must be in the Prepared state.

User response

Wait until the FlashCopy mapping or consistency group reaches the Stopped or idle_or_copied state and then prepare the FlashCopy mapping or consistency group before you start it.

CMMVC9017E **The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is already in the stopping state.**

Explanation

A Stop FlashCopy mapping or consistency group task was already submitted and is still in progress. When the task completes successfully, the FlashCopy mapping or consistency group state changes to Stopped.

User response

Wait for the existing task to complete. The group state changes to Stopped automatically.

CMMVC9018E **The FlashCopy mapping was not created because the source volume cannot be the target for a FlashCopy mapping.**

Explanation

A volume cannot simultaneously be both the source of a FlashCopy mapping and the target of a FlashCopy mapping. A source volume was specified that is currently defined as the target of a FlashCopy mapping.

User response

You have two options. One option is to specify a different source volume and resubmit the command. The other option is to delete the existing FlashCopy mapping that defines the source volume that was specified as the target volume, and resubmit the command.

CMMVC9019E **The FlashCopy mapping was not created because the source volume is already in the maximum number of FlashCopy mappings.**

Explanation

The number of FlashCopy mappings in which a volume can be defined as the source volume is limited. The source volume that was specified cannot be defined to another FlashCopy mapping because it is already defined as the source volume to the maximum number of FlashCopy mappings.

User response

You have two options. One option is to specify a different source volume and resubmit the command. The other option is to delete one of the existing FlashCopy mappings that contains the source volume and resubmit the command.

CMMVC9020E **The FlashCopy mapping was not created because the target volume is already a source volume in a FlashCopy mapping.**

Explanation

A volume cannot simultaneously be both the source of a FlashCopy mapping and the target of a FlashCopy mapping. The target volume that was specified is currently defined as the source of a FlashCopy mapping.

User response

You have two options. One option is to specify a different target volume and resubmit the command. The other option is to delete all of the existing FlashCopy mappings that contain the target volume that was specified and resubmit the command.

CMMVC9021E **The FlashCopy mapping was not created because the target volume is already a target volume in a FlashCopy mapping.**

Explanation

A volume cannot simultaneously be the target of more than one FlashCopy mapping. The target volume that was specified is currently defined as the target of another FlashCopy mapping.

User response

You have two options. One option is to specify a different target volume and resubmit the command. The other option is to delete the existing FlashCopy mapping that contains the target volume that was specified and resubmit the command.

CMMVC9022E **The command failed because the authorization table is full.**

Explanation

The command failed because the authorization table is full.

User response

Check with your administrator to verify the status of the authorization table. You must wait until the table is adjusted to resubmit the command.

CMMVC9023E **The command failed because the authorization record was not found or is already set to the default role.**

Explanation

The command failed because the authorization record was not found or is already set to the default role.

User response

Check your command for the role you specified. If the role is not set up, check with your administrator to verify the status. Resubmit the command after the role is set up in the system.

If you want to set the role as the default, no further action is required.

CMMVC9024E **The command failed because the authorization record is not set to the default role. Use rmauth to set the default role.**

Explanation

The command failed because the authorization record is not set to the default role.

User response

Use the **rmauth** command to set the default role.

CMMVC9025E **The command failed because the specified role was not found.**

Explanation

The command failed because the specified role was not found.

User response

Check the command and correct the specified role. If you believe the specified role exists, check with your administrator for clarification. Make the correction and resubmit the command.

CMMVC9026E **The command failed authorization because the session SSH key is invalid or was deleted.**

Explanation

The command failed authorization because the session SSH key is not valid or was deleted.

User response

Check the command and ensure that you specified a valid SSH key. Make the correction and resubmit the command.

CMMVC9027E **The task has failed because the user's role is not authorized to submit the command.**

Explanation

This error occurs when an attempt is made to run a command by a user who is not authorized to run the command or to run the command with certain parameters, based on the user's role. For example, a user that has a role of Monitor cannot create a volume. Some commands can be run only by users with the Administrator role.

User response

Log in as a user that has a role that is authorized to submit the task. Alternately, change the role of the user account that you are using to a role that is authorized to submit the task. Then, resubmit the task.

CMMVC9028E **The command failed because the specified SSH key was not found. NOTE This command must specify an admin key.**

Explanation

The command failed because the specified SSH key was not found. This command must specify an admin key.

User response

Recheck your command and ensure that you specified `admin` as the SSH user name. Make the correction and resubmit.

CMMVC9029E **The command cannot set the authorization record to the default role. Use rmauth to set the default role.**

Explanation

The command cannot set the authorization record to the default role.

User response

Use the **rmauth** command to set the default role.

CMMVC9030E **The command failed because the SSH key already exists or there is a duplicate SSH key.**

Explanation

An attempt was made to add an SSH key that already exists. A different authorization level might be associated with the key.

User response

If the existing SSH key of the same type does not have the authority level that you require, add a different SSH key.

CMMVC9031E **The command failed because one of the nodes in the specified I/O group was offline when a memory change was attempted.**

Explanation

All nodes in the I/O group must be online when you enter this command. This error indicates that at least one node is not currently online.

User response

Check the status of the nodes by using the **lsnodecanister** command for enclosure-based systems or the **lsnode** command for appliance-based systems. If any node in the specified I/O group is

not currently online, follow maintenance procedures for that node. When all nodes are online, repeat the command.

CMMVC9032E The Add E-mail User operation failed because there is no space left in the user list.

Explanation

The maximum number of email recipients is already configured.

User response

Remove an existing email recipient and retry.

CMMVC9033E The operation failed because there is already a user with that name.

Explanation

An email recipient with that address already exists.

User response

Make sure that you specified the correct user name. If not, make the correction and try again.

CMMVC9034E The operation failed because the specified user does not exist.

Explanation

An attempt was made to perform an operation on a user that does not exist.

User response

Retry the command with an existing user.

CMMVC9035E The Remove E-mail User operation failed because this is the last entry in the user list.

Explanation

Email services require that at least one participant is configured.

User response

Run the **stopemail** command to stop email services and then remove the email user.

CMMVC9036E Sendmail error EX_USAGE. A command or configuration line has been used incorrectly.

Explanation

The send email task failed because a command or a configuration line was used incorrectly.

User response

Ensure that the email settings are correct and resubmit the task.

CMMVC9037E Sendmail error EX_DATAERR. Address is wrong, or the message is too large for the mailbox.

Explanation

The send email task failed because the message sent is too large or a recipient address is incorrect.

User response

Ensure that all addresses are correct and that the message is not too large, and resubmit the task.

CMMVC9038E Sendmail error EX_NOINPUT. An input file (not a system file) did not exist or was not readable.

Explanation

The send email task failed because a file is missing or cannot be read.

User response

Ensure that the email system is configured correctly. Ensure that access permissions are specified correctly for all email configuration files. Resubmit the task.

CMMVC9039E Sendmail error EX_NOUSER. The sendmail command could not recognize a specified user.

Explanation

The send email task failed because the user and domain combination that you specified does not exist.

User response

Specify a defined user and domain combination, and resubmit the task.

CMMVC9040E Sendmail error EX_NOHOST. The sendmail command could not recognize the specified host name.

Explanation

The send email task failed because the host is not known to the email system.

User response

Ensure that you configured the SMTP environment correctly and that you specify a defined host. Resubmit the task.

CMMVC9041E Sendmail error EX_UNAVAILABLE. A required system resource is not available.

Explanation

The send email task failed because a required system resource is not available.

User response

Ensure that you configured the SMTP environment correctly, and then resubmit the task.

CMMVC9042E Sendmail error EX_SOFTWARE. An internal error occurred (including bad arguments).

Explanation

The send email task failed because an incorrect parameter or parameter value was detected.

User response

Ensure that you configured the SMTP environment correctly. Specify only supported parameters and parameter values, and resubmit the task.

CMMVC9043E Sendmail error EX_OSERR. A system resource error prevented the sending of an email.

Explanation

The send email task failed because a system resource error occurred.

User response

Ensure that you configured the SMTP environment correctly and resubmit the task.

CMMVC9044E Sendmail error EX_OSFILE. Failed to open a critical system file.

Explanation

The send email task failed because a required system file cannot be opened.

User response

Ensure that the email system is configured correctly, and that access permissions are specified correctly for all email configuration files. Resubmit the task.

CMMVC9045E Sendmail error EX_CANTCREAT. An output file could not be written to by sendmail.

Explanation

The send email task failed because the system cannot write to a required output file.

User response

Ensure that the email system is configured correctly, and that access permissions are specified correctly for all email configuration files. Resubmit the task.

CMMVC9046E Sendmail error EX_IOERR. A system I/O error occurred during a sendmail operation. This could be due to a disk failure.

Explanation

The send email task failed because a write or read I/O operation failed. This error might be caused by a disk device failure.

User response

Correct the root cause of the I/O failure, and resubmit the task.

CMMVC9047E Sendmail error EX_TEMPFAIL. The sendmail command could not create a connection to a remote system.

Explanation

The send email task failed because the sendmail application cannot establish a connection to the remote system.

User response

Ensure that the network connection to the remote system is functioning correctly, and resubmit the task.

CMMVC9048E Sendmail error EX_PROTOCOL. The remote system returned

something that was incorrect during a protocol exchange.

Explanation

The send email task failed because an error occurred in the protocol exchange.

User response

Ensure that the email system is configured correctly and that you configured the SMTP environment correctly. Resubmit the task.

CMMVC9049E **Sendmail error EX_NOPERM. The user does not have permission to perform the requested operation.**

Explanation

The send email task failed because the User ID does not have authorization to submit the task.

User response

Ensure that authorizations for your User ID in the email and SMTP configurations are correct, and resubmit the task.

CMMVC9050E **Sendmail error EX_CONFIG. There is a fatal problem with the sendmail configuration.**

Explanation

The send email task failed because the sendmail configuration is not correct.

User response

Ensure that the email system is configured correctly and that you configured the SMTP environment correctly. Resubmit the task.

CMMVC9051E **An unknown error occurred. Ensure your SMTP server is running.**

Explanation

The send email task failed because an unexpected error occurred.

User response

Ensure that the SMTP server is running, and resubmit the task.

CMMVC9052E **The email command timed out. Check your email server settings as listed on the cluster.**

Explanation

The send email task failed because a command timeout occurred.

User response

Ensure that your system settings match those that are recommended in the sendmail application documentation, and resubmit the task.

CMMVC9053E **The email service has not been enabled.**

Explanation

The send email task failed because the email application is not enabled.

User response

Enable the email application by using the **startemail** command and resubmit the task.

CMMVC9054E **The user specified does not exist.**

Explanation

You must specify a User ID that exists.

User response

Ensure that the User ID that you specify is defined, and resubmit the task.

CMMVC9055E **The command failed because a target volume has dependent FlashCopy mappings.**

Explanation

The target volume of the FlashCopy mapping, or the target volume of at least one of the FlashCopy mappings in the consistency group, has other FlashCopy mappings that are dependent on the data on the target volume.

User response

Use the **lsvdiskdependentmaps** command and specify the target volume to determine which FlashCopy mappings depend on the target volume. Either wait for these mappings to reach the idle_or_copied state, or stop these mappings. Resubmit the command that produced this error.

CMMVC9056E The create failed because the source and target volumes are members of FlashCopy mappings that have different grain sizes.

Explanation

All FlashCopy mappings that are in a tree of connected mappings must have the same grain size. An attempt was made to create a new FlashCopy mapping that links two existing trees that have different grain sizes.

User response

You have three options.

- Resubmit the command and specify a different source or target volume.
- Delete all of the existing mappings that contain the source volume and resubmit the command.
- Delete all of the existing mappings that contain the target volume and resubmit the command.

CMMVC9061E Cannot disable key server type because key server objects of this type exist.

Explanation

An attempt was made to disable a key server type for which a key server end point exists. The end point indicates that the key server type is actively being used for encryption and therefore cannot be disabled.

User response

Verify that you specified the correct key server type. If so, delete all key server objects of this type by using the **rmkeyserver** command and try again.

CMMVC9062E Cannot disable key server type because it is not currently enabled.

Explanation

An attempt was made to disable a key server type that is not currently enabled. The **-disable** parameter can only be used to disable a key server type that is currently enabled and has no online key server objects.

User response

Retry the command and specify an appropriate key server type, that is, one that is currently enabled and has no online key server objects.

CMMVC9064E This host is the only host in the host cluster. The host cluster will lose all of its shared mappings after removing this host. Use -force flag to continue.

Explanation

An attempt was made to remove the last host in a host cluster that still includes shared mappings. You must use the **-force** flag in these circumstances, which removes all shared mappings for the host cluster.

User response

If you are sure that you want to delete the host, retry the command and use the **-force** flag.

Note: Use of the **-force** flag can result in unintended loss of data.

CMMVC9065E Host cluster does not contain any hosts.

Explanation

An attempt was made to carry out one of the following actions on an empty host cluster:

- Modify the cluster
- Add or remove shared mappings

User response

Make sure that you specified the correct host cluster. If so, add hosts to the host cluster before you modify it, or before you add or remove shared mappings.

CMMVC9066E Volume already has a shared mapping to the host cluster.

Explanation

An attempt was made to map a volume to the same host cluster twice.

User response

Correct the volume name and retry the command.

CMMVC9067E Volume is mapped as a subsidiary LUN (Virtual Volume).

Explanation

An attempt was made to map a VMware vSphere Virtual Volume to a host cluster. This mapping is not permitted.

User response

Retry the command and use a volume that is not a Virtual Volume.

CMMVC9068E **Hosts in the host cluster have conflicting SCSI ID's for their private mappings.**

Explanation

An attempt was made to create a shared host cluster mapping to a volume, but a host within the host cluster is already privately mapped to that volume.

User response

Make sure that the host mappings are compatible with the mappings for the host cluster and the hosts it contains.

CMMVC9069E **Volume does not have a shared mapping to this host cluster.**

Explanation

An attempt was made to remove a mapping from a volume to a host cluster where the mapping did not exist.

User response

Make sure that you specified the correct volume and host cluster. If not, retry the command with the correct parameters. If so, no further action is necessary because the mapping does not exist.

CMMVC9070E **Either -keepmappings or -removemappings flag must be specified.**

Explanation

An attempt was made to remove a host from a host cluster that has one or more shared mappings, but no flag was set to determine the disposition of the mappings for that host.

User response

Retry the command and specify either the **-keepmappings** flag to retain the existing mappings or the **-removemappings** flag to delete them.

CMMVC9071E **Cannot remove a mapping because it is a shared mapping.**

Explanation

An attempt was made to remove a shared mapping from a host cluster with the **rmvdiskhostmap** command, which removes a single host-volume mapping.

User response

Use the **rmvolumehostclustermap** command to remove a shared mapping from a host cluster.

CMMVC9072E **The host cluster mapping was not created because a volume is already mapped to this host cluster with this SCSI LUN.**

Explanation

A SCSI LUN was specified that was already in use for a volume that is mapped to the current host cluster.

User response

Retry the command, specifying a different SCSI LUN.

CMMVC9073E **The host cluster mapping was not created because this volume is mapped to one of the hosts in the host cluster with a different SCSI LUN.**

Explanation

An incompatible SCSI LUN ID was specified for the mapping. Hosts within the host cluster have private mappings to volumes with this SCSI LUN.

User response

Retry the command and specify a compatible SCSI LUN ID, that is, an ID that is not in use for private mappings.

CMMVC9074E **The host cluster mapping was not created because this volume is mapped to a host not in this host cluster. Use -force to create the mapping.**

Explanation

An attempt was made to map a volume to a host cluster when that volume is already mapped to another host or host cluster.

User response

Make sure that you specified the correct volume and host cluster. If so, you can retry the command and use the **-force** flag to create the mapping. If you do so, the previous mapping will also be preserved. Always be cautious when you use the **-force** flag, which can have unforeseen consequences.

CMMVC9075E The host cluster mapping was not created because the limit of mapped volumes supported for a host in this host cluster has been reached.

Explanation

This error can occur if an attempt is made to map a new volume to a host cluster when one or more of the hosts in the host cluster already have the maximum number of mapped volumes. This error can also occur if an attempt is made to change the host cluster type to **adminlun**, but one of the mappings for one of the hosts in the host cluster has a SCSI LUN that exceeds the maximum for the **adminlun** type. The maximum number of mapped volumes for the default host type is 2048. For the **adminlun** host type, the maximum number of volume mappings that are not mappings to VMware vSphere Virtual Volumes is 512. Thus, the maximum SCSI LUN cannot exceed 512.

User response

Reduce the number of SCSI LUNs by using the **rmvdiskhostmap** command to remove at least one host mapping, then retry the command.

CMMVC9076E The host cluster mapping was not created because no free SCSI LUN could be found for this shared mapping.

Explanation

An attempt was made to create a shared mapping, but no free SCSI LUN slot was found that was available for all of the hosts in the host cluster.

User response

Unmap one or more volumes to free up at least one SCSI LUN slot and resubmit the command.

CMMVC9077E The host cluster mapping was not created because of conflicting SCSI LUNs. One of the hosts in the host cluster already has a mapping to this volume, however, another host in the host cluster has a mapping to a different volume but with the same SCSI LUN.

Explanation

An attempt was made to create a shared mapping, but hosts in the host cluster have mappings with the same SCSI LUN but different volumes.

User response

Remap the volume to a compatible SCSI LUN ID before you change to a shared mapping, or remove the private mappings for other hosts to eliminate the conflicts.

CMMVC9078E A volume listed in the -ignoreseedvolume list does not have a mapping to any of the hosts specified in the -seedfromhost list.

Explanation

A volume was specified as "non-seeding," but that volume is not mapped to the seeding host.

User response

Retry the command and specify a volume that is mapped to the seeding host.

CMMVC9079E Cannot delete all the hosts, because at least one of the hosts has mappings. After deleting all the hosts, all of their mappings will be deleted as well. Use -force flag to continue.

Explanation

An attempt was made to remove multiple hosts from a host cluster while at least one of the hosts was still mapped to a volume.

User response

- Delete the mappings for any hosts that you want to remove from a host cluster.
- Alternatively, use the **-force** flag to automatically delete all the mappings for a host when that host is removed.

  Note: Use of the **-force** flag can have unintended consequences.

CMMVC9080E The host cluster I/O group(s) cannot be removed because at least one volume in the I/O group(s) has received I/O within

the defined volume protection period.

Explanation

An attempt was made to remove an I/O group where global volume protection is enabled and I/O occurred within the specified quiet period.

User response

Make sure that you specified the correct I/O group. If so, make sure that no activity that involves the I/O group takes place during the required quiet period before you retry the command.

CMMVC9081E **Duplicate host IDs were entered.**

Explanation

An attempt was made to add or remove multiple hosts from a host cluster where the same host ID was specified twice.

User response

Correct the host IDs, specifying each ID only once, and retry the command.

CMMVC9082E **Too many hosts supplied in the list.**

Explanation

A list of hosts that are to be added to a host cluster was provided, but the number of hosts in the list is too large. A host cluster can include a maximum of 128 hosts.

User response

Retry the command with a shorter list of hosts.

CMMVC9083E **A new cloud snapshot could not be created because the volume is not ready.**

Explanation

A cloud snapshot cannot be created if any of the following conditions apply:

- A cloud snapshot, restore, or delete operation is already in progress on the volume
- An unfixed cloud snapshot error is logged against the volume

A new cloud snapshot can be started only when the volume **backup_status** is **ready**.

User response

Complete the following actions:

1. Fix any outstanding snapshot errors and retry the command.
2. If the error persists, use the management GUI or the **lsvolumebackupprogress** and **lsvolumerestoreprogress** commands to monitor the progress of existing cloud snapshot, delete, and restore operations. Wait for the existing operation to complete, or cancel the operation, and then retry the original command.

CMMVC9084W **Volumes that are enabled for cloud snapshots exist without a license for each enclosure.**

Explanation

Each enclosure that contains a volume that is enabled for cloud snapshots must have a valid transparent cloud tiering license. At least one enclosure is missing a license.

User response

Obtain a valid transparent cloud tiering license for any enclosures that require one.

CMMVC9085E **The command failed because the cloud account is still being initialized.**

Explanation

Initialization of the cloud account must complete before the cloud account can be used.

User response

Wait for initialization of the cloud account to complete, then retry the command.

CMMVC9086E **A new cloud snapshot could not be created because the cloud account configured for the volume is in import mode.**

Explanation

A new cloud snapshot can be created only when the cloud account is in normal mode.

User response

Change the mode of the cloud account or configure the volume to use a different cloud account.

CMMVC9087E **A new cloud snapshot cannot be created because the volume is not enabled for cloud snapshots.**

Explanation

A new cloud snapshot can be created only when the cloud snapshot function is enabled.

User response

Enable the cloud snapshot function for the volume by using the management GUI or the **chvdisk** command.

CMMVC9088E **A new cloud snapshot could not be created because the maximum number of cloud snapshots already exists for the volume.**

Explanation

The volume already has the maximum number of cloud snapshots.

User response

Remove any unwanted cloud snapshots for the volume and retry the command. Use the management GUI or the **lsvolumebackupgeneration** and **rmvolumebackupgeneration** commands to list and remove cloud snapshots.

CMMVC9089E **The command failed because no cloud snapshots for the specified volume exist.**

Explanation

An attempt was made to remove cloud snapshots where a valid volume ID was provided, but no cloud snapshots for the specified volume were found.

User response

Use the management GUI or the **lsvolumebackup** command to list the volumes for which cloud snapshots exist in the cloud. Specify one of these volumes when you retry the command.

CMMVC9090E **The command failed because the specified cloud snapshot does not exist.**

Explanation

An attempt was made to remove or restore a cloud snapshot, but the specified snapshot was not found in the cloud.

User response

Use the management GUI or the **lsvolumebackupgeneration** or **lsvolumebackup** command to list the cloud snapshots that exist in the cloud for a particular volume. Specify one of these snapshots when you retry the command.

CMMVC9091E **The command failed because an existing delete operation is in progress for this volume.**

Explanation

Only one delete operation is allowed at a time for a volume that has cloud snapshots. The delete process is asynchronous and runs in the background. While that process is running for one snapshot, you cannot delete another snapshot of the same volume.

User response

Use the management GUI or the **lsvolumebackupprogress** command to list the progress of current delete operations. When no other delete operations are in progress for the volume, retry the command.

CMMVC9092E **The command failed because the specified cloud snapshot is the most recent complete cloud snapshot for the volume.**

Explanation

An attempt was made to delete the most recent cloud snapshot for a volume. The most recent snapshot must remain available for use in restore operations.

User response

Use the management GUI or the **lsvolumebackup** to choose a different cloud snapshot to delete.

CMMVC9093E **The command failed because a restore operation is in progress from this cloud snapshot.**

Explanation

An attempt was made to delete a cloud snapshot while it was being used to restore a volume.

User response

Wait for the restore operation to complete, or cancel the restore operation. Use the management GUI or the **lsvolumerestoreprogress** command to display the progress of current restore operations.

CMMVC9094E **The command failed because a restore operation is in progress from one of the cloud snapshots for this volume.**

Explanation

An attempt was made to delete all cloud snapshots for a volume while one of the cloud snapshots is being used to restore a volume.

User response

Wait for the restore operation to complete, or cancel the restore operation. Use the management GUI or the **lsvolumerestoreprogress** command to list the progress of current restore operations.

CMMVC9095E **The command failed because a volume with the UID specified with the -fromuid parameter already exists on the local system.**

Explanation

The **-fromuid** parameter is used to restore a cloud snapshot that was taken from a different volume.

User response

Do not use the **-fromuid** parameter when the specified volume already exists on the local system.

CMMVC9096E **The command failed because an existing restore operation is in progress on this volume.**

Explanation

An attempt was made to restore a volume that was already being restored.

User response

Wait for the restore operation to complete. If you specified the wrong volume, you can display a list of all volumes for which a restore operation is in progress by using the management GUI or the **lsvolumerestoreprogress** command. Specify a volume for which no restore operations are in progress when you retry the original command.

CMMVC9097E **The command failed because the specified version of the cloud snapshot is not the latest cloud snapshot for the volume. The -deletelatergenerations parameter must be specified if the latest version of the cloud snapshot is not specified.**

Explanation

If cloud snapshots are enabled on the volume, and the cloud snapshot that is being restored is not the most recent cloud snapshot of the volume, then all later cloud snapshots for the volume must be deleted by specifying the **-deletelatergenerations** parameter.

User response

Complete one of the following actions:

- Retry the command with the **-deletelatergenerations** parameter.
- If you are restoring directly to the production volume, disable cloud snapshots on the volume and retry the command. If cloud snapshots are enabled later, the next cloud snapshot will be a full cloud snapshot.
- If you are committing a restore from a temporary volume, consider converting the temporary volume into an independent volume by using the **-detach** parameter. You can then keep the original volume and volume cloud snapshots unchanged.

CMMVC9098E **The command failed because an existing restore operation is in progress from this volume cloud snapshot.**

Explanation

An attempt was made to restore a volume where the specified cloud snapshot was already in use for a restore operation.

User response

Use the management GUI or the **lsvolumerestoreprogress** command to list the progress of current restore operations.

CMMVC9099E **The command failed because no restore operation is in progress for this volume.**

Explanation

An attempt was made to cancel a restore for a volume for which no restore operation is in progress.

User response

To display a list of all volumes for which a restore operation is in progress, use the management GUI or the **lsvolumerestoreprogress** command. Specify one of these volumes when you retry the original command.

CMMVC9100E **The command failed because there is no uncommitted restore operation for this volume.**

Explanation

An attempt was made to commit or detach a restore for a volume for which no uncommitted restore operation was found.

User response

To display a list of all volumes for which a restore operation is in progress, use the management GUI or the **lsvolumerestoreprogress** command. Specify one of these volumes when you retry the original command.

CMMVC9101E **The command failed because the -fromuid parameter was specified and the specified local volume has cloud snapshot enabled.**

Explanation

An attempt was made to restore a cloud snapshot from a different volume. However, cloud snapshots are enabled on the specified local volume.

When you use the **-fromuid** parameter, cloud snapshots cannot be enabled on the local volume.

User response

Complete one of the following actions:

- Disable cloud snapshots on the local volume and retry the command.
- Restore to a different local volume.

CMMVC9102E **The command failed because the -restoreuid parameter was specified when the existing volume has mappings to hosts.**

Explanation

An attempt was made to restore a cloud snapshot from a different volume, where the UID of the local volume was requested to be set to the UID of the volume cloud snapshot. However, the local volume has mappings to a host object.

When you use the **-restoreuid** parameter, the local volume cannot have any volume to host mappings.

User response

Complete one of the following actions:

- Remove the volume to host mappings for the local volume and retry the command.
- Restore to a different local volume.

CMMVC9103E **A restore operation could not be started because the volume is not ready.**

Explanation

A restore operation cannot be started if any of the following conditions apply:

- A cloud snapshot, restore, or delete operation is already in progress on the volume
- An unfixed cloud snapshot error is logged against the volume

A restore operation can be started only when the volume **restore_status** is **available**.

User response

Complete the following actions:

1. Fix any outstanding snapshot errors and retry the command.
2. If the error persists, use the management GUI or the **lsvolumebackupprogress** and **lsvolumerestoreprogress** commands to monitor the progress of existing cloud snapshot, delete, and restore operations. Wait for the existing operation to complete, or cancel the operation, and then retry the original command.

CMMVC9104E **A cloud snapshot could not be deleted because the volume is not ready.**

Explanation

A cloud snapshot cannot be deleted if any of the following conditions apply:

- A cloud snapshot, restore, or delete operation is already in progress on the volume
- An unfixed cloud snapshot error is logged against the volume

User response

Complete the following actions:

1. Fix any outstanding snapshot errors and retry the command.
2. If the error persists, use the management GUI or the **lsvolumebackupprogress** and **lsvolumerestoreprogress** commands to monitor the progress of existing cloud snapshot, delete, and restore operations. Wait for the existing operation to complete, or cancel the operation, and then retry the original command.

CMMVC9105E The command failed because the local volume is a different size than the specified cloud snapshot.

Explanation

An attempt was made to restore a volume that does not have the same provisioned capacity as the volume cloud snapshot.

User response

Resize the local volume and retry the command, or restore to an alternative local volume.

CMMVC9106E The command failed because the -createtemporaryvolume option can be used only if the volume is enabled for cloud snapshots.

Explanation

An attempt was made to create a temporary volume to hold the restored data, but cloud snapshots are not enabled on the target volume for the restore operation. This usage is not supported.

User response

Complete one of the following actions:

- Run the restore operation directly to the specified volume by omitting the **-createtemporaryvolume** parameter.
- Enable cloud snapshots on the target volume for the restore operation and retry the command.

CMMVC9107E A temporary volume could not be created for the restore process because there are insufficient free volume IDs or volume copy IDs in the system.

Explanation

An attempt was made to create a temporary volume to hold the restored data, which required that a new volume be created in the same caching I/O group as the target volume for the restore operation. The new volume was not created because no free volume IDs or volume copy IDs were available in the system.

User response

Determine whether the system includes unwanted volumes or volume copies. If so, remove unwanted volumes or volume copies by using the **rmvolume** or **rmvolumecopy** command, and then retry the command.

CMMVC9108E The volume cannot be resized because the volume is enabled for cloud snapshots.

Explanation

An attempt was made to change the size of a volume when the cloud snapshot feature was enabled on the volume.

User response

Make sure that you specified the correct volume. If so, disable the cloud snapshot feature on the volume by using the **chvdisk** command and retry the original command.

CMMVC9109E A temporary volume could not be created for the restore process because there is insufficient free capacity in the storage pool.

Explanation

An attempt was made to create a temporary volume to hold the restored data, which required that a new volume be created in the same caching I/O group as the target volume for the restore operation. The new volume was not created because no free volume IDs were available in the storage pool.

User response

Determine whether the system includes unwanted local volumes. If so, remove unwanted local volumes by using the **rmvdisk** command, and then retry the command.

CMMVC9110E The restore operation could not be started because the maximum number of FlashCopy mappings in the system has been reached.

Explanation

The restore operation requires an internal FlashCopy mapping to be created and the system limit has been reached.

User response

Remove unwanted FlashCopy mappings from the system by using the **rmfcmap** command and retry the **restorevolume** command.

CMMVC9111E **The command failed because the specified volume is part of a Metro Mirror or Global Mirror relationship.**

Explanation

An attempt was made to restore to a volume that is part of a remote copy relationship.

User response

Make sure that you specified the correct volume. If so, remove the remote copy relationship that uses this volume by using the **rmrcrelationship** command and retry the **restorevolume**command.

CMMVC9112E **The command failed because the specified volume is a HyperSwap volume.**

Explanation

An attempt was made to restore to a HyperSwap volume. This operation is not allowed.

User response

Make sure that you specified the correct volume. If so, convert the volume to a basic volume by removing a copy at one site, then retry the command.

CMMVC9113E **Cloud snapshots are already enabled for the specified volume.**

Explanation

An attempt was made to enable cloud snapshots on a volume where cloud snapshots are already enabled.

User response

Make sure that you specified the correct volume. If you want to enable cloud snapshots that use a different cloud account, first disable cloud snapshots for the current account, then retry the command.

CMMVC9114E **Cloud snapshots cannot be disabled because they are not enabled on this volume.**

Explanation

An attempt was made to disable the cloud snapshot function when it was not enabled.

User response

Make sure that you specified the correct volume.

CMMVC9115E **Cloud snapshots cannot be enabled because the maximum number of cloud snapshot-enabled volumes in the system has been reached.**

Explanation

The system limit was reached for the number of volumes that can be enabled for cloud snapshots.

User response

You must disable cloud snapshots for another volume before you can enable cloud snapshots for the current volume.

CMMVC9116E **Cloud snapshots cannot be enabled because the specified cloud account is in import mode.**

Explanation

The cloud account must be in normal mode when you enable cloud snapshots.

User response

Change the account to normal mode and retry the command.

CMMVC9117E **The volume was not deleted because a cloud snapshot operation is in progress.**

Explanation

The volume was not deleted because it would prevent a cloud snapshot operation from completing.

User response

Make sure that you specified the correct volume. If so, wait for the cloud snapshot operation to complete, or cancel the cloud snapshot. Use the management GUI or the **lsvolumebackupprogress** command to monitor the progress of the cloud snapshot, or the **rmvolumebackupgeneration** command to cancel the cloud snapshot. Alternatively, specify the **-cancelbackup** parameter to force-delete

the volume, canceling any active cloud snapshot operations on the volume.

CMMVC9118E **The volume was not deleted because a restore operation is in progress and it would allow access to an image mode copy that does not contain consistent data.**

Explanation

The volume was not deleted because an image mode copy might contain inconsistent data.

User response

Make sure that you specified the correct volume. If so, wait for the restore operation to complete, or cancel the restore. Use the management GUI or the **lsvolumerestoreprogress** command to monitor the progress of the restore, or the **restorevolume** command to cancel the restore. Alternatively, specify the **-discardimage** parameter to force-delete the volume, canceling any active restore operation.

CMMVC9119E **The volume copy was not deleted because a restore operation is in progress and it might allow access to an image mode copy that does not contain consistent data.**

Explanation

The volume copy was not deleted because an image mode copy might contain inconsistent data.

User response

Make sure that you specified the correct volume. If so, wait for the restore operation to complete, or cancel the restore. Use the management GUI or the **lsvolumerestoreprogress** command to monitor the progress of the restore operation, or the **restorevolume** command to cancel the restore operation. Alternatively, specify the **-discardimage** parameter to force-delete the volume copy.

CMMVC9120E **The remote copy relationship was not created because the master or auxiliary volume is enabled for cloud snapshots.**

Explanation

A volume where cloud snapshots are enabled cannot be part of a remote copy relationship.

User response

Make sure that you specified the correct volume. If so, disable the cloud snapshot feature by using the **rmcloudaccount** command and retry the original command.

CMMVC9121E **The change volume could not be associated because the volume is enabled for cloud snapshots.**

Explanation

A volume cannot be configured as a change volume for a remote copy relationship if the cloud snapshot feature is enabled on the volume.

User response

Make sure that you specified the correct volume. If so, disable the cloud snapshot feature on the volume by using the **rmcloudaccount** command and retry the original command.

CMMVC9122E **The FlashCopy mapping was not created because the source or target volume is enabled for cloud snapshots.**

Explanation

A volume cannot be the source or target volume in a FlashCopy mapping if the cloud snapshot feature is enabled on the volume.

User response

Make sure that you specified the correct volume. If so, disable the cloud snapshot feature on the volume by using the **chvdisk** command and retry the original command.

CMMVC9123E **The command failed because a key server type has not been enabled.**

Explanation

An attempt was made to create a key server without enabling a key server type.

User response

Enable a key server type by using the **chkeyserverisklm** command. Then, retry the **mkkeyserver** command.

CMMVC9124E **The command failed because a primary key server already exists.**

Explanation

An attempt was made to create a key server object by using the **mkkeyserver -primary** command where the primary key server already exists. After the primary key server is created, subsequent key server objects cannot be created with the **-primary** flag. The primary key server can be changed later by using the **chkeyserver** command.

User response

Retry the command without specifying the **-primary** flag.

CMMVC9125E **The command failed because the -primary flag must be set when creating the first ISKLM key server.**

Explanation

An attempt was made to create the primary ISKLM key server object without specifying the **-primary** flag. The first key server that you create of the ISKLM type must have the **-primary** flag set. In other words, you must create your primary key server first. Subsequent key server objects cannot be created with the **-primary** flag. The primary ISKLM key server can be changed later by using the **chkeyserver** command.

User response

Specify the **-primary** flag when you retry the command.

CMMVC9126E **The command failed because the -primary flag can only be set for ISKLM key servers.**

Explanation

An attempt was made to create a key server object where the **-primary** flag was specified. The **-primary** flag is valid only for ISKLM key servers, but the ISKLM key server type is not currently enabled.

User response

Retry the command without specifying the **-primary** flag.

CMMVC9127E **The command failed because an SSL certificate is required.**

Explanation

An attempt was made to create a key server object without specifying a self-signed SSL certificate. The parent key server type object does not have a CA certificate configured either. As a result, no SSL certificate would be available on the system for communicating with the key server.

User response

Configure a CA certificate for this key server type or supply a self-signed certificate for the key server with the **-sslcert** parameter.

CMMVC9128E **Cannot enable key server type because it would exceed the permitted number of enabled key server types.**

Explanation

One key server type is typically enabled at a time. Migration between key server types is a special case where a second type can be enabled until the migration is complete. This error means that an attempt was made to enable a third key server type during a migration, which is not permitted.

User response

If you are currently migrating between two key server types, try the command again and this time include the **-disable** parameter to disable the migration target.

CMMVC9129E **The command has failed because an IPv4 address was specified and each node does not have an IPv4 service IP address.**

Explanation

An attempt was made to create the key server object with an IPv4 address, which requires each node in the system to have an IPv4 service IP address set.

User response

Use the **lsservicestatus** command to ensure that each node in the system has an IPv4 service IP address and then retry the command. Alternatively, if each node has an IPv6 service address, specify an IPv6 service IP address when you retry the command.

CMMVC9130E **The command has failed because an IPv6 address was specified and each node does not have an IPv6 service IP address.**

Explanation

An attempt was made to create the key server object with an IPv6 address, which requires each node in the system to have an IPv6 service IP address set.

User response

Use the **lsservicestatus** command to ensure that each node in the system has an IPv6 service IP address and then retry the command. Alternatively, if each node has an IPv4 service address, specify an IPv4 service IP address when you retry the command.

CMMVC9131E The command failed because the key server reported an error.

Explanation

An attempt was made to create or test a key server on the system. During both the **mkkeyserver** and **testkeyserver** tasks, the system attempts to validate the key server. During this validation, the key server reported an error. More detailed information about the error might be found in the additional sense data in the event log.

User response

Check the event log for key server errors. Fix any key server errors and run the task again.

CMMVC9132E This will change the site of at least one host in the host cluster that already has a site defined. Use -force flag to continue.

Explanation

An attempt was made to change the site of the host cluster, which would change the site of at least one host in the host cluster.

User response

Make sure you specified the correct information in the command. If so, you can use the **-force** flag to forcibly change the site of one or more hosts in the host cluster. Using the **-force** flag can have unexpected results.

CMMVC9134E The command failed because the key server is not supported.

Explanation

An attempt was made to create or test a key server on the system. During both the **mkkeyserver** and **testkeyserver** tasks, the system attempts to validate the key server. During this validation, the key server reported unsupported vendor information. More detailed information about the server might be found in the additional sense data in the event log.

User response

Check the event log for key server errors. Fix any key server errors and run the task again.

CMMVC9135E The command failed because there was a problem establishing a connection to the key server.

Explanation

An attempt was made to create or test a key server on the system. During both the **mkkeyserver** and **testkeyserver** tasks, the system attempts to validate the key server. An error occurred in establishing a connection to the key server by using the supplied IP address, IP port, and SSL certificate for the key server. This error might be caused by a network problem, incorrect IP address or port details, or a problem with the SSL certificate. More detailed information about the error might be found in the additional sense data in the event log.

User response

Confirm that the correct IP address, IP port, and SSL certificate were supplied for the key server. Confirm that each node in the system has access to the key server. Confirm that the key server is fully operational and run the task again.

CMMVC9136E The command failed because the key server's response was not understood.

Explanation

An attempt was made to create or test a key server on the system. During both the **mkkeyserver** and **testkeyserver** tasks, the system attempts to validate the key server. The system was unable to process the response from the key server.

User response

Fix any key server errors and run the task again.

CMMVC9137E The command failed because no primary key server exists.

Explanation

An attempt was made to create new key server master keys without a primary key server object. The key

server object that is marked as primary is responsible for creating new keys. A primary key server object must exist before key server master keys can be created.

User response

Use the **mkkeyserver** command to designate one key server as the primary key server and run the **chencryption** command again.

CMMVC9138E **The command failed because no key servers exist.**

Explanation

An attempt was made to create new key server master keys when no key server was known to the system. At least one key server object must exist before key server master keys can be created.

User response

Use the **mkkeyserver** command to create a key server object and run the **chencryption** command again.

CMMVC9139E **The command failed because not all key servers are online.**

Explanation

An attempt was made to create new key server master keys when the status of a key server object shows that it is not online. All key server objects must be online before key server master keys can be created.

User response

Check the status of all key servers and fix any issues. Run the command again when the status of all key servers is online.

CMMVC9140E **The key server could not be deleted because it is the only key server of the enabled key server type.**

Explanation

An attempt was made to delete the last remaining key server object for a key server type with a status of **enabled_active**. Deleting this key server would result in loss of access to encryption keys and might cause encrypted objects to go offline.

User response

Use the **chencryption** command to disable key server encryption.

CMMVC9141E **The validate command is not valid for key servers. Use the testkeyserver command to validate a key server.**

Explanation

An attempt was made to use the **chencryption -keyserver validate** command for key servers, which is not permitted. Only a single key server can be tested at once.

User response

Use the **testkeyserver** command to test individual key server objects.

CMMVC9142E **The command failed because another encryption function is already enabled.**

Explanation

An attempt was made to enable a second encryption function when another function is already enabled on the system. For example, key server encryption might have been requested when USB encryption was already enabled.

User response

To enable the new encryption function, first use the **chencryption** command to disable the function that is currently enabled.

CMMVC9143E **The command failed because the key server reported that the primary has been misconfigured on the system.**

Explanation

An SKLM key server reported a server type that conflicted with the value defined on the system. The key server reported it is not the primary, but the server is defined to be the primary on the system.

User response

Make sure that the correct key server is designated as primary. For more information, see the documentation for error code 1785.

CMMVC9144E The command failed because the system cannot retrieve the current key from the primary key server.

Explanation

An attempt was made to complete one of the following actions on a system that is configured to use key server encryption:

- Disable the USB flash drives encryption function that is enabled along with key server encryption.
- Prepare a USB flash drives encryption key for the first time.
- Prepare an encryption recovery key encryption key for the first time.

To safely allow USB flash drives encryption to be disabled, or to allow USB flash drives or encryption recovery keys to be prepared for the first time, the system checks that it has access to the current encryption key on the primary key server. This error message is displayed because the system was not able to fetch the key. Possible reasons for the failure include the following situations:

- A network issue prevented connection with the key server.
- The encryption key does not exist on the key server.
- No primary key server is configured on the system.

User response

Fix any key server problems in the event log to bring the primary key server online, then retry the command. If the problem persists, verify that the current key exists on the primary key server.

CMMVC9145E Unable to perform LBA lookup due to unreadable metadata.

Explanation

A reverse lookup failed due to media errors.

User response

Use your host application to validate the data on the volume. After you have located the hard disk errors, restore the missing data from backup.

If the volume is mirrored, and the affected volume is one of the copies, you can use the **repairvdiskcopy** command with the **-validate** flag to compare the faulty copy to the good one. Use the **-resync** option to restore the data from the good copy to the faulty one.

CMMVC9146E Cannot add or remove support center objects when support assistance is enabled.

Explanation

Support center configurations cannot be modified while secure remote access is enabled.

User response

Disable support assistance by using the **chsra -disable** command and try the original command again.

CMMVC9148E The command has failed because the default support center cannot be deleted.

Explanation

Some support center objects are added when you first install the software or update to the current version. These objects cannot be deleted.

User response

Specify a non-default support center object for deletion.

CMMVC9157E The command has failed because the specified secondary expansion module is offline.

Explanation

The specified secondary expansion module is offline.

User response

This error starts a directed maintenance procedure (DMP) that gives instructions for reseating or replacing the module. If the DMP is not displayed, contact IBM Support.

CMMVC9158E The command has failed because the specified display panel is offline.

Explanation

The specified display panel is offline.

User response

This error starts a directed maintenance procedure (DMP) that gives instructions for replacing the display panel. If the DMP is not displayed, contact your support representative.

CMMVC9159E Host is invalid or does not exist

Explanation

The ID or name of the host that is to be throttled was not valid or was not found.

User response

Use the **lshost** command to create a list of valid hosts. Retry the command with a valid host.

CMMVC9160E Throttle is already associated with this host.

Explanation

The host that was specified in the **mkthrottle** command already has an associated throttle.

User response

Make sure that you specified the correct host. If necessary, use the **lshost** command to create a list of valid hosts and retry the command with a valid host. If you want to change the throttle parameters for this host, use the **chthrottle** command.

CMMVC9161E Host already has an associated host cluster throttle.

Explanation

A host cannot have a throttle if its parent host cluster defines a host cluster throttle.

User response

Make sure that you specified the correct host. If necessary, use the **lshost** command to create a list of valid hosts and retry the command with a valid host. You can tune the system by changing the host cluster throttle to accommodate the requirements of the member hosts by using the **chthrottle** command.

CMMVC9162E Hostcluster is invalid or does not exist.

Explanation

The ID or name of the host cluster that is to be throttled was not valid or was not found.

User response

Use the **lshostcluster** command to create a list of valid host clusters, and then retry the command with a valid host cluster.

CMMVC9163E Throttle is already associated with this hostcluster.

Explanation

The host cluster that was specified in the **mkthrottle** command already has an associated throttle.

User response

Make sure that you specified the correct host cluster. If necessary, use the **lshostcluster** command to create a list of valid host clusters and retry the command with a valid host cluster. If you want to change the throttle parameters for this host cluster, use the **chthrottle** command.

CMMVC9164E Host cluster has a member host which has a throttle defined for it.

Explanation

An attempt was made to define a throttle for a host cluster where throttles are already defined for one or more member hosts. This action is not permitted.

User response

Make sure that you specified the correct host cluster. If so, use the **lsthrottle** command to review the hosts that have throttles. If the information shown is correct and you do not need to define a throttle for the host cluster, no further action is required. If you want to define a throttle for the host cluster, you must remove the individual host throttles by using the **rmthrottle** command. When no throttles are defined for any of the member hosts, retry the command.

CMMVC9165E Mdiskgroup is invalid or does not exist.

Explanation

The ID or name of the storage pool that is to be throttled was not valid or was not found.

User response

Use the **lsmdiskgrp** command to create a list of valid storage pools. Retry the command with a valid storage pool.

CMMVC9166E Throttle is already associated with this mdiskgroup.

Explanation

The storage pool that was specified in the **mkthrottle** command already has an associated throttle.

User response

Make sure that you specified the correct storage pool. If necessary, use the **lsmdiskgrp** command to create a list of valid storage pools and retry the command with a valid storage pool. If you want to change the throttle parameters for this storage pool, use the **chthrottle** command.

CMMVC9167E **Host and hostcluster already have associated throttles.**

Explanation

An attempt was made to add a host to a host cluster when throttles were already defined for both the host and the host cluster. Throttles can be defined for the host or for the host cluster, but not for both.

User response

Use the **rmthrottle** command to remove either the host throttle or the host cluster throttle, then retry the **addhostclustermember** command.

CMMVC9168E **Seeding host has a throttle associated with it.**

Explanation

An attempt was made to create a host cluster when a throttle was defined for one or more of the seeding hosts.

User response

Use the **rmthrottle** command to remove throttles from any seeding hosts, then retry the **mkhostcluster** command.

CMMVC9173E **The pool specified is a data reduction pool. Thin provisioned or compressed volume copies created from a data reduction pool must use -autoexpand.**

Explanation

When you create a volume copy from a data reduction pool, you must enable the autoexpand capability. The command fails if you try to create a thin-provisioned or compressed volume copy from a data reduction pool and do not specify the **-autoexpand** parameter.

User response

Use the **-autoexpand** parameter when you create a thin-provisioned or compressed volume copy from a data reduction pool.

CMMVC9175E **The pool specified is a data reduction pool. Volumes or volume copies which are thin provisioned or compressed and created from a data reduction pool can not be striped specifying an mdisk, sequential or image mode.**

Explanation

When you create a thin-provisioned or compressed volume or volume copy from a data reduction pool, do not specify striped with **mdisk**, **sequential**, or **image** mode.

User response

Retry the command with valid options. You can use striped mode if you do not specify an MDisk.

CMMVC9176E **The pool specified is a data reduction pool. For volumes or volume copies which are thin provisioned or compressed and created from a data reduction pool, the volume cannot have a cache mode of none or readonly.**

Explanation

A thin-provisioned or compressed volume or volume copy that you create from a data reduction pool must have a cache mode of **readwrite**.

User response

Change the cache mode on the volume to **readwrite** and retry the command.

CMMVC9177E **The pool specified is a data reduction pool. Thin provisioned or compressed volume copies created from a data reduction pool can not use -noautoexpand.**

Explanation

A thin-provisioned or compressed volume copy created from a data reduction pool must have the autoexpand option enabled.

User response

Retry the command without the **-noautoexpand** parameter.

CMMVC9179E **The command failed as the resource required for a data reduction pool could not be allocated.**

Explanation

Insufficient system resources are available to create a data reduction pool.

User response

Delete enough volume copies, volumes, or pools to free the needed system resources, then retry the command.

CMMVC9180E **The data reduction pool can not be created as the maximum number of data reduction pools already exists.**

Explanation

A maximum of four data reduction pools can be created per system.

User response

Use the **rmmdiskgrp** command to delete an existing data reduction pool and then retry the original command.

CMMVC9181E **Data reduction can only be set to on for a parent pool.**

Explanation

A child pool cannot be designated as a data reduction pool.

User response

Use the **chmdiskgrp** command to change the parent pool to use data reduction and then use the **chmdiskgrp** command to change the child pool to an inherited data reduction pool.

CMMVC9184E **The volume specified is a thin or compressed volume in a data reduction pool. Autoexpand must be on and can not be changed.**

Explanation

Thin-provisioned or compressed volumes cannot have an autoexpand option that is turned off.

User response

Make sure the **-autoexpand** parameter is set to on and then retry the command.

CMMVC9185E **The volume or volume copy specified is a thin or compressed volume in a data reduction pool and the action requested is not supported on volumes of this type.**

Explanation

The command cannot be run if either copy is a thin-provisioned or compressed volume in a data reduction pool.

User response

Retry the command, specifying volumes that meet at least one of the following requirements:

- The volumes are fully allocated.
- The volumes are in a regular pool.

CMMVC9186E **The pool specified is a data reduction pool. It is not possible to migrate thin or compressed volumes to a data reduction pool.**

Explanation

You cannot migrate thin-provisioned or compressed volumes to a data reduction pool by using the **migratevdisk** command.

User response

Complete the following steps to migrate a thin-provisioned or compressed volume to a data reduction pool:

1. Create a copy of the thin-provisioned or compressed volume in the data reduction pool by using the **addvdiskcopy** command.
2. Delete the original version of the volume copy by using the **rmvdiskcopy** command.

CMMVC9190E **The pool specified is a data reduction pool, creating thin or compressed volumes from this type of pool must have autoexpand on.**

Explanation

You must enable the autoexpand option for all thin-provisioned or compressed volumes in data reduction pools.

User response

For the **mkvdisk** command, make sure that the **-autoexpand** parameter is set to on.

For the **mkvolume** command, make sure that the **-noautoexpand** parameter is *not* included.

Retry the command.

CMMVC9191E **The command failed because the volume is enabled for cloud snapshots.**

Explanation

An attempt was made to migrate a volume between storage pools or to add a copy of a volume to different storage pool when the cloud snapshot feature was enabled on the volume. These actions are not permitted.

User response

Make sure that you specified the correct volume. If so, disable the cloud snapshot feature by using the **chvdisk** command and retry the original command.

CMMVC9193E **The command failed because the system requested for import does not have data in the account.**

Explanation

An attempt was made to configure a cloud account to import data from another system, but the other system does not have data in its account.

User response

Check the list of systems with data in the account by using the **lscloudaccountimportcandidate** command and retry the original command.

CMMVC9194E **The command failed because at least one volume is using the account.**

Explanation

An attempt was made to change the mode of a cloud account when at least one of the volumes in the system is configured to use the account. For example, a volume might be configured to use cloud snapshots. The account mode cannot be changed until these volumes are reconfigured to not use the account.

User response

Reconfigure the volumes that are using the account and retry the command.

CMMVC9195E **The account cannot be deleted because it is being used by volumes.**

Explanation

An attempt was made to delete a cloud account, but volumes on the system are configured to use that account. For example, a volume might be configured to use the cloud snapshot function.

User response

Make sure that you specified the correct account. If so, verify that you want to disconnect the volumes that use the account. If so, disconnect the volumes and retry the command.

CMMVC9196E **All available quorum disks are dependent on the sem that you have specified.**

Explanation

All quorum disks depend on the specified secondary expander module (SEM). You must configure the system so that at least one of the drives that are allocated to hold quorum remains online when the canister goes offline.

User response

Assign one or more of the drives in the control enclosure as a quorum drive by using the **chquorum** command. After you configure the quorum drives, test for dependencies by using the **lsdependentvdisks** command with the **-sem** option.

CMMVC9197E **Cloud snapshot and restore operations are not allowed because of a generation delete abandon. Use rmvolumebackupgeneration to retry the delete operation.**

Explanation

A previous generation delete operation was abandoned. You must complete this operation before further snapshot or restore operations can be run.

User response

Rerun the abandoned generation delete operation for the volume. When the operation is complete, new snapshot and restore operations are possible. Use the management GUI or the **lsvolumebackupgeneration** and **rmvolumebackupgeneration** commands to list and remove snapshot generations.

CMMVC9198E **A previous rmvolumebackupgeneration -all command was abandoned. Use rmvolumebackupgeneration -all to retry the delete operation.**

Explanation

A previous generation delete operation that used the **-all** option was abandoned. The operation must complete before further snapshots or restores are allowed.

User response

Rerun the abandoned generation delete operation for the volume and include the **-all** option. When the operation is complete, new snapshot and restore operations are possible. Use the management GUI or the **rmvolumebackupgeneration -all** command to remove all snapshot generations for a volume.

CMMVC9199E **A previous rmvolumebackupgeneration or cloud restore with -deletelatergenerations command was abandoned. Use rmvolumebackupgeneration specifying the lowest deleting generation to retry the delete operation.**

Explanation

One of the following operations was abandoned:

- A generation restore operation with the **-deletelatergenerations** option
- A **rmvolumebackupgeneration** operation with the **-generation** option

The operation must complete before further snapshots or restores are allowed.

User response

Run the **lsvolumebackupgeneration** command and look for one or more generations that are labeled **deleting**. Run the **rmvolumebackupgeneration** operation and specify those generations. When the operation is complete, new snapshot and restore operations are possible.

CMMVC9202E **The task failed because the volume is part of a Metro Mirror or Global Mirror relationship to a system which is running software which does not support resize of volumes in a relationship.**

Explanation

You can change the size of the volumes that are part of Metro Mirror or Global Mirror relationships only if both systems involved in that relationship are running software that supports the resizing function.

User response

Either upgrade the remote system or remove the Metro Mirror or Global Mirror relationship, then retry the command.

CMMVC9203E **The expandvdisksize task failed because there is not enough memory available for this feature.**

Explanation

The volume is part of Metro Mirror or Global Mirror relationship and requires more memory to accommodate the change recording map.

User response

Increase the memory space available for Remote Copy in the caching I/O group of the volume you want to expand, then retry the task.

CMMVC9204E **The task failed because the volume is part of a relationship that is configured to operate in Global Mirror in cycling mode.**

Explanation

An attempt was made to change the size of a volume that is part of a Global Mirror relationship while that relationship is configured to operate in multiple cycling mode. This change is not permitted.

User response

Complete one of the following procedures:

- Delete the Global Mirror relationship.
- Remove the relationship from any consistency group, convert the single relationship to Global Mirror non-cycling mode, start the relationship, and wait until

the relationship achieves a consistent_synchronized state.

Retry the volume size change.

CMMVC9205E The task failed because the volume is part of a HyperSwap relationship.

Explanation

Volumes cannot be resized if they are part of a HyperSwap relationship.

User response

Make sure that you specified the correct volume. If so, convert the volume to a basic volume by removing a copy at one site, then retry the command.

CMMVC9206E The task failed because the volume is part of a relationship that is not consistent_synchronized.

Explanation

An attempt was made to resize a volume that is part of a remote copy relationship that is not synchronized. A volume that is part of a Metro Mirror or Global Mirror relationship can be resized only when the relationship is synchronized.

User response

Start the relationship and wait for it to synchronize, and then retry the resize command.

CMMVC9207E The task failed because the volume being expanded is a primary in a Metro Mirror or Global Mirror relationship and would make the primary volume size different than the secondary volume size.

Explanation

The volume that is being expanded is a primary in a Metro Mirror or Global Mirror relationship, and the request to expand is making the volume either larger or smaller than the associated secondary. Volumes in relationships can be expanded by first expanding the secondary by the required amount, and then expanding the primary volume to the same size.

User response

Make sure that the secondary is expanded first, and then adjust the requested size so that the primary volume becomes the same size as the associated secondary volume.

CMMVC9208E The task failed because the volume being expanded is a secondary in a Metro Mirror or Global Mirror relationship and has already been expanded so is larger than the associated primary volume.

Explanation

An attempt was made to expand a secondary volume that was already larger than its associated primary. You cannot expand the secondary volume unless it is the same size as the primary volume.

User response

Either expand the primary volume to the same size as the secondary, or shrink the secondary volume to the same size as the primary. After both volumes are the same size, you can retry the expansion of the secondary volume.

CMMVC9209E The task failed because the volume being shrunk is a secondary in a Metro Mirror or Global Mirror relationship and would make the secondary volume size different than the associated primary volume.

Explanation

The volume that is being shrunk is a secondary in a Metro Mirror or Global Mirror relationship, and the request to shrink would make the volume either larger or smaller than the associated primary volume. Volumes in relationships can be shrunk by first shrinking the primary by the required amount, and then shrinking the secondary volume to the same size.

User response

If you want to shrink the secondary volume, the secondary volume must be larger than the primary volume, and you can shrink it to the size of the primary volume only. When both volumes are the same size, you can shrink the primary volume first, then shrink the secondary volume to the same size.

CMMVC9210E The task failed because the volume being shrunk is a primary in a Metro Mirror or Global Mirror relationship and has already been shrunk so is already smaller than the associated secondary volume.

Explanation

The volume that is being shrunk is a primary volume that is already smaller than the associated secondary. Volumes in relationships can be shrunk by first shrinking the primary by the required amount, and then shrinking the secondary volume to the same size.

User response

Either shrink the secondary volume to be the same size as the primary volume, or expand the primary volume to be same size as the secondary volume. When both volumes are the same size, you can retry shrinking the primary volume.

CMMVC9211E **Cannot perform task because primary is different size to the secondary.**

Explanation

The specified command can be completed only on a relationship where the primary and secondary volumes are the same size.

User response

Make sure that you specified the correct relationship. If so, either expand the primary to the same size as the secondary, or shrink the secondary to the same size as the primary. You can then retry the command.

CMMVC9212E **Cannot perform task because one or more of the primaries is a different size to the corresponding secondary**

Explanation

The specified command can be completed only on a consistency group where all of the relationships consist of primary and secondary volumes that are the same size.

User response

Ensure that all primary volumes are the same size as their corresponding secondary volumes, then retry the command.

CMMVC9215E **The action failed as the tier cannot be changed for MDisks from this controller.**

Explanation

The tier is fixed for MDisks from this controller.

User response

Make sure that you specified the correct MDisk. If not, specify an MDisk from a controller that allows the tier to be changed.

CMMVC9218E **The action failed because remote support is either in connected or active state.**

Explanation

An attempt was made to test remote support assistance when it was in a connected or active state.

User response

Test only when remote support assistance is in a disconnected, connecting, or failure state.

CMMVC9219E **The command failed because the remote support feature has already been enabled.**

Explanation

The **chsra -remotesupport enable** command was entered when remote support assistance was already enabled.

User response

If all you wanted was to enable remote support assistance, no further action is required.

If you ran the **chsra** command to modify the **-idletimeout** parameter, complete the following steps:

1. Disable remote support by running the **chsra -remotesupport disable** command.
2. Enable remote support with an idle timeout, as in the following example:

   ```
   chsra -remotesupport enable -idletimeout 60
   ```

   This example sets the idle timeout to 60 minutes.

If your support representative advised you to add a new support center, or if you want to configure a proxy server, complete the following steps:

1. Disable support assistance by using the **chsra -disable** command.
2. Configure the new support center or proxy server by using the **mksystemsupportcenter** command.
3. Enable local and remote support assistance as required by using the **chsra -enable** and **chsra -remotesupport enable** commands.

CMMVC9220E **The command failed because the name of the support center started with the prefix associated with default support centers.**

Explanation

An attempt was made to create a support center whose name began with "default_support_center". This prefix is reserved for use by default support centers.

User response

Try the command again and use a name for the support center that does not begin with "default_support_center".

CMMVC9224E **The command failed because the specified node is not a part of the site or I/O group specified, or both.**

Explanation

An attempt was made to change the authentication information for a node that is not a member of the specified site or I/O group.

User response

Complete one of the following actions:

- Retry the command and specify a node that is a member of the specified site or I/O group.
- Use the management GUI or the **addnode** or **chnode** command to add the node to the appropriate site and I/O group, then retry the **chiscsistorageport** command.

Note: The **site_id** or **site_name** parameters can be specified for stretched or HyperSwap topologies only.

CMMVC9225E **The command failed because the maximum limit for the initiator node-specific authentication credentials was exceeded.**

Explanation

When you use the initiator node-specific credentials mode, a maximum of 32 characters is allowed for the user name and 32 characters for the chapsecret. This mode is triggered by the **-node** parameter.

User response

Retry the command and specify a user name and chapsecret that each have 32 or fewer characters.

CMMVC9226E **The command failed because the specified iSCSI storage port is configured for single username/ password credentials across all nodes. Per SVC node credentials cannot be specified.**

Explanation

An attempt was made to change the mode of authentication from single credentials system wide to per-node credentials. This iSCSI storage port is configured to use a single user name and password across all nodes.

User response

Retry the command and omit the **-node** parameter.

CMMVC9227E **The command failed because the specified iSCSI session does not exist.**

Explanation

An attempt was made to change the authentication for an iSCSI session that does not exist.

User response

Retry the command and specify a valid iSCSI session.

CMMVC9228E **The command failed because some parameters are missing or invalid parameters were entered.**

Explanation

The command contained one or more invalid parameters, or failed to include one or more required parameters, or both.

User response

Refer to the product documentation for the command syntax and retry the command with the correct parameters.

CMMVC9230E **The command failed because the system cannot retrieve the current key from a USB flash drive.**

Explanation

An attempt was made to complete one of the following actions:

- Disable the key server encryption function while the system is configured to also use USB encryption.

- Prepare a new key server key when key server encryption is being enabled for the first time.

To safely allow key server encryption to be disabled, or to allow key server keys to be prepared for the first time, the system checks that it has access to the current encryption key on at least one of the USB flash drives. This error message is displayed because the system was not able to read the key. Possible reasons for the failure include the following situations:

- The USB flash drive that is installed is defective.
- The USB flash drive belongs to a different system.

User response

Verify that at least one working USB flash drive that contains the correct key for the system is installed in the system, then retry the command.

CMMVC9231E **The command failed because a USB rekey operation is in progress.**

Explanation

USB encryption is enabled on the system and is currently in the prepared state. An action was attempted that is not permitted in the current state.

User response

Complete the rekey operation by entering one of the following commands:

- `chencryption -usb newkey -key commit`
- `chencryption -usb newkey -key cancel`

Once the rekey operation is complete, retry the original command.

CMMVC9232E **The command failed because a key server rekey operation is in progress.**

Explanation

Key server encryption is enabled on the system and is currently in the prepared state. An action was attempted that is not permitted in the current state.

User response

Complete the rekey operation by entering one of the following commands:

- `chencryption -keyserver newkey -key commit`
- `chencryption -keyserver newkey -key cancel`

Once the rekey operation is complete, retry the original command.

CMMVC9233E **The command failed because a USB rekey operation is required.**

Explanation

USB encryption is enabled on the system. An action was attempted that is not permitted until a new USB key has been created.

User response

Create a new USB key for the system by completing the instructions in the topic "Rekeying an encryption-enabled system using a USB flash drive."

CMMVC9234E **The command failed because the cloud account is in import mode.**

Explanation

An attempt was made to enable or disable a key provider when the cloud account was in import mode. In import mode, the account is read-only and no changes can be made.

User response

Use the change cloud account command for the appropriate provider (such as **chcloudaccountswift**) to change the account mode to **normal**, then retry the original command.

CMMVC9235E **The pool specified is a data reduction pool. Volumes or volume copies which are thin provisioned and created from a data reduction pool can not use the -grainsize parameter.**

Explanation

You cannot create thin-provisioned volumes or volume copies that use the **-grainsize** parameter. This type of volume or volume copy is created with a mandatory grain size of 8 KB.

User response

Retry the command without the **-grainsize** parameter.

CMMVC9236E **The pool specified is a data reduction pool. Volumes or volume copies which are thin provisioned and created from a data reduction**

pool can not use the -warning parameter.

Explanation

This type of volume or volume copy is created without a warning threshold.

User response

Retry the command without the **-warning** parameter.

CMMVC9237E **The volume specified is a thin or compressed volume in a data reduction pool. Cache can not be set to none or readonly.**

Explanation

You must enable the cache for thin-provisioned or compressed volumes in a data reduction pool.

User response

Retry the command with a value for the **-cache** parameter other than none or `readonly`.

CMMVC9238E **The volume specified is a thin volume in a data reduction pool. The -warning parameter can not be used.**

Explanation

An attempt was made to use the **-warning** parameter of the **chvdisk** command to set a warning threshold for a thin-provisioned volume in a data reduction pool. This action is not permitted. However, you can set a warning threshold at the pool level.

User response

To set a warning threshold at the pool level, use the **chmdiskgrp -warning** command.

CMMVC9240E **The action requested is not supported on volumes or volume copies in data reduction pools.**

Explanation

You cannot run this command on volumes that are in data reduction pools.

User response

Retry this command and specify a regular pool.

CMMVC9241E **The action cannot be completed because the Quorum device is unavailable.**

Explanation:
An issue exists in communicating with the Quorum device.

User response:
View the event log and resolve any outstanding issues with the Quorum device.

CMMVC9242E **The volume or volume copy specified is a thin or compressed volume in a data reduction pool. Thin provisioned or compressed volume copies created from a data reduction pool can not use the -rsize parameter.**

Explanation

You cannot specify an **-rsize** parameter for thin-provisioned or compressed volume copies that are created in a data reduction pool. Data reduction pools automatically manage the physical capacity that is used across the volumes in the pool, so you cannot shrink or expand the **-rsize** for this type of volume copy.

User response

Retry the command without the **-rsize** parameter.

CMMVC9243E **Node VPD is unavailable for inactive spare nodes. Use sainfo lsservicenodes instead.**

Explanation

The **lsnodevpd** command does not display the vital product data for inactive spare nodes. For most applications, you can use the **sainfo lsservicenodes** command instead.

User response

Use the **sainfo lsservicenodes** command to display vital product data for the node.

CMMVC9245E **The pool specified is a data reduction pool and the IO group selected already contains compressed volumes within a regular pool. An IO group using software compression can not contain compressed volumes from regular and data reduction pools at the same time.**

Explanation

An attempt was made to add a compressed volume from a data reduction pool to an I/O group that already contains at least one compressed volume from a regular pool. I/O groups that use software compression can contain compressed volumes from data reduction pools or regular pools, but not both.

User response

Retry the command and specify a different I/O group or a volume from a regular pool.

CMMVC9246E **The pool specified is a regular pool and the IO group selected already contains compressed volumes within a data reduction pool. An IO group using software compression can not contain compressed volumes from regular and data reduction pools at the same time.**

Explanation

An attempt was made to add a compressed volume from a regular pool to an I/O group that already contains at least one compressed volume from a data reduction pool. I/O groups that use software compression can contain compressed volumes from data reduction pools or regular pools, but not both.

User response

Retry the command and specify a different I/O group or a volume from a data reduction pool.

CMMVC9247E **The pool specified is a data reduction pool. Thin or compressed volumes or volume copies in data reduction pools cannot have the easytier status set or changed.**

Explanation

You cannot set the easytier mode on a volume or volume copy basis for thin-provisioned or compressed volumes in a data reduction pool. The easytier mode is an attribute of the pool and must be set or changed at the pool level.

User response

Use the **mkmdiskgrp** command to set or change the value for the **-easytier** parameter at the pool level.

CMMVC9248E **The volume or volume copy specified is a thin or compressed volume in a data reduction pool. The easytier status cannot be changed.**

Explanation

You cannot specify the easytier mode of a volume or volume copy that you create from a data reduction pool.

User response

Complete one or both of the following actions:

- Retry the command and omit the **-easytier** parameter.
- Use the **mkmdiskgrp** command to set the value for the **-easytier** parameter at the pool level.

CMMVC9249E **Data reduction pools can not have child pools.**

Explanation

An attempt was made to create a child pool within a data reduction pool.

User response

If you need child pools, you must create them within a regular pool.

CMMVC9250E **The pool specified is a data reduction pool. Thin provisioned or compressed volume copies created from a data reduction pool can not use the -buffersize parameter.**

Explanation

You cannot use the **-buffersize** parameter when you create thin-provisioned or compressed volume copies from a data reduction pool.

User response

Complete one of the following actions:

- Retry the command without the **-buffersize** parameter.
- Retry the command and specify a pool that is not a data reduction pool.

CMMVC9251E **The command failed because the -deletelatergenerations parameter was specified and the specified local volume has cloud snapshot not enabled.**

Explanation

An attempt was made to restore a cloud snapshot. However, cloud snapshots are not enabled on the specified local volume. When you use the **-deletelatergenerations** parameter, cloud snapshots must be enabled on the local volume.

User response

Complete one of the following actions:

- Restore the volume without using the **-deletelatergenerations** parameter.
- Enable cloud snapshots on the local volume by using the **chvdisk** command and then retry the original command.

CMMVC9252E **The command failed because the I/O group contains compressed volumes. The new hardware configuration does not support compressed volumes.**

Explanation

An attempt was made to submit a hardware configuration that does not support compressed volumes, where the I/O group that this node belongs to contains compressed volumes.

User response

Complete one of the following actions:

- Submit a hardware configuration that supports compressed volumes.
- Use the **rmvdisk** command to remove all compressed volumes from the I/O group, and then retry the command.

CMMVC9253E **An IO group using software compression can not contain compressed volumes from regular and data reduction pools at the same time.**

Explanation

The I/O group that was selected supports software compression only. You cannot create a compressed volume that has copies in both a regular pool and a data reduction pool.

User response

Complete one of the following actions:

- Retry the command and select a different I/O group.
- Create the compressed volume with both copies in regular pools or both copies in data reduction pools. Then retry the command.

CMMVC9254E **The command failed because the maximum number of spare nodes have already been assigned to this cluster.**

Explanation

The system already has the maximum number of assigned spare nodes, and additional spares cannot be added.

User response

Use the **rmnode** command to remove an existing spare, then add the new node

CMMVC9256E **The specified node is not an active spare node.**

Explanation

The **swapnode** command was not completed because the specified node is not a spare node.

User response

Retry the command and specify a valid spare node.

CMMVC9257E **The command failed because -deactivatespare was not specified.**

Explanation

An attempt was made to remove a node when a spare node was active in its place. You must deactivate the spare node before the specified node can be removed. The **-deactivatespare** parameter performs the required deactivation as part of the command.

User response

Retry the command and include the **-deactivatespare** parameter.

CMMVC9258E **The command failed because -deactivatespare was specified.**

Explanation

The **-deactivatespare** parameter was used to indicate that a spare node was currently taking the place of the node to be deleted. No spare is actually being used in place of the designated node.

User response

Retry the command and omit the **-deactivatespare** parameter.

CMMVC9259E **The upgrade is currently not paused**

Explanation

An attempt was made to continue a software update (**svctask applysoftware -continue**) when the update was not paused.

You can continue a software update when the **lsupdate** command shows one of the following statuses:

- system_updating_pausing
- system_restoring_pausing

User response

The most likely scenario is that the continue command was run by mistake. No user response is necessary in this case; the update continues automatically unless you see one of the "pausing" messages in **lsupdate**.

CMMVC9260E **The command failed because the specified node has been replaced by a spare.**

Explanation

The specified node was replaced by a spare, so you cannot currently complete the command. You must restore the original node first.

User response

Restore the original node by entering the **swapnode -failback** command, then retry the original command.

CMMVC9261E **The command failed because the specified node does not have a status of candidate.**

Explanation

The node does not have the required status of "candidate" in the display from the **sainfo lsservicenodes** command.

User response

If the specified node has a status of "service" in **sainfo lsservicenodes** and a node error of 690, run the **satask stopservice** command to exit service mode, then retry the original command.

If your situation does not match these conditions, or if the **stopservice** command is unsuccessful, contact your support representative.

CMMVC9262E **The pool specified is a data reduction pool. The action requested is not supported with pools of this type.**

Explanation

You cannot run this command on a data reduction pool.

User response

Retry the command and specify a regular pool.

CMMVC9264E **The task failed because unable to format the newly expanded area in the original volume's last grain**

Explanation

An I/O operation related to formatting the last grain of the original volume failed. The expansion process cannot continue. This error can occur when the volume is offline, or when I/O operations are delayed.

User response

Ensure that the volume is online and that I/O operations are not delayed. Try the command again.

CMMVC9265E **The volume specified has a thin or compressed volume copy in a data reduction pool. Thin provisioned or compressed volumes created from a data reduction pool can not use -size parameter to reduce the size of the volume.**

Explanation

You cannot use the **-size** parameter to shrink thin-provisioned or compressed volumes that were created from a data reduction pool.

User response

Retry the command and specify a regular disk or a disk from a regular pool.

CMMVC9266E **The command has failed because specified host port group id is different from the failover port's host port group id.**

Explanation

The system administrator used the **-hpgid** parameter to assign a host port group ID that did not match the host port group ID of the failover port. This error occurs only when the administrator manually edits the `svc_config.backup.xml` file that is used in T3 or T4 recovery.

User response

Retry the command and assign the same host port group ID that is used for the failover port.

CMMVC9267E **The command has failed because more than 4 ports are being assigned the same host port group id.**

Explanation

The system administrator used the **-hpgid** parameter to assign the same host port group ID to more than four ports. This error occurs only when the administrator manually edits the `svc_config.backup.xml` file that is used in T3 or T4 recovery.

User response

Retry the command and assign a host port group ID that is not already used for four ports.

CMMVC9268E **Network connection refused.**

Explanation

The network cannot connect to api.service.softlayer.com.

User response

Configure the network so that you can reach api.service.softlayer.com. Then, retry the command.

CMMVC9269E **The specified storage was not found.**

Explanation

An attempt was made to configure IBM Cloud backend storage where the **-storage** parameter referred to an invalid IBM Cloud storage name.

User response

Retry the command and use a valid IBM Cloud storage name.

CMMVC9270E **Invalid username or key.**

Explanation

An attempt was made to configure IBM Cloud backend storage where the value of the **-username** parameter or the **-key** parameter was not valid.

User response

Refer to the IBM Cloud user portal to obtain the correct IBM Cloud API user name or IBM Cloud API key. Retry the command with valid values.

CMMVC9271E **Storage access control error, possibly due to duplicate iSCSI qualified name (IQN) on IBM Cloud.**

Explanation

A storage access control error occurred, possibly due to duplicate IQNs on different bare metal servers that were detected on IBM Cloud.

User response

Ensure that the IQN on each bare metal server is unique, and then retry the command.

CMMVC9272E **The specified storage has already been configured.**

Explanation

The **-storage** parameter refers to IBM Cloud storage that is already configured.

User response

If you specified the correct IBM Cloud storage name, no further action is necessary. If the name you specified was incorrect, retry the command with the correct storage name.

CMMVC9273E **The specified storage is already in use by another cluster.**

Explanation

The **-storage** parameter refers to IBM Cloud storage that is already in use by another cluster.

User response

Retry the command and specify a valid storage name.

CMMVC9274E **Some of the serial numbers are not correct.**

Explanation

The product panel name is a unique serial number from the IBM Cloud web page. This name is a mandatory parameter for the **initnode** command. The **initnode** command was run with an incorrect serial number.

User response

Retry the **initnode** command with the correct serial number. Then, retry the **cfgcloudstorage** command.

CMMVC9277E **The command failed because there are out of sync volume mirrors and some copies are in data reduction pools that may also be deleted.**

Explanation

The **rmmdiskgrp** command was attempted when at least one volume mirror is out of sync and the out-of-sync copy is in a data reduction pool.

User response

Synchronize the volume mirrors in all cases where the volume copy is in a data reduction pool. Then retry the **rmmdiskgrp** command.

CMMVC9279E **The task failed because there is no up-to-date node in the IO group that contains the volume.**

Explanation

This error can occur when both nodes in the I/O group of the volume are offline, or when some nodes in the I/O group are online but they have stale information.

User response

Ensure that each node in the I/O group is online and has no stale information, and then retry the command. Ensure that the nodes stay online while the command runs.

CMMVC9280E **The command did not complete because creating mirror bitmap is not finished within the allotted time period.**

Explanation

The mirror bitmap must be created successfully before the command can continue. The maximum amount of time that is allotted for the command expired while the bitmap was being created.

User response

Run the command again. If the problem persists, contact IBM Support.

CMMVC9281E **One or more of the nodes in the selected I/O group has no working compression hardware.**

Explanation:
It is not possible to create a compressed volume in an I/O group that has nodes with no working compression hardware. Before you create a compressed volume, the issue with the hardware needs to be fixed.

User response:
View the event log and follow the directed maintenance procedure to fix the issue with the compression hardware.

CMMVC9282E **The specified iSCSI Storage port uses per node credentials. Resetting credentials for all nodes can temporarily cause MDisks to go offline. Use the -force option to make this change.**

Explanation

An attempt was made to reset the authentication credentials for all initiator nodes in the system by running a single command on the command line. This action can be completed only when the **-f** (force) option is used.

User response

Complete one of the following actions:

- Retry the command for a single initiator node.
- Retry the command for all initiator nodes in the system and include the **-f** (force) option.

CMMVC9283E **The command failed because one of the storage pools for the volume is a data reduction pool that has insufficient capacity.**

Explanation

Insufficient capacity exists within the pool that is specified to be able to create a thin or compressed volume within a data reduction pool.

User response

Add extra capacity to the specified storage pool and then run the command again.

CMMVC9284E **The command failed because the pool specified is a data reduction pool that contains offline volumes.**

Explanation

An attempt was made to create a volume in a data reduction pool where one or more volumes are offline. Volumes can be offline for any of the following reasons:

- The data reduction pool is out of capacity.
- The data reduction pool is corrupt.
- The data reduction pool contains corrupt volumes.
- The data reduction pool contains offline MDisks.

User response

Resolve the problem that is causing the volume or volumes to be held offline and retry the command. You might need to complete any of the following tasks:

- Add capacity to the data reduction pool by using the **addmdisk** command, or use the **mkarray** command if you use internal drives.
- Repair corrupt volumes.
- Repair a corrupt data reduction pool by using the **recovervdisk** command.

CMMVC9285E **A thin or compressed volume in a data reduction pool cannot be created within this I/O group because the amount of memory configured for FlashCopy is greater than 1.5 gigabytes (GB) and there is insufficient free memory in the I/O group.**

Explanation

For I/O groups that contain 8 GB node types, thin or compressed volumes cannot be created if the FlashCopy bitmap is larger than 1.5 GB.

User response

Complete one of the following actions:

- Upgrade the node types within the I/O group to contain more than 8 GB of memory.
- Remove data reduction volumes within the specified I/O group.
- Select another I/O group to create the thin or compressed volumes in a data reduction pool.
- Reduce the amount of FlashCopy bitmap memory that is in use for the I/O group to below 1.5 GB.

Then, retry the original command.

CMMVC9286E **The amount of memory configured for FlashCopy cannot be greater than 1.5 gigabytes (GB) because data reduction volumes exist in this I/O group**

Explanation

For I/O groups that contain 8 GB node types and thin or compressed volumes, the FlashCopy bitmap cannot be increased beyond 1.5 GB.

User response

Complete one of the following actions:

- Upgrade the node types within the I/O group to contain more than 8 GB of memory.
- Remove any data reduction volumes within that I/O group.
- Set the FlashCopy bitmap size under 1.5 GB.
- Specify a different I/O group to use for flash copies.

Then, retry the original command.

CMMVC9287E **The volume or volume copy cannot be deleted because the data reduction pool is corrupt.**

Explanation

An attempt was made to delete a volume when the data reduction pool where it was located was marked as corrupted.

User response

Complete one of the following actions:

- Repair the data reduction pool by using the **recovervdisk** command.
- Delete the entire pool by using the **rmmdiskgrp -force** command.

CMMVC9288E **The node could not be added as a data reduction pool containing a thin or compressed volume exists within this I/O group and the new node does not meet the minimum CPU requirements.**

Explanation

When the first thin or compressed volume in a data reduction pool is created for an I/O group, the I/O group sets a minimum CPU threshold. This threshold is based on the lowest number of CPU resources available to any node in the I/O group. A new node with fewer CPU resources cannot be added to the I/O group.

User response

Complete one of the following actions:

- If possible, add another CPU to the node.
- Specify a different node with the same type and CPU as the nodes that are currently in the system.
- If neither of the previous options is available, delete all thin and compressed volumes in all data reduction pools for the specified IO group.

Then, retry the **addnode** command.

CMMVC9289E **The command failed because the drives selected have an incompatible combination of technology type and transport protocol.**

Explanation

This error occurs when an attempt is made to mix NVMe drives in an array with drives that use a different transport protocol or technology type.

User response

Create an array that uses only SAS drives, or only NVMe drives of a single technology type. Ensure that any drives added to the array are compatible with the existing array members.

CMMVC9290E **The mkarray command failed because -level raid5 and raid6 are not supported by this IO group.**

Explanation

This error occurs because the hardware for the I/O group does not support the specified RAID level.

To create a distributed RAID array, use the **mkdistributedarray** command.

User response

Create an array that uses a supported RAID level.

CMMVC9291E **Task failed because volume is non-mirrored with fully allocated copy and is part of a Metro Mirror or Global Mirror relationship with -nofmtdisk specified**

Explanation

To expand a non-mirrored volume that is fully allocated and that is part of a Metro Mirror or Global Mirror relationship, the region of expansion must be formatted. In this case, the **-nofmtdisk** parameter cannot be used to skip formatting.

User response

Retry the command, omitting the **-nofmtdisk** parameter.

CMMVC9292E **An IO group can not contain deduplicated volumes or volume copies in data reduction pools at the same time as compressed volumes or volume copies in regular storage pools.**

Explanation

Deduplicated volumes and volume copies cannot be created in an I/O group when compressed volumes or volume copies are in regular storage pools.

User response

Migrate all compressed volumes in the I/O group from regular storage pools to data reduction pools, and then retry the action. Alternatively, create the deduplicated volume or volume copy in a different I/O group.

CMMVC9294E **Unable to add host to storage.**

Explanation

The bare metal server is not authorized as a valid host to the block storage in IBM Cloud.

User response

On the IBM Cloud user portal, authorize the bare metal server as a valid host to the block storage. Then, retry the command.

CMMVC9295E **The command failed because this raid level is not supported with compressed drives.**

Explanation

This error occurs if an attempt is made to create a non-distributed array that uses compressed drives. Non-distributed arrays do not support compressed drives.

User response

Create a distributed array with compressed drives. Ensure that any drives that are used by the array are compatible.

CMMVC9296E **The command has failed because the volume logical block address (LBA) supplied is not recoverable and contains a virtual medium error.**

Explanation:
The operation has failed because the data at the supplied volume LBA cannot be recovered and the volume LBA cannot look up the physical address.

User response:
Restore the data to the virtual LBA from a backup or from a previously generated backup of the volume or repair the volume from the host application.

CMMVC9298E **The command has failed because there is already an analysis in progress on the specified volume.**

Explanation

An analysis cannot be queued on the specified volume because an analysis is already in progress.

User response

Check whether the correct volume was specified. Check the status of the volume that was specified.

CMMVC9299E **The command has failed because there is not an analysis in progress on the specified volume.**

Explanation

An analysis cannot be canceled on the specified volume because an analysis is not in progress or queued.

User response

Check whether the correct volume was specified. Check the status of the volume that was specified.

CMMVC9300E **The specified IO group does not support deduplication.**

Explanation

Deduplicated volumes and volume copies can be created only in an I/O group where both nodes have a minimum of 128 GB of memory.

User response

Specify an I/O group where both nodes have a minimum of 128 GB of memory.

CMMVC9301E **One of the pools specified is not a data reduction pool and cannot be used to create a deduplicated volume or volume copy.**

Explanation

Deduplicated volumes and volume copies must be created from a data reduction pool.

User response

Use a data reduction pool to create the deduplicated volume or volume copy.

CMMVC9302E **The command failed because it is not possible to mix compressed and non-compressed drives in an array.**

Explanation

Arrays must contain members that are either all compressed drives or all non-compressed drives.

User response

Select a new drive that has the same compression characteristics as the existing array members.

CMMVC9303E **The command failed because one or more required drives did not complete format.**

Explanation

This error occurs when one or more drives that are candidates for use in an array do not complete the formatting process in the required time. All NVMe drives in the system must be formatted before an array can be created.

User response

Wait for the formatting process to complete for all drives. You can use the **`lsdriveprogress`** command to monitor the formatting process.

CMMVC9304E **The command failed because the driveclass does not support the specified number of rebuild areas.**

Explanation

An FCM drives array cannot be created with more than 1 rebuild area. This error can occur when an attempt is made to expand a thin-provisioned array and to add rebuild areas to the array simultaneously, which is not supported.

User response

If the error occurred during array creation, Retry the command, specifying the parameter **-totalrebuildareas 1**.

If the error occurred during expansion of a thin-provisioned array, retry the command but do not attempt to add rebuild areas to the array.

CMMVC9305E The Cloud Call Home feature is not supported on this platform.

Explanation

This system does not support sending Call Home information directly to a server in the cloud.

User response

Not applicable.

CMMVC9306E The Cloud Call Home feature is not enabled.

Explanation

The Call Home function to send notifications to the support center via cloud services is not enabled.

User response

Enable the Call Home function for cloud services by issuing the following command:

```
svctask chcloudcallhome -enable
```

CMMVC9307E No response from the Cloud Call Home service.

Explanation

Call Home information cannot be sent via cloud services. This error can occur if network issues exist or if the cloud services are currently unavailable.

User response

Ensure that port 443 is open in any firewalls to allow secure HTTP connections to the Call Home servers in the cloud.

- If you are using DNS, allow outbound traffic over port 443 to `esupport.ibm.com`.
- If you are not using DNS, allow outbound traffic over port 443 to the following IP addresses: 129.42.54.189, 129.42.56.189, and 129.42.60.189.

CMMVC9308E The command failed because the requested certificate is incompatible with the current SSL protocol level.

Explanation

The requested certificate type does not match the certificate types that are allowed by the current security level setting on the system.

User response

Specify a certificate type that is supported by the current SSL security level. Alternately, use the **chsecurity** command to reduce the SSL security level, and then retry the command.

CMMVC9309E The command failed because the SSL protocol level is incompatible with the current system certificate.

Explanation

The specified level of SSL protocol allows only ciphers that the current certificate does not allow because of the way the certificate is signed. This error occurs when the combination of the specific security level and the currently stored certificate would prevent the GUI from working. For example, this error occurs when a security level is specified that prohibits RSA key exchange, but the current certificate is RSA-signed.

User response

Use the **chsystemcert** command to generate a new system certificate. Specify an SSL certificate key type which the system is using (ECDSA and RSA).

CMMVC9310E The command failed because one or more drive format was still in progress.

Explanation

This error can occur when a drive is being formatted and the **mkarray** command is run, if the formatting process did not complete in the allotted time.

User response

Retry the command after the format completes. You can use the **lsdriveprogress** command to view the progress of the formatting process.

CMMVC9312E **The command failed because a compressed drive cannot be exchanged with a compressed drive of a different drive class.**

Explanation

In an array of compressed drives, all members must have the same drive class.

User response

Exchange the member drive with a compressed drive that has the same drive class.

CMMVC9315E **The create failed because there are no IO groups supporting deduplication.**

Explanation

Deduplicated volumes and volume copies can be created only in an I/O group where both nodes have a minimum of 32 GB of memory.

User response

Specify an I/O group where both nodes have at least 32 GB of memory.

CMMVC9316E **The create failed because there are no IO groups supporting deduplication in site**

Explanation

Deduplicated volumes and volume copies can be created only in an I/O group where both nodes have a minimum of 32 GB of memory.

User response

Specify a site that contains an I/O group where both nodes have at least 32 GB of memory.

CMMVC9318E **Compressed volumes in the selected I/O group must be in a data reduction pool.**

Explanation

An operation was attempted that would result in storing a real-tim compressed volume copy in an I/O group that does not support real-time compression.

User response

Retry the command and specify either an I/O group that supports real-time compression, or specify a data reduction storage pool.

CMMVC9319E **Volumes or volume copies which are thin provisioned or compressed and created from a data reduction pool can only be repaired with the assistance of support.**

Explanation

This error is displayed when the **repairsevdiskcopy** command is run on a disk in a data reduction pool.

User response

Contact IBM Support Center for assistance.

CMMVC9321E **Inventory data is not yet available. Please retry the command in a few minutes.**

Explanation

Immediately after cluster creation, manual update, or failover, inventory data is not available for a short time because it is being gathered in the background.

User response

Wait a few minutes and then retry the command.

CMMVC9323E **The command failed because the specified host protocol is not compatible with the specified port type.**

Explanation

A host that uses the Non-Volatile Memory express (NVMe) protocol can contain only NMVe Qualified Names (NQNs). A host that uses the SCSI protocol cannot contain NQNs. Host type must match the protocol.

User response

- For the host that uses port type **-saswwpn**, the protocol must be SAS or empty.
- For the host that uses port type **-fcwwpn**, the protocol must be FCSCSI or empty.
- For the host that uses port type **-iscsiname**, the protocol must be iSCSI or empty.

- For the host that uses port type **-nqn**, the protocol must be FCNVMe or RDMANVMe.

CMMVC9324E **The command failed because the specified host type is not compatible with the specified protocol.**

Explanation

A host that uses the Non-Volatile Memory express (NVMe) protocol must be of type **generic**. Other host types are not supported.

User response

Change the host protocol or type, and then retry the command.

CMMVC9325E **The command failed because the maximum number of host ports for the system would be exceeded.**

Explanation

Adding the specified ports would result in a total number of host ports that exceed the limit that can be configured on the system.

User response

Remove unused hosts or host ports, and then retry the command.

CMMVC9326E **The command failed because the maximum number of host ports for an I/O group would be exceeded.**

Explanation

Adding the specified ports would result in a total number of host ports that exceeds the limit that can be configured in an I/O group.

User response

Identify I/O groups where the limit would be exceeded and remove unused hosts or host ports, then retry the command. Alternately, specify a different set of I/O groups that the host can access volumes from.

CMMVC9327E **The command failed because the maximum number of NQNs for an I/O group would be exceeded.**

Explanation

Adding the specified NVMe Qualified Names (NQNs) would result in a total number of NQNs that exceeds the limit that can be configured in an I/O group.

User response

Identify I/O groups where the limit would be exceeded and remove unused NVMe hosts or NQNs, and then retry the command. Alternately, specify a different set of I/O groups that the host can access volumes from.

CMMVC9328E **The command failed because one of the specified NQNs is already assigned to a host.**

Explanation

A Non-Volatile Memory express Qualified Name (NQN) can be assigned to only one host. NQNs that differ only by case are not considered to be unique.

User response

Identify the existing host that contains the specified NQN. Remove the NQN from that host, and then retry the command.

CMMVC9329E **The command failed because the specified host type is not compatible with the current host protocol.**

Explanation

A host that uses the Non-Volatile Memory Express (NVMe) protocol must be of type **generic**.

User response

You cannot change the host type for a host that uses the NVMe protocol.

CMMVC9330E **The command failed because the specified host protocol is not compatible with the current host type.**

Explanation

A host that uses the Non-Volatile Memory express (NVMe) protocol must be of type **generic**.

User response

Change the host type to **generic**, and then retry the command to change the host protocol.

CMMVC9331E **The host protocol cannot be changed because the host has ports configured.**

Explanation

The host protocol cannot be changed while ports are configured for the host.

User response

Remove all ports from the host, and then retry the command.

CMMVC9332E **The host protocol cannot be changed because there are volumes mapped to the host.**

Explanation

The host protocol cannot be changed while volumes are mapped to the host.

User response

Remove all mappings to the host, and then retry the command.

CMMVC9334E **The command failed as the volume is already mapped to host/s with a protocol inconsistent with the specified host.**

Explanation

The volume cannot be mapped to hosts that use the SCSI protocol and to hosts that use the NVMe protocol.

User response

Ensure that all the hosts that are mapped to the volume use the same protocol.

CMMVC9335E **The command failed as the -scsi parameter is not supported for hosts with protocol set as NVMe.**

Explanation

The SCSI ID cannot be specified when you map a volume to a host that uses the NVMe protocol.

User response

Retry the command without specifying the **-scsi** parameter.

CMMVC9336E **The command failed because the system does not support NVMe host attachment.**

Explanation

The system does not support host connections that use the Non-Volatile Memory express (NVMe) protocol.

User response

Use an alternative protocol for host connections.

CMMVC9337E **The command failed because the specified port type is not compatible with the host protocol.**

Explanation

A host that uses the Non-Volatile Memory express (NVMe) protocol can contain only NMVe Qualified Names (NQNs). A host that uses the SCSI protocol cannot contain NQNs. Host type must match the protocol.

User response

- For the host that uses port type **-saswwpn**, the protocol must be SAS or empty.
- For the host that uses port type **-fcwwpn**, the protocol must be FCSCSI or empty.
- For the host that uses port type **-iscsiname**, the protocol must be iSCSI or empty.
- For the host that uses port type **-nqn**, the protocol must be FCNVMe or RDMANVMe.

CMMVC9338E **The command failed because the maximum number of ports for the host would be exceeded.**

Explanation

Adding the specified ports would result in a total number of ports that exceed the limit that can be configured for a host.

User response

Identify and remove any unused ports from the host before you retry the command. Alternately, create a host that uses the specified ports.

CMMVC9339E **The command failed because the maximum number of NQNs for the host would be exceeded.**

Explanation

Adding the specified Non-Volatile Memory express (NVMe) Qualified Names would result in the total

number of NQNs exceeding the limit that can be configured for a host.

User response

Identify and remove any unused NQNs from the host before you retry the command. Alternately, create a new host that uses the specified NQNs.

CMMVC9340E **The command failed because the specified NQN is not assigned to the specified host.**

Explanation

The specified NVMe Qualified Name (NQN) cannot be removed because it is not assigned to the specified host.

User response

Check the specified NQN and host, and then retry the command and specify an NQN that is assigned to the host.

CMMVC9341E **The command failed because Source IP Address specified via -src_ip parameter does not exist or port id specified via -src_port_id parameter does not exist.**

Explanation

This error can occur if the specified source IP address or port ID does not exist. The source IP address and port ID must be in the list that is displayed by the **svcinfo lsportip** command.

User response

Run the **svcinfo lsportip** command to determine the correct source IP address and port ID to specify.

CMMVC9342E **The command failed because Node ID specified via -node_id parameter does not exist or is in OFFLINE state.**

Explanation

This error can occur if the specified node is not active.

User response

Specify the ID of a node that has active or online status, as shown by the **svcinfo lsnode** or **svcinfo lsnodecanister** commands.

CMMVC9343E **The command failed because the host protocol does not match the existing hosts in the host cluster.**

Explanation

All hosts in a host cluster must use the same protocol.

User response

Specify a host cluster that uses a protocol that matches the protocol of the host.

CMMVC9344E **The command failed because the protocols of the specified hosts are different.**

Explanation

All hosts in a host cluster must use the same protocol.

User response

Specify hosts that use the same protocol.

CMMVC9345E **The host protocol cannot be changed because the host is a member of a host cluster.**

Explanation

All hosts in a host cluster must use the same protocol.

User response

Remove the host from the host cluster, and then retry the command.

CMMVC9352E **The command failed as the volume is already mapped to hosts with a protocol inconsistent with the specified host cluster.**

Explanation

A volume cannot be mapped to host clusters that use the SCSI protocol and to host clusters that use the NVMe protocol at the same time.

User response

Make sure that all of the host clusters that are mapped to the volume use the same protocol.

CMMVC9353E **The command failed as the -scsi parameter is not supported for host clusters with protocol set as NVMe.**

Explanation

The SCSI ID cannot be specified when you map a volume to a host cluster that uses the NVMe protocol.

User response

Retry the command, omitting the **-scsi** parameter.

CMMVC9354E The command failed because the array has an expansion ongoing.

Explanation

This error can occur when an attempt is made to start another array expansion or to change members of a managed disk while an array expansion is already in progress.

User response

- To remove a drive from membership in the array, use the **chdrive -failed** command.
- To change the membership of the array, use the **charraymember -immediate** command.
- To expand the array, wait until the current expansion process completes, and then retry the command. You can monitor the expansion process by using the **lsarrayexpansionprogress** command.

CMMVC9355E The command failed because the array has a member resync ongoing.

Explanation

This error can occur when an attempt is made to run a command while array member synchronization is happening because of high latency.

User response

Wait for the synchronization process to complete, and then retry the command. You can monitor the synchronization process by using the **lsarraymemberprogress** command.

CMMVC9356E The command failed because an array in the parent pool has an expansion ongoing.

Explanation

This error can occur when an attempt is made to expand an array when there is already an array in the parent pool that is being expanded. Only one array per parent pool can be expanded at a time.

User response

Wait for the current expansion process to complete, and then retry the command. You can monitor the expansion process by using the **lsarrayexpansionprogress** command.

CMMVC9357E The command failed because the array does not support expansion.

Explanation

This error can occur when an attempt is made to expand an unsupported array. Array expansion is supported only on distributed arrays of SAS drives.

User response

Run the command on a distributed array of SAS drives.

CMMVC9358E An unknown error condition occurred while setting up the self encrypting drives.

Explanation

An unspecified error occurred while encryption was being configured on one or more drives.

User response

Format the candidate drives, and then retry the command. If the issue persists, contact IBM Support Center.

CMMVC9359E The command failed because the system has reached the maximum number of ongoing expansions.

Explanation

This error can occur when an attempt is made to expand too many arrays. Only four arrays and only one array per storage pool can be expanded at the same time.

User response

Wait for an array expansion to complete, and then retry the command. You can monitor the expansion process by using the **lsarrayexpansionprogress** command.

CMMVC9360E The command failed because the array specified is offline.

Explanation

This error can occur when an attempt is made to expand an array that is offline.

User response

Bring the array online, and then retry the command. You can use the **lsarray** command to check the status of an array.

CMMVC9361E **The command failed because the array has rebuild areas in use**

Explanation

This error can occur when an attempt is made to expand an array that has rebuild areas that are in use. To run an expansion, all rebuild areas must be inactive.

User response

- Fix any errors that are associated with the drives in the array.
- Run the **lsarray** command until the number of rebuild areas available (`rebuild_areas_available`) is the same as the total number of rebuild areas (`rebuild_areas_total`). When the number of available rebuild areas is the same as the total number of rebuild areas, the rebuild areas are inactive.
- Retry the command.

CMMVC9362E **The command failed because authentication credentials are not set by chiscsiportauth CLI.**

Explanation

If the authentication information is not set for a node port, an empty source iSCSI qualified name (IQN) is used. An iSCSI login message cannot contain an empty source IQN. This error applies to IBM Cloud only.

User response

Use the **chiscsiportauth** command to set the authentication information for the port on the configuration node.

CMMVC9363E **The command failed because the array specified has a component in doubt.**

Explanation

This error can occur when an attempt is made to expand an array that has components that might be about to fail.

User response

Fix all errors that are related to the specified array, and then retry the command.

CMMVC9364E **The command failed because the array status was not stable when the command was executed. Please retry the command if you want to carry on with the expansion.**

Explanation

This error can occur when the status of an array changes during an expansion process.

User response

Wait until the array status is stable, and then retry the command. You can use the **lsarray** command to determine the status of an array.

CMMVC9365E **The node cannot be removed because it would result in an I/O group outage. Use -force flag to override.**

Explanation

The **rmnode** command cannot be used to remove the last node in an I/O group unless you specify the **-force** parameter.

User response

Ensure that you want to remove the last node in the I/O group. If you are certain that you want to remove the last node in the I/O group, enter the **rmnode** command with the **-force** parameter.

CMMVC9366E **The command failed because the additional drive count must exceed the additional rebuild area count.**

Explanation

This error occurs when an attempt is made to expand an array, but the number of rebuild areas to add is greater than or equal to the number of drives to add.

User response

Retry the command, specifying a number of drives to add that is greater than the number of rebuild areas to add.

CMMVC9368E **The command failed because the additional rebuild area**

count would result in the array exceeding the maximum rebuild area count.

Explanation

This error occurs when an attempt is made to expand an array, but the specified number of rebuild areas to add is not supported. The maximum number of rebuild areas is four.

User response

Retry the command, specifying a number of rebuild areas to add that does not exceed the maximum.

CMMVC9369E **The command failed because there are no IO groups that support both compression and deduplication**

Explanation

A compressed and deduplicated volume cannot be created in the clustered system because no I/O group that supports both deduplication and compression is available.

This message is displayed when the clustered system contains I/O groups that support compression but not deduplication and I/O groups that support deduplication but not compression, but contains no I/O groups that support both deduplication and compression.

User response

Add an I/O group that supports both compression and deduplication to the system.

CMMVC9370E **The command failed because there are no IO groups that support both compression and deduplication in site *site_id*.**

Explanation

A compressed and deduplicated volume cannot be created in the clustered system at the specified site because no I/O group that supports both deduplication and compression is available.

This message is displayed when the clustered system contains I/O groups that support compression but not deduplication and I/O groups that support deduplication but not compression, but contains no I/O groups that support both deduplication and compression.

User response

Add an I/O group that supports both compression and deduplication to the site.

CMMVC9371E **The command failed because the I/O group specified is not in the same site as the storage pool.**

Explanation

Creating a volume with a thin or compressed copy in a data reduction pool on a different site from the site of the I/O group is not allowed.

This message is displayed when an attempt is made to create a thin or compressed volume with a copy in a data reduction pool on a system that uses a HyperSwap topology.

User response

Retry the command, specifying storage pools that are on the same site as the I/O group of the volume.

CMMVC9372E **The command failed because the storage pool specified is not in the same site as the volume.**

Explanation

Creating a thin or compressed volume copy in a data reduction pool on a different site from the I/O group of the volume is not allowed.

This message is displayed when an attempt is made to create a thin or compressed volume copy in a data reduction pool on a system that uses a HyperSwap topology.

User response

Retry the command, specifying a storage pool on the same site as the I/O group of the volume.

CMMVC9373E **The command failed because the drive is only supported for Distributed Arrays, which do not use spares.**

Explanation

An attempt was made to change the role of a FlashCore Module (FCM) so that it can be used as a hot-spare drive. This error can occur when the `chdrive -use spare` *`drive_id`* command is run on an FCM. Only distributed arrays are supported on FCMs. FCMs cannot be used as hot-spare drives.

User response

Do not attempt to change the role of an FCM so that it can be used as a hot-spare drive. If you are replacing failed drives in an array, insert the replacement drives directly into the same slot as the failed drive, which causes the system to automatically configure the FCM as a replacement member.

If the role of the replacement drive was already changed, inserting it in the same slot does not cause the system to automatically replace the previous member drive in the array. In this situation, use the **charraymember** command to add the replacement candidate drive directly into the array. Change the role of the old drive to **unused** to remove it from the system.

If the member drive was removed from the array by changing its role from **failed** to **unused**, you must use the **charraymember** command to service the array.

CMMVC9374E **The command failed because volume mappings exist and the host is a member of a host cluster.**

Explanation

This error can occur when an attempt is made to delete mappings between I/O groups and a host when the host is a member of a host cluster and is mapped to volumes in the I/O groups that are being removed.

User response

Remove the host from the host cluster, and then retry the command.

CMMVC9375E **The command failed because the volume is already mapped to another host. Use -force to create the mapping.**

Explanation

This error can occur when an attempt is made to create a mapping between a volume and a host, when the volume is already mapped to another host. If a volume is mapped to multiple hosts, the hosts must coordinate access to avoid corrupting data.

User response

To map the volume to multiple hosts, create the mapping by retrying the command and specifying the **-force** flag. Otherwise, remove the existing host mapping to the volume and then retry the command.

CMMVC9376E **The node could not be added as the code level of new node is different with the code level of the cluster in this platform.**

Explanation

This error can occur when an attempt is made to add a cloud node that is running a different version of software from the existing system.

User response

Install the same version of software on the new cloud node as is installed on the existing system, and then add the node to the system.

CMMVC9377W **There are dependent FlashCopy mappings not in your ownership group.**

Explanation

This warning occurs when all dependent FlashCopy mappings cannot be displayed because some of the mappings are not in the ownership group of the user who ran the command.

User response

There is no user response. If you delete a FlashCopy mapping that you own, it can affect other objects outside your ownership group.

CMMVC9378E **The command failed because the specified ownership group is associated with one or more objects. To continue specify -keepobjects.**

Explanation

This error can occur if an attempt is made to delete an ownership group that is still associated with objects.

User response

Complete one of the following steps:

- To remove the ownership group without removing the associated objects from the system, retry the command and specify the **-keepobjects** option.
- Change or delete the ownership group for all associated objects, and then retry the command.

CMMVC9379E **The command failed because it cannot be executed by a user in an ownership group.**

Explanation

This error can occur if a command is specified that affect objects that are not in the user's ownership group.

User response

Contact someone with the Administrator or Security Administrator role and have them run the command.

CMMVC9380E **The command failed because quorum site cannot be set in standard quorum mode.**

Explanation

This error can occur when an attempt is made to assign a quorum site while in `standard` quorum mode. Quorum mode can be set to `standard`, `winner`, or `preferred`.

User response

Use the **`chsystem -quorummode`** command to change the quorum mode to `winner` or `preferred`, and then retry the command.

CMMVC9381E **The required update package can not be found and downloaded from the support server in this platform.**

Explanation

When the **`svctask applysoftware`** command runs, the software package is downloaded from the support server. This error can occur if the software package is not downloaded from the support server to the configuration node when the **`svctask applysoftware`** command runs.

User response

Check the network connectivity between the configuration node and the support server, and then retry the command. If the error message is displayed again, contact IBM Support Center.

CMMVC9384E **The command failed because the specified ownership group does not exist.**

Explanation

This error occurs when a command is run that specifies an ownership group that does not exist.

User response

Check list of valid ownership groups using CLI lsownershipgroup or create a new ownership group using CLI mkownershipgroup.

CMMVC9385E **The command failed because one or more parameters cannot be specified by a user in an ownership group.**

Explanation

This error can occur if a parameter is specified that would affect objects that are not in the user's ownership group.

User response

Contact someone with the Administrator or Security Administrator role and have them run the command.

CMMVC9386E **The command failed because it would result in inconsistent ownership groups and the -inconsistentownershipgroup parameter was not specified.**

Explanation

This error can occur if an attempt is made to run a command that would result in an inconsistent ownership groups, and the **`-inconsistentownershipgroup`** parameter is not specified. For example, this error is displayed if the result of a command would be that related objects are in different ownership groups.

User response

Retry the command, specifying the **`-inconsistentownershipgroup`** parameter to use different ownership groups for related objects.

CMMVC9387E **The command failed because the seeding hosts are in more than one ownership group.**

Explanation

This error occurs when an attempt is made to create a host cluster that specifies seeding hosts that are in different ownership groups.

User response

Retry the command, specifying a list of seeding hosts that are in only one ownership group.

CMMVC9388E **The command failed because seeding hosts were owned and that ownership group needs to be specified using the -ownershipgroup parameter.**

Explanation

This error can occur when an attempt is made to create a host cluster and the seeding hosts are in an ownership group but the ownership group is not specified.

User response

Retry the command, specifying the **-ownershipgroup** parameter.

CMMVC9389E **The command failed because there are already max iSCSI (initiator + target) sessions present on node.**

Explanation

This error can occur when an attempt is made to establish more than 1024 iSCSI connections on a node, including host attachment and external storage connections.

User response

Remove unneeded iSCSI connections, and then retry the command. You might need to remove host attachment connections to create more iSCSI connections to external storage.

CMMVC9390E **The command failed because the SecurityAdmin group and ownership groups are mutually exclusive.**

Explanation

This error can occur if an attempt is made to run a command that would result in an ownership group that is also a user group with SecurityAdmin privileges.

User response

Retry the command, specifying a user group that either has no ownership group or no SecurityAdmin privileges.

CMMVC9391E **The command failed because the thin provisioned array is write protected.**

Explanation

This error can occur when an attempt is made to expand an array that is in a read-only state because the array has no free physical capacity.

User response

Free up space on the array by moving volumes to other arrays or by deleting volumes, and then retry the command. If volumes cannot be moved or deleted, contact IBM Support Center.

CMMVC9392E **The command failed because the cache state for a volume or copy is corrupt.**

Explanation

This error can occur if the cache for a volume or volume copy is corrupted.

User response

Use the **recovervdisk** command to recover the volume and bring it back online, and then retry the command.

CMMVC9393E **The network is unreachable.**

Explanation

This error occurred because the system is unable to access the Public Cloud API services.This error can occur when the system is unable to communicate with the Amazon Virtual Private Cloud (VPC) endpoint to access the Amazon Elastic Compute Cloud (EC2).

User response

Check the network settings.

- Check network settings and modify them as needed to establish a connection with the VPC endpoint.
- Ensure that DNS is configured correctly on the system.

CMMVC9394E **The node is not authorized to complete this operation.**

Explanation

This error can occur when the Identity and Access Management (IAM) role does not have the permissions that are required to run the command. Access in Amazon Web Services (AWS) is managed by creating policies and attaching them to IAM roles. IAM roles are then assigned to Elastic Compute Cloud (EC2) instances. An IAM role with the required permissions is bundled with the product.

User response

Log on to the Public Cloud web console, then check and add the required permission policy to this node.Follow the Directed Maintenance Procedures (DMP) to log on to the AWS management console and add the required permissions to the IAM role of the EC2 instance.

CMMVC9395E **The specified local managed disk is not in a valid state.**

Explanation

This error can occur if an attempt is made to add local managed disks to a storage pool or image volume when the managed disks are not in the `available` or `in-use` state and in the `unmanaged` access mode. Local managed disks are cloud storage that is located on a cloud service provider.

User response

Select a local managed disk that is in the `available` or `in-use` state and in the `unmanaged` access mode, and then retry the command. You can use the **`lslocaldisk`** command to determine the states of local managed disks.

CMMVC9396E **The system already manages the maximum number of local managed disks.**

Explanation

This error can occur when an attempt is made to manage more than 16 local managed disks on a system. Local managed disks are cloud storage that is located on a supported cloud service provider.

User response

Do not attempt to manage more than 16 local managed disks on a system.

CMMVC9397E **The local managed disks were not attached to the node in the allotted time.**

Explanation

This error can occur when an attempt is made to add multiple Amazon Elastic Block Storage (EBS) volumes as local managed disks. If all of the EBS volumes are not attached to the system node within 160 seconds of entering the command, this error message is displayed. Local managed disks are cloud storage that is located on Amazon Elastic Block Storage (EBS) volumes.

User response

Retry the command, adding fewer managed disks. If the error occurs when you attempt to add a single local managed disk, check the event log, resolved all unfixed errors, and then resubmit the command.

CMMVC9399E **The command failed because the EasyTier FCM over allocation max ratio is out of range.**

Explanation

This error occurs when the specified IBM Easy Tier overallocation maximum ratio for a FlashCore Module (FCM) is out of range. The supported range for the overallocation maximum ratio is between 100% and 400%.

User response

Retry the command and specify an overallocation maximum ratio between 100% and 400%.

CMMVC9402E **The SCSI LUN ID currently allocated to the volume mapping is not available in all of the new access IO groups.**

Explanation

This error can occur when a SCSI LUN ID is not available in all of the access I/O groups that are being added. The specified volume mapping has a SCSI LUN ID that is allocated to it in the access I/O groups in the current access I/O group set. The SCSI LUN ID is the same for all access I/O groups in the access I/O group set, but is not available in all of the new access I/O groups that are being added.

User response

Remove the host mapping in the new access I/O groups that is using the required SCSI LUN ID and then retry the command. Alternately, if the host supports mismatched SCSI LUN IDs, use the **`-allowmismatchedscsiids`** option to assign a different SCSI LUN ID to the volume.

CMMVC9403E **The SCSI LUN IDs currently allocated to the volume mapping are different in existing access IO groups.**

Explanation

This error can occur when the system cannot allocate the same SCSI LUN IDs in all access I/O groups because different SCSI LUN IDs are allocated to the

volume mapping in the access I/O groups in the current access I/O group set.

User response

Complete one of the following steps.

- Remove access from the access I/O groups that have mismatched SCSI LUN IDs and then retry the command, including the removed access I/O group and the new access I/O group to allocate matching SCSI LUN IDs. Make sure that the required SCSI LUN ID is available in all access I/O groups.
- Remove the host mapping to the volume that is using different SCSI LUN IDs and then re-create the host mapping, including the new access I/O group with matching SCSI LUN IDs.
- Use the **-allowmismatchedscsiids** option to assign the volume a different SCSI LUN ID from the existing IDs. Not all hosts support this configuration.

CMMVC9404E **An identical SCSI LUN ID is not available in all access IO groups.**

Explanation

This error can occur when the system is not able to allocate an identical SCSI LUN ID in all the volume access I/O groups because identical SCSI LUN IDs are not available.

User response

Remove host mappings or I/O group access for one or more volumes to free up at least one SCSI LUN ID across all access I/O groups, and then retry the command.

Alternately, specify the **-allowmismatchedscsiids** option to allocate SCSI LUN IDs that are not the same in all the access I/O groups. Not all hosts support this configuration.

CMMVC9405E **The command failed as the -allowmismatchedscsiids parameter is not supported for hosts with protocol set as NVMe.**

Explanation

This error can occur when the **-allowmismatchedscsiids** option is specified for a volume that is mapped to a host that uses the Non-Volatile Memory express (NVMe) protocol. The **-allowmismatchedscsiids** option is not supported for hosts that use the NVMe protocol.

User response

Retry the command but do not specify the **-allowmismatchedscsiids** option.

CMMVC9406E **The FlashCopy mapping or consistency group could not be started as a target volume has received I/O within the defined volume protection period.**

Explanation

This error can occur when an attempt is made to start a FlashCopy mapping or consistency group when volume protection is enabled, and the protected volume received I/O operations within the defined volume protection time period.

User response

Ensure that no host I/O operations are being sent to the volume and ensure that you waited for the defined volume protection time period, or disable volume protection.

CMMVC9407E **The Remote Copy relationship could not be created as the specified secondary volume has received I/O within the defined volume protection period.**

Explanation

This error can occur when an attempt is made to create a remote copy relationship with a secondary volume when volume protection is enabled, and the protected volume received I/O operations within the defined volume protection time period.

User response

Ensure that no host I/O operations are being sent to the secondary volume and ensure that you waited for the defined volume protection time period, or specify a different secondary volume.

Alternately, temporarily disable volume protection on the secondary system.

CMMVC9408E **The volume snapshot could not be restored as the volume has received I/O within the defined volume protection period.**

Explanation

This error can occur when an attempt is made to restore a snapshot to a volume when volume protection is enabled, and the protected volume

received I/O operations within the defined volume protection time period.

User response

Ensure that no host I/O operations are being sent to the volume and ensure that you waited for the defined volume protection time period.

Alternately, temporarily disable volume protection.

CMMVC9409E The volume caching IO group must be included in the volume access IO group set.

Explanation

This error can occur when the specified volume access I/O group set does not include the volume caching I/O group.

User response

Retry the command, specifying a volume access I/O group set that includes the volume caching I/O group.

CMMVC9410E The new caching IO group for the volume must be in the volume access set.

Explanation

This error can occur when the specified caching I/O group for the volume is not in the volume access I/O group set.

User response

Retry the command, specifying a volume caching I/O group that is in the volume access I/O group set.

CMMVC9411E The volume caching IO group cannot be removed from the volume access IO group set.

Explanation

This error can occur when an attempt is made to remove the volume caching I/O group from the volume access I/O group set. The volume access I/O group set must include the volume caching I/O group.

User response

Retry the command, specifying an I/O group that is not the volume caching I/O group. Alternately, change the volume caching I/O group and then retry the command.

CMMVC9412E The command failed because cloud call home requires an active connection.

Explanation

This error can occur when an attempt is made to enable support assistance, but the cloud call home function is not enabled, the network connection to the cloud service is blocked or unavailable, or the cloud service is unavailable.

User response

Complete the following steps:

1. Use the **lscloudcallhome** command or the management GUI to check that the cloud call home function is enabled.
2. If the cloud call home function is disabled, use the **chcloudcallhome -enable** command or the management GUI to enable the function.
3. If the cloud call home function is enabled, ensure that port 443 is open in any firewalls so that secure HTTP connections can be established between the system and the call home servers in the cloud. On systems that use DNS, allow outbound traffic over port 443 to esupport.ibm.com. On systems that do not use DNS, allow outbound traffic on port 443 to the following IP addresses: 129.42.54.189, 129.42.56.189, and 129.42.60.189.

You can use the **sendcloudcallhome -connectiontest** command to test the connection to the cloud service, and then view the test results by using the **lscloudcallhome** command.

CMMVC9413E Cloud call home failed to send remote support assist information.

Explanation

This error occurs when call home information cannot be sent to the cloud service. Typically this error occurs when network issues exist, or when the cloud service is unavailable.

User response

Complete the following steps:

1. Ensure that port 443 is open in any firewalls so that secure HTTP connections can be established between the system and the call home servers in the cloud. On systems that use DNS, allow outbound traffic over port 443 to esupport.ibm.com. On systems that do not use DNS, allow outbound traffic on port 443

to the following IP addresses: 129.42.54.189, 129.42.56.189, and 129.42.60.189.

2. Use the **sendcloudcallhome -connectiontest** command to test the connection to the cloud service, and then view the test results by using the **lscloudcallhome** command.

CMMVC9420E **The storage pool cannot be deleted because it has an out of sync volume copy in a storage pool with volume protection active, this copy will be deleted and the volume has received I/O within the defined volume protection period.**

Explanation

This error can occur when an attempt is made to delete a storage pool that contains at least one mirrored volume that has an unsynchronized copy in another pool where volume protection is enabled, and the protected volume received I/O operations within the defined volume protection time period.

User response

Ensure that no host I/O operations are being sent to the mirrored volume and ensure that you waited for the defined volume protection time period.

CMMVC9439E **The command failed because the thin provisioned array does not have sufficient free physical capacity to start expansion.**

Explanation

This error can occur when an attempt is made to expand an array that does not have enough physical free capacity to complete the expansion process.

User response

Free up space on the array by moving volumes to other arrays or by deleting volumes, and then retry the command.

CMMVC9440E **The command failed because the snmp engine id format is invalid.**

Explanation

This error can occur when an attempt is made to use a Simple Network Management Protocol (SNMP) engine ID that has an invalid format. SNMP engine IDs are composed of a series of octets (bytes). They must be at least five octets long and not more than 32 octets long. Thus, an SNMP engine ID is an even number of hexadecimal characters, 10 - 64 hexadecimal characters long.

User response

Retry the command, specifying an SNMP engine ID in a valid format.

CMMVC9441E **The command failed because a post-update background task is in progress.**

Explanation

This error can occur during a software update, when the use of some CLI commands is restricted.

User response

Wait for the update process to complete. Typically, the update process takes no longer than 1 hour.

CMMVC9442E **The command failed because the specified volume is not a HyperSwap volume.**

Explanation

This error occurs when an attempt is made to use the **expandvolume** command to expand a non-HyperSwap volume.

User response

Retry the command, specifying a HyperSwap volume. You can use the **expandvdisksize** command to expand a non-HyperSwap volume.

CMMVC9443E **The command failed because the volume is part of a Remote Copy consistency group.**

Explanation

This error can occur when an attempt is made to change the size of a HyperSwap volume while the volume is part of a remote copy consistency group.

User response

Remove the active-active relationship for the volume from the remote copy consistency group, and then retry the command. The active-active relationship can be added back to the consistency group after the volume is expanded.

CMMVC9444E **The command failed because the volume is part of a FlashCopy mapping.**

Explanation

This error can occur when an attempt is made to expand a HyperSwap volume when the volume is part of a FlashCopy mapping. Only the system-controlled FlashCopy mappings for the change volumes are supported when a HyperSwap volume is expanded.

User response

Remove the FlashCopy mappings for the HyperSwap volume, leaving only the system-controlled FlashCopy mappings for the change volumes, and then retry the command.

CMMVC9445E **The command failed because not all copies of the volume are synchronized.**

Explanation

This error can occur when an attempt is made to expand a HyperSwap volume when not all copies of the volume are synchronized and online.

User response

Resolve any unfixed errors. Wait until all copies of the volume are synchronized, or remove unsynchronized copies if they are no longer required, and then retry the command.

CMMVC9446E **The command failed because there is not enough memory available for the Remote Copy feature in I/O group ID [%1].**

Explanation

This error can occur when an attempt is made to expand a volume but insufficient memory is available for remote copy processing.

User response

Use the **chiogrp -feature remote** command to increase the amount of memory that is available for remote copy processing in the I/O group, and then retry the command.

CMMVC9447E **The command failed because there is not enough memory available for the FlashCopy feature in I/O group ID [%1].**

Explanation

This error can occur when an attempt is made to expand a volume but insufficient memory is available for FlashCopy services.

User response

Use the **chiogrp -feature flash** command to increase the amount of memory that is available for FlashCopy services in the I/O group, and then retry the command.

CMMVC9448E **The command failed because there is not enough memory available for the Mirroring feature in I/O group ID [%1].**

Explanation

This error can occur when an attempt is made to expand a volume but insufficient memory is available for volume mirroring operations.

User response

Use the **chiogrp -feature mirror** command to increase the amount of memory that is available for volume mirroring operations in the I/O group, and then retry the command.

CMMVC9449E **The command failed because the volume has one or more fully allocated copies.**

Explanation

This error can occur when an attempt is made to expand a HyperSwap volume that has volume copies that are not thin-provisioned or compressed. HyperSwap volumes can be expanded only if all copies of the volume are thin-provisioned or compressed.

User response

Remove any volume copies that are not thin-provisioned or compressed, and then retry the command.

CMMVC9450E **The command failed because the volume has multiple copies in a site.**

Explanation

This error can occur when an attempt is made to expand a HyperSwap volume that has multiple volume copies that are in a single site. HyperSwap volumes

can be expanded only when there is one copy of the volume in each site.

User response

Remove one of the mirrored volume copies and then retry the command.

CMMVC9451E The command failed because the last grain of a volume copy has not been fully formatted.

Explanation

This error can occur when the volume was created by using a previous version of the system software and the volume has at least one fully allocated copy where the last grain is not fully formatted.

User response

Convert the volume to a basic volume by removing one of the HyperSwap volume copies. Expand the size of the volume and then add a HyperSwap volume copy.

Alternately, convert all fully allocated volume copies of the volume to use thin-provisioning or compression, and then retry the command.

CMMVC9452E The command failed because the volume has one or more image mode copies.

Explanation

This error can occur when the HyperSwap volume has image mode copies.

User response

Convert the volume to a basic volume by removing one of the HyperSwap volume copies. Expand the size of the volume and then add a HyperSwap volume copy.

Alternately, convert all image mode copies of the volume to a different volume type, and then retry the command.

CMMVC9453E The command failed because there are nodes offline or because there are unrecovered Remote Copy relationships or FlashCopy mappings.

Explanation

This error can occur when the HyperSwap volume has unresolved remote copy or FlashCopy errors or a node is offline.

User response

Resolve the errors and ensure that all nodes are online, and then retry the command.

CMMVC9454E The command failed because one or more HyperSwap volumes have a copy or a change volume in this pool that must be removed first.

Explanation

This error can occur when an attempt is made to force the deletion of a storage pool that contains HyperSwap volume copies.

User response

Use the **rmvolumecopy** command to remove all HyperSwap volume copies in the pool, and then retry the command.

CMMVC9455E The command failed because both [-engineid] and [-securityname] are required to change to version 3 security.

Explanation

This error can occur when an attempt is made to change from Simple Network Management Protocol (SNMP) version 2 to SNMP version 3 but the engine ID or security name was not specified. To change to SNMP version 3, both the engine ID and security name must be specified.

User response

Retry the command, specifying both the **-engineID** and **-securityname** parameters.

CMMVC9456E The command failed because both [-authprotocol] and [-authpassphrase] are required to set up snmp authentication.

Explanation

This error can occur when an attempt is made to set up Simple Network Management Protocol (SNMP) authentication but the protocol or passphrase was not specified. To set up authentication, both the protocol and passphrase must be specified.

User response

Retry the command, specifying both the **-authprotocol** and **-authpassphrase** parameters.

CMMVC9457E **The command failed because authentication must be setup if privacy is to be used.**

Explanation

This error can occur when an attempt is made to set up Simple Network Management Protocol (SNMP) privacy but authentication is not configured. Authentication must be configured before SNMP privacy can be set up.

User response

Use the **chsnmpserver -authprotocol** *protocol* **-authpassphrase** *passphrase* command to set up authentication, and then set up privacy.

CMMVC9458E **The command failed because both [-privprotocol] and [-privpassphrase] are required to set up snmp privacy.**

Explanation

This error can occur when an attempt is made to set up Simple Network Management Protocol (SNMP) privacy but the protocol or passphrase was not specified. To set up SNMP privacy, both a protocol and passphrase must be specified.

User response

Retry the command, specifying both the **-privprotocol** and **-privpassphrase** parameters.

CMMVC9459E **The command failed because the total drive count adds more drives than it is supported to add in a single expansion at this distributed RAID array level**

Explanation

A single expansion of a distributed RAID-1 array supports no more than 12 additional drives. A single expansion of other distributed RAID array levels supports no more than 42 additional drives.

User response

Retry the command, specifying an integer for the **-totaldrivecount** parameter that is no greater than the supported number of drives. For distributed RAID 5 or 6, the maximum number of additional drives in a single expansion is 42. For distributed RAID 1, the maximum number of additional drives in a single expansion is 12.

CMMVC9460E **The command failed because the total drive count is less or equal to the current drive count.**

Explanation

This error can occur when an attempt is made to reduce the number of drives or to add 0 drives. To expand an array, drives must be added.

User response

Retry the command, specifying an integer for the **-totaldrivecount** parameter that is greater than the current number of drives.

CMMVC9461E **The command failed because the total rebuild area count is less than the current rebuild area count.**

Explanation

This error can occur when an attempt is made to reduce the number of rebuild areas. Reducing the number of rebuild areas is not supported during array expansion.

User response

Retry the command, specifying an integer for the **-totalrebuildareas** parameter that is equal to or greater than the current number of rebuild areas.

CMMVC9463E **The command failed because the specified driveclass does not meet or exceed the goals of the array. Please select a driveclass with lscompatibledriveclass.**

Explanation

This error can occur when an attempt is made to add drives that are not from the same drive class or a superior drive class.

User response

Retry the command, specifying drives that are from the same drive class as the drives currently in the array, or a superior drive class. You can use the **svcinfo lscompatibledriveclasses** command to display the drive classes that are compatible with a specified drive class.

CMMVC9470E **The command failed because the FlashCopy map concerned is associated with an access point.**

Explanation

This error can occur when you attempt to remove an FlashCopy map that is providing 3-site functionality.

User response

Only a 3-site administrator can manipulate the 3-site objects.

CMMVC9471E **The command failed because remote users cannot specify a password.**

Explanation

This error occurs when an attempt is made to set a password for a remote user.

User response

Create a remote user but do not attempt to specify a password for the remote user.

CMMVC9473E **The command failed because the superuser password must be changed before the system can be configured.**

Explanation

This error occurs when an attempt is made to run an **svctask** command other than **chuser** or **chcurrentuser** when the superuser password is still set to the default value.

User response

Change the default password, and then retry the command. Passwords must contain more than 5 and fewer than 65 ASCII characters.

CMMVC9474E **The command failed because an insecure password was chosen.**

Explanation

This error can occur when an attempt is made to reset the password to the default password, or to use a password that does not comply with the password policy.

User response

Specify a password that complies with the password policy, and is not the default password.

CMMVC9475E **The node cannot be added because it does not support compressed volume copies in a standard storage pool.**

Explanation

This error occurs when an attempt is made to add a node that does not support real-time compression to an I/O group that includes volumes that use real-time compression. Certain node types support compressed volumes only in data reduction pools.

User response

Complete one of the following tasks:

- Retry the command, selecting an I/O group that does not include volumes that use real-time compression in standard storage pools.
- Decompress the compressed volumes in the I/O group, and then retry the command.
- Migrate the compressed volumes to a data reduction pool, and then retry the command.

CMMVC9476E **The command failed because one or more options are not allowed for 3site RC Relationships.**

Explanation

This error occurs when an attempt is made to use the **chrcrelationship** command to remove a relationship from a three-site configuration.

User response

Ensure that the correct relationship is specified, and then retry the command.

CMMVC9477E **The command failed because one or more options are not allowed for 3site RC Consistency Groups.**

Explanation

This error occurs when an attempt is made to use the **chrcrelationship** command to add a relationship to a three-site configuration.

User response

Ensure that the correct relationship is specified, and then retry the command.

CMMVC9479E **All mutually inclusive parameters should be present.**

Explanation

This error occurs when one of the parameters from a group of mutually inclusive parameters is missing.

User response

Retry the command, specifying all required parameters.

CMMVC9480E **Object list must contain exactly one object from each of the three sites.**

Explanation

This error can occur when an incorrect number of objects are specified. Some command parameters require a list of objects that are separated by colons. The list must contain exactly one object from each of the three sites.

User response

Retry the command, specifying a list of three objects separated by colons.

CMMVC9482E **Specified IP address [%s] is not a valid SVC cluster IP address.**

Explanation

This error occurs when an invalid IP address is specified.

User response

Retry the command, specifying a valid IP address of a clustered system.

CMMVC9483E **Specified 3-site remote copy consistency group does not exist.**

Explanation

This error occurs when a specified 3-site remote copy consistency group does not exist.

User response

Retry the command, specifying a valid 3-site remote copy consistency group.

CMMVC9486E **This CLI action is not supported by orchestrator.**

Explanation

This error occurs when an attempt is made to use the **ch3siterconsistgrp** command to make unsupported changes to a remote copy consistency group in a three-site topology.

User response

Retry the command, specifying a supported configuration change.

CMMVC9491E **Specified 2-site remote copy relationship does not exist.**

Explanation

This error occurs when a specified 2-site remote copy relationship does not exist.

User response

Retry the command, specifying a valid 2-site remote copy relationship.

CMMVC9492E **Specified volume name already exists on the auxiliary-far site.**

Explanation

This error occurs when a specified volume name on a command is already present on the auxiliary-far site.

User response

Retry the command, specifying a new volume name or after removing the volume on the auxiliary-far site.

CMMVC9496E **Specified cluster IPs [%1] and [%2] do not belong to same SVC cluster.**

Explanation

This error occurs when the specified IP addresses are not used by the same clustered system.

User response

Retry the command, specifying IP addresses that are both used by the same clustered system.

CMMVC9498E **Specify IP addresses for 3 different SVC clusters.**

Explanation

This error occurs when one or more specified IP addresses are in use by the same clustered system. The list of IP addresses must include an IP address for each of the three sites.

User response

Retry the command, specifying a valid IP address list that includes a different IP address for each clustered system.

CMMVC9499E Specified user is not a valid 3SiteAdmin user.

Explanation

This error occurs when the specified username is not defined on the clustered system or the user account does not have the 3SiteAdmin role.

User response

Retry the command, specifying a username that is defined on each system and that has the 3-site administrator role.

CMMVC9501E 3-site configuration does not exist.

Explanation

This error occurs when an attempt is made to run a command but the 3-site configuration is not defined on the system.

User response

Use the **mk3siteconfig** command to define a 3-site configuration. Then, retry the command.

CMMVC9503E 3-site configuration cannot be removed if it has 3-site consistency groups.

Explanation

This error occurs when an attempt is made to remove a 3-site configuration when 3-site consistency groups still exist.

User response

Remove the 3-site remote copy consistency groups. Then, retry the command.

CMMVC9510E Enter at least one parameter.

Explanation

This error occurs when no parameters are specified for a command that requires at least one parameter.

User response

Retry the command, specifying at least one parameter.

CMMVC9513E Specified cycling period is invalid.

Explanation

This error occurs when the replication cycling period for the remote site is invalid or is not specified.

User response

Retry the command, specifying a valid replication cycling period for the remote site. The **cycleperiod** parameter must be greater than 300 seconds and less than 86400 seconds.

CMMVC9514E A partially converted relationship already exists in 3-site consistency group.

Explanation

This error occurs when a 3-site consistency group already contains a partially converted relationship.

User response

Complete the conversion of the partially converted relationship. Then, retry the command.

CMMVC9515E Specified periodicsrc or cycleperiod does not match.

Explanation

This error occurs when the specified cycling period or periodic source does not match the value that was specified when the first 2-site consistency relationship was converted to a 3-site consistency relationship.

User response

Retry the command, specifying a cycling period and a periodic source that matches the target 3-site consistency group.

CMMVC9516E Invalid periodic source or target specified.

Explanation

This error occurs when a specified replication source or target is invalid or not part of a 3-site remote copy consistency group or 3-site configuration.

User response

Retry the command, specifying a valid replication source or target.

CMMVC9518E **State or type of specified 3-site consistency group is not valid.**

Explanation

This error occurs when the state or the type of the specified 3-site consistency group is not valid to exclude for maintenance or to rejoin after maintenance. The group must be in the `stopped` or `3site_consistent` or `3site_inconsistent` state to exclude for maintenance, and in the `stopped` or `2site_periodic` state to rejoin after maintenance.

User response

Retry the command, specifying a valid 3-site consistency group.

CMMVC9519E **Specified relationship must be part of a remote copy consistency group.**

Explanation

This error occurs when an attempt is made to convert a stand-alone relationship for use in a 3-site consistency group. Only 2-site remote copy relationships can be converted for use in a 3-site consistency group.

User response

Add a relationship to the remote copy consistency group so that it is a 2-site remote copy relationship. Then, retry the command.

CMMVC9520E **Specified 3-site remote copy consistency group is in partial state.**

Explanation

This error occurs when an attempt is made to start or stop 3-site data cycling when the 3-site configuration is not complete. All of the 2-site consistency groups are not yet converted to 3-site consistency groups.

User response

Complete the 3-site configuration. Then, retry the command.

CMMVC9521E **3-site consistency group is processing another configuration change command.**

Explanation

This error occurs when a configuration change request for the specified consistency group is already being processed.

User response

Wait for the current operation to complete before requesting a new configuration change.

CMMVC9522E **3-site remote copy consistency group was partially converted.**

Explanation

This error occurs when a 3-site remote copy consistency group is only partially converted. This error can occur because of connectivity issues or because of other errors on the clustered systems.

User response

Verify the Secure Shell (SSH) connection from the host to all three sites. Then, retry the command.

CMMVC9523E **State or type of specified relationship is not valid.**

Explanation

This error occurs when an attempt is made to convert an invalid type of 2-site relationship for use in a 3-site remote copy consistency group. The 2-site relationship must be in the `consistent_synchronized` state and must be a Metro Mirror relationship with no change volumes configured.

In a HyperSwap configuration, this error occurs when an attempt is made to convert an invalid type of 2-site relationship for use in a 3-site remote copy consistency group. The 2-site relationship must be in the `consistent_synchronized` state and must be an active-active relationship.

User response

Retry the command, specifying a valid 2-site relationship.

CMMVC9524E **Another relationship is being added to 3-site consistency group.**

Explanation

This error occurs when an attempt is made to add a relationship to a 3-site consistency group when another relationship is already being added to the consistency group.

User response

Wait for the previous relationship to join the consistency group. Then, retry the command.

CMMVC9525E **Specified 2-site remote copy relationship is not a standalone relationship.**

Explanation

This error occurs when an attempt is made to add a 2-site remote copy relationship that is already part of a consistency group to a 3-site consistency group.

User response

Retry the command, specifying a 2-site stand-alone remote copy relationship.

CMMVC9526E **Specified relationship attributes do not match with target 3-site consistency group.**

Explanation

This error occurs when attributes of the specified relationship, such as the master system and auxiliary system, do not match the attributes for the target consistency group.

User response

Retry the command, specifying a relationship with attributes that match the attributes of the target 3-site consistency group.

CMMVC9527E **Specified relationship must be present between near sites in 3-site configuration.**

Explanation

This error occurs when a specified 2-site relationship does not exist between two near sites.

User response

Retry the command, specifying a 2-site relationship that is between two near sites.

CMMVC9528E **Specified site is not a suitable candidate.**

Explanation

This error occurs when a change to a specified site is not valid for the 3-site configuration.

User response

Retry the command, specifying a valid site.

CMMVC9529E **Specified relationship is already being added to given 3-site consistency group.**

Explanation

This error occurs when an attempt is made to add a relationship while that relationship is already being added to the specified 3-site consistency group.

User response

No response is needed. Wait for the relationship to be added to the 3-site consistency group.

CMMVC9546E **The host port cannot be added as a host cannot have both iSCSI names and WWPNs.**

Explanation

This error occurs when an attempt is made to add an Internet Small Computer System Interface (iSCSI) port to a host that already has Fibre Channel or SAS ports, or to add Fibre Channel or SAS ports to a host that already has iSCSI or iSER ports. A host cannot be used for svc_mkarrayat the same time as Fibre Channel or SAS protocols.

User response

Use the **mkhost** command to create another host object, and then add the new host ports to the new host object.

CMMVC9568E **The command failed because only `svctask applysoftware -abort` is allowed.**

Explanation

This error can occur when an attempt is made to resume an upgrade that stalled during the committing phase. When an upgrade stalls in the committing phase, the update must be stopped, and the system must be returned to the original code level.

User response

Run **`applysoftware -abort`** command, and then retry the upgrade.

CMMVC9571E **Two channel start failed because of interruption.**

Explanation

This error occurs when connectivity between nodes is interrupted while a two channel start command is active. The interruption does not necessarily mean that overall connectivity changes.

User response

Prepare the three site consistency group, and then retry the command.

CMMVC9573E **Cannot add accesspoint to HyperSwap volume because change volumes haven't been configured.**

Explanation

This error occurs when an attempt is made to add an access point to a HyperSwap volume when change volumes are not configured on the HyperSwap volume.

User response

Add change volumes to the HyperSwap volume, and then retry the command.

CMMVC9574E **This command failed because the links that are specified in the partnership are not valid. .**

Explanation

For the IP partnership, the mapped portset type must match. For example, only replication portset types are mapped to the IP partnership.

User response

Speak to the system administrator.

CMMVC9575E **The command failed because [-encrypt] is not applicable for a data reduction child pool**

Explanation

This error occurs when an attempt is made to specify the **-encrypt** parameter when a data reduction child pool is created. A data reduction child pool must inherit the encryption status of the parent pool.

User response

Do not use the **-encrypt** parameter when you create a data reduction child pool.

CMMVC9576E **Cannot create a child pool of a different type to the parent pool**

Explanation

This error occurs when an attempt is made to create a data reduction child pool in a regular parent pool or a regular child pool in a data reduction parent pool. Data reduction child pools can be created only in data reduction parent pools. Regular child pools can be created only in regular parent pools.

User response

Create a child pool that matches the type of the parent pool, or create a child pool in a different parent pool.

CMMVC9577E **The command failed because the argument [-noquota] is not applicable**

Explanation

This error occurs when an attempt is made to use the **-noquota** parameter when a regular child pool is created. The **-noquota** parameter can be used only when a data reduction child pool is created.

User response

Use the **-size** parameter to specify the capacity of a regular child pool.

CMMVC9578E **The argument [-size] is not applicable for a data reduction child pool**

Explanation

This error occurs when an attempt is made to use the **-size** parameter when a data reduction child pool is created or changed. The **-size** parameter can be used only when a regular child pool is created or changed.

User response

Use the **-noquota** parameter when you create a data reduction child pool that has no capacity quota.

CMMVC9579E **Cannot remove a data reduction child pool with associated volumes.**

Explanation

This error occurs when an attempt is made to remove a data reduction child pool when the pool still contains volumes or volume copies.

User response

Remove all associated volumes and volume copies from the data reduction child pool, and then use the **rmmdiskgrp** command without the **-force** parameter. Do not use the **-force** parameter with the **rmmdiskgrp** command to remove a data reduction child pool.

CMMVC9580E **The command failed because the password does not meet the minimum length requirement.**

Explanation

This error occurs when an attempt is made to use a password that is shorter than the minimum length requirement.

User response

Retry the command, entering a password that meets the minimum length requirement. You can use the **lssecurity** command to find the minimum password length for the system.

CMMVC9581E **The command failed because strip value 128 is no longer supported, strip value 256 should be used instead.**

Explanation

This error occurs when an attempt is made to use a strip size of 128 KiB, which is no longer supported.

User response

Create an array that uses a strip size of 256 KiB.

CMMVC9583E **The command failed because a user cannot lock their own account.**

Explanation

This error occurs when an attempt is made to run the **chuser -lock** command on the same account where the user is logged in.

User response

Retry the command, specifying a different account to lock.

CMMVC9584E **The command failed because the specified user cannot be locked.**

Explanation

This error occurs when an attempt is made to run the **chuser -lock** command on an account that cannot be locked. Remote user accounts cannot be locked, and locking the superuser account can be disabled.

User response

Retry the command, specifying a different account to lock.

CMMVC9585E **The command failed because the identical link portset IDs or portset names are specified for a partnership.**

Explanation

The error occurred when the same portset IDs or the names are used for WAN link1 and WAN link2.

User response

Specify different portset IDs for both the links, and resubmit the command.

CMMVC9587E **The command failed because a host can only be added to a host cluster with the same ownership group.**

Explanation

This error occurs when an attempt is made to add a host to a host cluster that is in a different ownership group. Hosts can be added to host clusters only if they are in the same ownership group.

User response

Change the ownership group of the host to match the host cluster, and then retry the command. You might need to contact a another user who has the Security Administrator role to change the ownership group.

CMMVC9588E **The command failed because a parent pool cannot be assigned an ownership group.**

Explanation

This error occurs when an attempt is made to assign an ownership group to a parent storage pool. Parent storage pools, as opposed to child pools, cannot be assigned to ownership groups.

User response

Create a child pool in the parent storage pool, and then retry the command on the child pool.

CMMVC9589E **Invalid flag, must be 3SiteAdmin user**

Explanation

This error occurs when an attempt is made to specify a parameter that can only be used by the 3-site administrator.

User response

No response is needed. Alternately, log in as the 3-site administrator to use the parameter.

CMMVC9590E **The command failed because the argument [-warning] is not applicable for a quotaless child pool.**

Explanation

This error occurs when an attempt is made to set the warning level for a quotaless child pool. Quotaless child pools inherit the warning level from the parent pool.

User response

Change the warning level on the parent pool.

CMMVC9591E **The command failed because the partnership between the specified systems does not support high availability.**

Explanation:
The partnership between the specified systems does not support high availability. Both systems must be running a code version and hardware type that supports high availability between systems and the partnership type must be high bandwidth with a low round trip time.

User response:
Recreate the partnership type to one that meets the requirements.

CMMVC9597E **The command failed because the invalid portset ID or the portset name was specified for the IP partnership.**

Explanation

Invalid portset ID or the portset name is specified that to be used by the partnership.

User response

Correct the portset ID and the name, and resubmit the command.

CMMVC9598E **The command failed because an both the portsets cannot be removed from the partnership.**

Explanation:
The **chpartnership** command modifies a partnership that is created over IP links. The IP partnership must be associated with at least one port set ID or name, so both *link1* and *link2* cannot be deleted.

User response:
Do not delete both *link1* and *link2* for an IP Partnership.

CMMVC9601E **The command failed because volume has image mode copy.**

Explanation

This error occurs when an attempt is made to resize a volume that has an image mode copy.

User response

Check that the volume does not have an image mode copy. If the volume has an image mode copy, delete the image mode volume copy and then retry the command.

CMMVC9602E **The command failed because volume has formatting corrupt copy.**

Explanation

This error occurs when an attempt is made to resize a volume that has a copy that is corrupted and is being reformatted as part of a recovery process.

User response

Wait until the volume copy is recovered, and then retry the command.

CMMVC9603E **The command failed because unable to resize the bitmap for the volume.**

Explanation

This error occurs when the bitmap space for the volume cannot be resized.

User response

Check that the volume does not have an image mode copy. If the volume has an image mode copy, delete the image mode volume copy and then retry the command.

CMMVC9604E **The command failed because the current user's password has expired and must be changed.**

Explanation

The password expired for the current user. All actions on the clustered system are blocked until the password is changed.

User response

Use the management GUI or the **chcurrentuser -password** command to change the password for the current user.

CMMVC9605E **The command failed because superuser locking is only supported on platforms with a dedicated technician port.**

Explanation

This error occurs when an attempt is made to enable locking the superuser account on a system that does not have a dedicated technician port. The superuser account cannot be locked on a system that does not have a dedicated technician port because allowing the superuser to be locked out might result in loss of access to the system to run service or recovery actions.

User response

Ensure that the system has a dedicated technician port, and then retry the command. If the system has a dedicated technician port, contact IBM Support Center.

CMMVC9609E **The command failed because the specified user is locked.**

Explanation

This error occurs when an attempt is made to change attributes of an account that cannot be changed while that user account is locked. For example, a user cannot be changed to use remote authentication while their account is locked. A user cannot be moved to a new user group while their account is locked.

User response

Unlock the user account, and then retry the command.

CMMVC9610E **The command failed because users cannot remove themselves.**

Explanation

This error occurs when an attempt is made to remove the account where the user is logged in.

User response

Ensure that the correct user is logged in and the correct user is specified for removal, and then retry the command.

CMMVC9611E **The command failed because the maximum limit for IP addresses with the gateway per physical port is reached. Use another Ethernet port to configure an IP address with gateway.**

Explanation

Only portset 0 supports a gateway IP address.

User response

Select one of the following steps to resolve the issue:

1. Select portset 0.
2. Remove the gateway parameter.
3. Select a different port ID to configure routable IP address.

CMMVC9612E **The command failed because the port is disabled for the connection type specified by the portset.**

Explanation:
IP address allocation failed because the Ethernet port is not enabled for the connection types that are indicated by the portset.

User response

Select the following steps to resolve the issue:

1. Check with **lsportethernet** if the specified port has enabled portset functions.
2. Enable the portset functions using **chportethernet** command.

CMMVC9613E **Sharing IP addresses across different Ethernet ports are not allowed.**

Explanation

Shared IP addresses must be configured on the same Ethernet port.

User response:
Share the IP addresses that are configured in the same Ethernet port.

CMMVC9614E **Shared IP addresses must have an identical VLAN.**

Explanation

Previous IP address was configured with a different VLAN. The same VLAN must be used.

User response:
Specify identical VLAN ID to share IP addresses.

CMMVC9615E **`-shareip` applies to 2 different portsets. You cannot share IP address within the same portset.**

Explanation

The parameter **`-shareip`** indicates that an IP address is intentionally being shared between multiple portsets. The existing IP address (IPv4 or IPv6) is configured with the same portset. An IP address cannot be shared within the same portset.

User response

Specify a different portset from the existing to share an IP address.

CMMVC9616E **The command failed because the maximum limit of IP addresses on a node is reached. Use another *node_id* for the IP allocation.**

Explanation:
The maximum IP address limit per node is reached.

User response:
Use different node to configure IP addresses.

CMMVC9617E **The command failed because the maximum limit of IP addresses that can be associated to an Ethernet port ID is reached. Use another portset ID for IP allocation.**

Explanation

The maximum number of IP addresses per physical Ethernet port (64) is exceeded.

User response

Select the following steps to resolve the issue:

1. select **`lsip`** parameter to view the IP addresses that are currently assigned to the port.
2. Select a different port ID to create a new IP address or select **`rmip/mkip`** parameter to modify an existing IP address to same triplets.

CMMVC9618E **The command failed because the maximum limit of IP addresses that can be associated to a portset ID is reached. Use another portset ID for IP association.**

Explanation

Maximum four IP addresses can be assigned to a single portset.

User response

Select the following steps to resolve the issue:

1. select **`lsip`** parameter to view the IP addresses that are currently assigned to the portset.
2. Select a different portset ID or a portset name.
3. Select **`rmip`** parameter to modify the IP addresses assigned to the portset.

CMMVC9619E **Mixing IP address types within an IP replication type portset is not allowed.**

Explanation

All the IP addresses within an IP replication type portset must be of identical type. For example, either all are IPv4 type or all are IPv6 type.

User response:
Select either IPv4 or IPv6 version within an IP replication type portset.

CMMVC9620E **The command failed because multiple IP addresses from a port are not allowed with the same portset.**

Explanation

Only one IP address can be created for a node, port, or a portset.

User response

Select the following steps to resolve the issue:

1. Use **`lsip`** parameter to view the existing IP addresses.

2. Select a new node, port or a portset to create a new IP address, or select **rmip/mkip** parameter to modify an existing IP address to same triplets.

CMMVC9623E **The system update test utility has not been run recently. Run the system update test utility using the drive test options and ensure that no fatal drive errors are outstanding before attempting this command again.**

Explanation

This error occurs when an attempt is made to update drive firmware but the software update test utility was not run first.

User response

Run the software update test utility, and then within 2 hours retry the **applydrivesoftware** command. For more information about the software update test utility, see http://www.ibm.com/support/docview.wss?uid=ssg1S4000585.

CMMVC9624E **The system update test utility has found fatal errors related to all the drives selected for upgrade.**

Explanation

This error occurs when an attempt is made to update drive firmware but the drives that were specified for upgrade failed the tests run by the software update test utility.

User response

Check the event log. When the **applydrivesoftware** command is run with the **-all** or **-drive** parameters and drives fail some tests, an event is generated in the event log. To determine which drives are affected, use the **lsdrive** command.

CMMVC9625E **The command failed because the password entered matches a previously used password.**

Explanation

This error occurs when an attempt is made to use a password that was already used. Depending on the system configuration, users might be restricted from reusing the current password and up to 10 of the most recently used passwords for the account.

User response

Retry the command, specifying a password that was not previously used.

CMMVC9626E **The command failed because the password has been recently changed. Users must wait [%1] days between password changes.**

Explanation

This error occurs when an attempt is made to change a password before the minimum allowed number of days elapse. If password history checking is enabled, users must wait a minimum number of days between password changes. The **lssecurity** command displays the minimum number of days in the `min_password_age_days` attribute.

User response

Wait the minimum number of days, or wait until the password expires, and then retry the command.

CMMVC9627E **The command failed because the password history feature is not supported on this platform.**

Explanation

This error occurs when an attempt is made to enable password history checking on a system that does not support it.

User response

No user action is required.

CMMVC9628E **The command failed because superuser locking cannot be enabled while superuser password reset is disabled.**

Explanation

This error occurs when an attempt is made to enable locking of the superuser account while the superuser password reset method is disabled.

User response

Run the **svctask setpwdreset -enable** command to enable the superuser password reset method, and then retry the command.

CMMVC9629E **The command failed because superuser password reset cannot be disabled while superuser locking is enabled.**

Explanation

This error occurs when an attempt is made to disable the superuser password reset method while the ability to lock the superuser password is enabled.

User response

Run the **chsecurity -superuserlocking** command to disable locking of the superuser password, and then retry the command.

CMMVC9630E **The command failed because this IP address is the last live connection to backend storage controller.**

Explanation

When an IP address is unconfigured from a source Ethernet port, the system checks whether any sessions are established from the selected port to any back-end controller. You must remove sessions to the back-end iSCSI controller by using the command before you attempt to unconfigure the port.

User response

Use **rmiscsistorageport** command to disconnect back-end controller before you remove the IP address.

CMMVC9631E **The command failed because last IP address in the host-attach portset is linked with the host object.**

Explanation

The last IP (IPv4 or IPv6) address in a host-attach portset cannot be removed when it is linked with the host object.

User response

To delete the IP address, either use **rmhost** command to delete a host object or use **chhost** command to move the host object to another host-attach portset.

CMMVC9632E **The command failed because an IPv4 or IPv6 address is associated with remote service portset that is used by an active partnership.**

Explanation

Any IP addresses that are assigned to remote copy portset with an existing active IP partnership cannot be removed.

User response:

Stop the active partnership and remove the IP address.

CMMVC9633W **Your password has expired. It must be changed now.**

Explanation

The password expired for the current user and must be changed.

User response

Use the management GUI or the **chcurrentuser -password** command to change the password.

CMMVC9634W **Your password will expire in %1 days.**

Explanation

The password for the current user expires in the specified number of days.

User response

Use the management GUI or the **chcurrentuser -password** command to change the password before the password expires.

CMMVC9635E **The command failed because the number of lower case characters supplied does not meet the system policy.**

Explanation

This error occurs when an attempt is made to use a password where the specified number of lowercase characters does not meet the system policy. Users with the SecurityAdmin role can change the password policy by using the **chsecurity** command. The **lssecurity** command displays the password policy. The specified characters must be in the ASCII character set.

User response

Retry the command, specifying a password that meets the minimum requirements for lowercase letters.

CMMVC9636E **The command failed because the number of upper case characters supplied does not meet the system policy.**

Explanation

This error occurs when an attempt is made to use a password where the specified number of uppercase

characters does not meet the system policy. Users with the SecurityAdmin role can change the password policy by using the **chsecurity** command. The **lssecurity** command displays the password policy. The specified characters must be in the ASCII character set.

User response

Retry the command, specifying a password that meets the minimum requirements for uppercase letters.

CMMVC9637E **The command failed because the number of numeric characters supplied does not meet the system policy.**

Explanation

This error occurs when an attempt is made to use a password where the specified number of numeric characters does not meet the system policy. Users with the SecurityAdmin role can change the password policy by using the **chsecurity** command. The **lssecurity** command displays the password policy. The specified numeric characters must be in the ASCII character set.

User response

Retry the command, specifying a password that meets the minimum requirements for numeric characters.

CMMVC9638E **The command failed because the number of special characters supplied does not meet the system policy.**

Explanation

This error occurs when an attempt is made to use a password where the specified number of special characters (symbols) does not meet the system policy. Users with the SecurityAdmin role can change the password policy by using the **chsecurity** command. The **lssecurity** command displays the password policy. The special characters must be in the ASCII character set.

User response

Retry the command, specifying a password that meets the minimum requirements for special characters.

CMMVC9640E **The command failed because the requested total rebuild areas is not supported for expansion.**

Explanation

This error occurs when an attempt is made to expand a distributed RAID 1 array with a number of rebuild areas other than 1.

User response

Retry the command, specifying the parameter **-totalrebuildareas 1**.

CMMVC9641E **The command failed because the requested total drive count is not supported for expansion.**

Explanation

This error occurs when an attempt is made to expand a distributed RAID 1 array with a total number of drives smaller than 3 or greater than 16.

User response

Retry the command, specifying 3 - 16 total drives for the **-totaldrivecount** parameter.

CMMVC9642E **The command failed because non-distributed RAID arrays are not supported for this IO group.**

Explanation

This error occurs when an attempt is made to create a traditional RAID array on a system that does not support traditional RAID.

User response

Retry the command, specifying distributed RAID instead of traditional RAID.

CMMVC9643E **The command failed because the requested raid level requires a larger extent size.**

Explanation

This error occurs when an attempt is made to create a distributed RAID 1 array in a storage pool where the extent size is less than 1024 MiB (mebibytes).

User response

Retry the command, specifying a storage pool where the extent size is 1024 MiB or greater.

CMMVC9644E **The command failed because provided count of drives is too small.**

Explanation

This error occurs when an attempt is made to create a distributed RAID 1 array that uses only two drives for a drive class that requires at least three drives.

User response

Retry the command, increasing the number of drives or specifying drives from a different drive class.

CMMVC9645E **The command failed because provided count of rebuild areas is too small.**

Explanation

This error occurs when an attempt is made to create a distributed RAID 1 array and the specified number of rebuild areas does not support the specified number of drives.

User response

Retry the command, specifying at least one rebuild area.

CMMVC9646E **The command failed because the capacity of the provided drives is too large for the requested RAID level.**

Explanation

This error occurs when an attempt is made to create a distributed RAID array by using drives with a raw capacity that is too large.

User response

Retry the command, specifying drives with a smaller raw capacity.

CMMVC9647E **The command failed because distributed RAID-1 arrays are not supported for this IO group.**

Explanation

This error occurs when an attempt is made to create a distributed RAID 1 array on a system that does not support distributed RAID 1 arrays.

User response

Retry the command, specifying a traditional RAID 6 array rather than a distributed RAID 1 array.

CMMVC9651E **Cannot activate new hardware configuration because additional memory is installed.**

Explanation

This error occurs when the **chnodehw** command is run after memory is added to a node. The **chnodehw** command is no longer supported for use during memory upgrades.

User response

To add memory to a node, use the **rmnode** command to remove the node from the clustered system. Add the memory, and then use the **addnode** command to add the node back to the clustered system.

CMMVC9652E **The command failed because no recent node statistics file was found.**

Explanation

This error occurs when an attempt is made to run the **svcinfo lsportstats** command to display port statistics, but no node statistics file is available.

User response

Check that statistics collection is enabled by examining the output of the **lssystem** command. The `statistics_status` attribute indicates whether statistics collection is on or `off` and the `statistics_frequency` attribute shows the collection interval in minutes.

If statistics collection is disabled, enable it by using the **startstats -interval** *minutes* command.

If statistics collection is enabled, ensure that the node is online. Statistics are collected only for nodes that are online. You can use the **lsnode** or **lsnodecanister** command to determine whether the node is online.

CMMVC9653E **Cluster topology for specified master is not hyperswap.**

Explanation

This error occurs when an attempt is made to create a 3-site HyperSwap configuration when the specified master system does not have a HyperSwap topology.

User response

Ensure that topology of the master system matches the topology of the specified 3-site configuration.

You can use the **chsystem** command to change the topology of a system.

CMMVC9654E **Shared IP addresses must have identical gateway.**

Explanation

The parameter **-shareip** indicates that an IP address is intentionally being shared between multiple portsets. The same gateway must be used.

User response:
Specify identical gateway to share the IP addresses.

CMMVC9655E **Shared IP addresses must have identical prefix.**

Explanation

The parameter **-shareip** indicates that an IP address is intentionally being shared between multiple portsets. The same prefix must be used.

User response:
Specify identical **-prefix** parameter to share the IP addresses.

CMMVC9656E **The IP address that is not configured in the system cannot be shared.**

Explanation:
The IP address cannot be shared because the IP address is not yet configured in the system.

User response

Remove the **-shareip** parameter and resubmit the command.

CMMVC9657E **The command failed because only a single IP partnership is allowed with the remote-copy systems at software version before 8.4.2.0.**

Explanation

Only a single IP Partnership is allowed with the peer system software version before 8.4.2.0.

User response

Update the peer system to create an IP Partnership with the system software version 8.4.2.0 and later.

CMMVC9658E **The command failed because the VDisk is the source volume of an incremental FlashCopy mapping and the target VDisk of such mapping cannot be smaller than the source.**

Explanation

The VDisk cannot be expanded because it is the source VDisk of an incremental FlashCopy mapping whose target VDisk virtual capacity is smaller than the proposed virtual capacity of the source VDisk.

User response

Select the incremental FlashCopy and expand the virtual capacity of the target VDisk.

CMMVC9659E **The command failed because a non-mirrored volume cannot be expanded with the -nofmtdisk parameter if the volume is the source or target volume of an incremental FlashCopy mapping.**

Explanation

The error occurred because an attempt was made to expand the VDisk that is the source or target VDisk of an incremental FlashCopy mapping and **-nofmtdisk** was specified.

User response

Run the command without **-nofmtdisk** parameter.

CMMVC9660E **One or more of the volumes are the source VDisk of an incremental FlashCopy mapping and the target VDisk of such mapping cannot be smaller than the source VDisk.**

Explanation

The volume cannot be expanded because one of its copies is the source VDisk of an incremental FlashCopy mapping whose target VDisk virtual capacity is smaller than the proposed virtual capacity of the source VDisk.

User response

Select the incremental FlashCopy and expand the virtual capacity of the target VDisk.

CMMVC9661E **The VDisk cannot be expanded because it is associated with a FlashCopy mapping that cannot be expanded.**

Explanation

A VDisk cannot be expanded if it is the source or target of a FlashCopy mapping that is not directly controlled by the user.

User response

Remove the FlashCopy mapping, and resubmit the command.

CMMVC9662E **The command failed because the VDisk is the target volume of more than one incremental FlashCopy mapping.**

Explanation

The error occurred when an attempt was made to expand the VDisk, which is the target volume of two or more incremental FlashCopy mappings.

User response

Reconfigure the system so that the VDisk is the target of at most one incremental FlashCopy mapping.

CMMVC9663E **The command failed because the FlashCopy mapping cannot be started if the source volume and the target volume have different virtual sizes.**

Explanation

The FlashCopy function creates copies of content on a source volume to a target volume when they both have the same virtual capacity.

User response

Expand either the source or the target volume capacity, and resubmit the command.

CMMVC9664E **The FlashCopy consistency group cannot be started because one or more mappings in the group have different virtual capacity.**

Explanation

A FlashCopy consistency group is a container for FlashCopy mappings to manage related copies and ensure consistency. The FlashCopy mappings in the group have different volume sizes.

User response

Using **lsfcmap** to find the FlashCopy mappings in the consistency group that have a size mismatch. Expand either the source or the target volume capacity of the mapping, and resubmit the command.

CMMVC9665E **The action failed because the -activeactive parameter is not supported.**

Explanation

The **mkrcrelationship** command cannot be used to create an active-active relationship.

User response

Use the **mkvolume** command to create a new HyperSwap volume, or use the **addvolumecopy** task to add a HyperSwap copy to an existing volume.

CMMVC9666E **The command failed because the specified parent storage pool already has a Safeguarded backup location.**

Explanation

The error occurred when an attempt was made create a Safeguarded backup location in a parent pool that already has a Safeguarded backup location. Only one Safeguarded backup location can be created in a parent pool.

User response

To create a Safeguarded backup location, assign it to another parent pool that does not have a Safeguarded backup location.

CMMVC9667E **The -owner parameter cannot be specified when a new child pool created into a Safeguarded backup location.**

Explanation

The error occurred when an attempt was made to specify **-owner** parameter value.

User response

To create a Safeguarded backup location, use the **-safeguarded** parameter.

CMMVC9668E **The command has failed because the format of the FQDN or IP address is not valid**

Explanation

This error occurred because the hostname or IP address entered is invalid.

User response

Enter a valid hostname or IP address and retry the command.

CMMVC9669E **The action failed because the virtual disk is a part of a HyperSwap volume.**

Explanation

The **rmvdisk** or the **rmvdiskcopy** command cannot be used to remove a virtual disk that is part of a HyperSwap volume.

User response

Use the **rmvolume** command to remove a HyperSwap volume, or use the **rmvolumecopy** command to remove a HyperSwap volume copy.

CMMVC9670E **A portset must be empty and unmapped from the host or partnership objects before deletion.**

Explanation

The error occurred when an attempt was made to delete a portset that has the IP addresses or the hosts that are attached to it.

User response

Select the following steps to resolve the issue:

1. Use **lsportset** command to find the port and hosts count that are associated with the portset.
2. Use **lshost -filtervalue** *portset_id=<portset>* parameter to list hosts associated with the portset.
3. Use **lsip -filtervalue** *portset_id=<portset>* parameter to view IP addresses associated with the portset.

CMMVC9671E **The command failed because portset of type replication can be owned and created by global owner only.**

Explanation:
IP replication type portset can be created and owned only by the global owner.

User response

Select the one of the following steps to resolve the issue:

1. Log in as a user that is authorized to run the command. Then, resubmit command.
2. Speak to a system administrator.

CMMVC9672E **The command failed because new portset of type storage cannot be created.**

Explanation:
Portset 3 is the default portset that is automatically configured when the system is upgraded or created, and it is reserved for the storage virtualization traffic. Only a single storage portset can exist on the system.

User response

Use pre-defined portset 3 for storage virtualization portset.

CMMVC9673E **The command failed because the specified parameters of portset type storage or replication cannot be modified.**

Explanation:
The error occurred as an attempt was made to change the specific parameters of storage or replication portset.

User response

Create another portset with a different type and the desired ownership group. Add IP addresses into the new portset.

CMMVC9674E **The command failed because the port is in use by one or more portsets of this type.**

Explanation:
The port is in use for specified functions and cannot be disabled by using **chportethernet** command.

User response:
Stop the specified portset functions on that port and retry the same command.

CMMVC9675E **The command failed because not all the volumes in the volume groups belong to parent pools that have a Safeguarded backup location (child pool).**

Explanation

Safeguarded copy requires that all volumes in the volume group belong to the parent pools that contains

a Safeguarded backup location (child pool) to be safeguarded.

User response:
Ensure that all the volumes in the specified volume groups have Safeguard backup location.

CMMVC9677E **The command failed because the selected parameters cannot have their values as 0 if the specified volume groups are Safeguarded backup locations.**

Explanation

This error occurs when either of the **-size** or **-rsize** parameter values are specified as 0. When the specifies pool is a Safeguarded backup location, then these parameters cannot have their values as 0.

User response

Select a nonzero value when you specify the parameters values.

CMMVC9678E **The command failed because the selected parameters cannot be used if the specified volume groups are Safeguarded backup locations.**

Explanation

When the pool is a Safeguarded backup location, then the specified parameters cannot be specified.

User response

Resubmit the command without specifying the parameters that are not allowed.

CMMVC9679E **The command failed because some of the required parameters are not specified when the storage pool is a Safeguarded backup location.**

Explanation:
Parameters such as **-size**,**-preferrednode**, and **-iogrp** must be specified when the specified pool is a Safeguarded backup location.

User response:
Specify the required parameters and resubmit the command.

CMMVC9681E **The command failed because multiple I/O groups cannot be configured when the preferred access node is detected.**

Explanation

If a system is created that is on a single site, each of the detected nodes needs to be configured as part of a single I/O group.

User response:
Specify only one I/O group to the preferred access node.

CMMVC9682E **The command failed because default portsets cannot be removed.**

Explanation:
The system-defined Ethernet or Fibre Channel host attach, storage and management portsets cannot be removed.

User response:
No action needed.

CMMVC9683E **The VDisk cannot be shrunk because it is associated with a FlashCopy mapping that cannot be shrunk.**

Explanation

A VDisk cannot be shrunk if it is the source or target of a FlashCopy mapping that is not directly controlled by the user.

User response

Remove the FlashCopy mapping, and resubmit the **shrinkvdisksize** command.

CMMVC9684E **A VDisk cannot be shrunk if it is the target of a FlashCopy mapping.**

Explanation

A volume cannot be shrunk if it is the target of a FlashCopy mapping. If the volume is a source of a FlashCopy mapping, it can be shrunk provided no FlashCopy mapping for which it is the source has a starting capacity that is greater than the proposed size of the volume. The starting capacity can be found in the **lsfcmap** view.

User response

Remove the FlashCopy mapping, and resubmit the **shrinkvdisksize** command.

CMMVC9685E **The VDisk cannot be shrunk because it has dependent FlashCopy mappings with a starting virtual capacity that is**

larger than the proposed virtual capacity.

Explanation

The volume is a source of the dependent FlashCopy mappings. It can be shrunk provided no FlashCopy mappings for which it is the source has a starting capacity that is greater than the proposed size of the volume. The starting capacity can be found in the **lsfcmap** view.

User response

Use **lsfcmap** command to generate the list of all the dependent FlashCopy mappings for the VDisk. Stop the mappings that have starting capacity greater than proposed virtual capacity.

CMMVC9686E **The pool specified is a Safeguarded backup location. The action requested is not supported with pools of this type**

Explanation:
This error occurs when a command that is not supported for a pool that is a Safeguarded backup location is attempted to run.

User response:
User need to choose a different storage pool and retry the command.

CMMVC9687E **The command failed because the maximum limit of per node IP addresses associated with the replication portset is reached.**

Explanation:
The maximum IP address limit per node to a replication portset is reached.

User response:
Use a different replication portset to assign the IP address on the same node.

CMMVC9688E **The addvolumecopy command cannot be used to add a mirror copy to Safeguarded backup if the topology system is not standard.**

Explanation

The error occurs when the system has HyperSwap or Stretched topology and the specified volume is a Safeguarded backup.

User response

Run the **addvdiskcopy** command. Do not run **addvolumecopy** command.

CMMVC9690E **The volume cannot be expanded as it is in a migration relationship.**

Explanation

The volumes in a migration relationship do not support being expanded.

Note: To expand the volume, either complete the migration or remove the relationship and start the migration again after expanding the volume.

User response

N/A

CMMVC9691E **The relationship type was not changed as migration relationships cannot be changed.**

Explanation

This type of migration relationships cannot be changed.

Note: To change the type, you will need to delete the secondary volume and recreate the relationship.

User response

N/A

CMMVC9692E **The relationship was not added to the consistency group as migration relationships do not support consistency groups.**

Explanation

The migration relationships do not support being added to consistency groups as they are not suitable for disaster recovery scenarios.

User response

N/A

CMMVC9693E **The change volume was not associated as migration relationships do not support change volumes.**

Explanation

The migration relationships do not support change volumes being associated.

User response

N/A

CMMVC9694E **The command failed because a username and password are required.**

Explanation

This error can occur when you enter the username and password that do not exist already in the system.

User response

You must specify both, a username and password.

CMMVC9695E **The command failed because the Safeguarded policy is associated with one or more volume groups.**

Explanation

After the volume group is created, the volume group is assigned to a Safeguarded policy. When the volume groups are assigned to a Safeguarded policy, the Safeguarded policy cannot be removed without specifying **-removefromvolumegroups** parameter.

User response

Select the one of the following steps to resolve the issue:

1. Remove the volume groups associated to a Safeguarded policy by using **-rmvolumegroup** command.
2. Contact someone with the Security Administrator role and have them run the command with **-removefromvolumegroups** parameter value.

CMMVC9696E **The command failed because an IP partnership requires at least one link to be specified.**

Explanation:
This command is invalid because at least one link parameter must be specified in the command.

User response:
Provide at least one portset of type replication to link1 or link2.

CMMVC9697E **The relationship cannot not be created as migration relationships are only allowed to be created to a remote system.**

Explanation

The migration relationship is only allowed to be created if the aux system is in a different system.

Note: To migrate a volume within the same system, use the non-disruptive VDisk move procedure (movevdisk).

User response

Rerun the command using a remote cluster name or ID for the **-cluster** parameter.

CMMVC9698E **The relationship could not be started as migration relationships do not support the -clean flag**

Explanation

The migration relationships are not supported to be started with the -clean flag as a full synchronization is required to perform a migration.

User response

Re-run the command without the -clean flag.

CMMVC9699E **The command failed because the provided portset object is not a type of host attach.**

Explanation:
The error occurred because the specified portset is not of type host.

User response:
Specify the portset of type host and resubmit the command.

CMMVC9700E **The command failed because the argument [-removefromvolumegroups] is not allowed for the current user role.**

Explanation:
The argument [-removefromvolumegroups] is not allowed for the current user role.

User response

Have a user with the appropriate authority i.e. only Security Administrator role user can execute the command or remove the volume group associations and execute the command without the -removefromvolumegroups parameter.

CMMVC9703E **Host and portset object ownership must match. Restricted owner of the host object can be mapped to globally owned portset.**

Explanation

A portset object can belong to one ownership group. Users in an ownership group are restricted to view and

manage portset objects outside their ownership group. Restricted owner of the host object can map the host to a globally owned portset.

User response

Map the host object to the portset with the same ownership ID.

CMMVC9704E **The system is unable to generate encryption keys.**

Explanation

This action failed because the system is unable to generate encryption keys.

User response

Contact IBM support.

CMMVC9706E **The command failed because ODX cannot be set whilst a migration relationship exists.**

Explanation

ODX commands from a host are not allowed while a vdisk is being migrated to another system.

User response

Remove or wait for the migration relationship to be completed before retrying the command.

CMMVC9707E **The command failed because the specified provisioning policy does not exist**

Explanation

The specified provisioning policy does not exist.

User response

Verify the ID or name of the provisioning policy with the **lsprovisioningpolicy** command and reissue the command with the correct provisioning policy.

CMMVC9708E **The command failed because a host using the NVMe protocol was not mapped to a migration relationship volume.**

Explanation

The hosts connected using the NVMe protocol cannot be mapped to volumes in a migration relationship.

User response

Remove or wait for the migration relationship to be completed and re-run the command.

CMMVC9709E **The relationship could not be stopped as migration relationships do not support the -access flag.**

Explanation

The migration relationships are not supported to be stopped with the -access flag.

User response

Re-run the command without the -access flag.

CMMVC9710E **The migration relationship could not be created as the master system has ODX enabled.**

Explanation

ODX is not supported with migration relationships and must be disabled in both the systems.

User response

Disable ODX in the master system.

CMMVC9711E **The migration relationship could not be created as the aux system has ODX enabled.**

Explanation

ODX is not supported with migration relationships and must be disabled in both the systems.

User response

Disable ODX in the aux system.

CMMVC9712E **The migration relationship could not be created as the master volume is using persistent reservations and migration of volumes using persistent reservations is unsupported.**

Explanation

SCSI persistent reservations are not supported with the migration relationships.

Note: You may be able to configure your application not to use SCSI persistent reservations, refer to the host or application vendor's documentation for information on how to change the configuration.

User response

The migration operation cannot be performed.

CMMVC9713E **The migration relationship could not be created as the master volume is mapped to a host using the NVMe protocol.**

Explanation

Volumes using the NVMe protocol are not supported with migration relationships. Only SCSI is supported for migration relationships.

User response

The migration operation cannot be performed.

CMMVC9714E **The command failed because the new portset does not have an overlapping port so this action can cause a host outage. Specify a portset with one or more common ports or turn off VDisk protection to force the change.**

Explanation

Mapping a new portset to a host can cause host outage without an overlapping port. The newly mapped portset must contain at least one port that is overlapping with all the old portsets.

User response

Specify a portset with one or more common ports or turn off the VDisk volume protection to apply the change.

CMMVC9715E **The command failed because the auxiliary volume has host mappings.**

Explanation

The auxiliary volume must not have host mappings before the migration relationship is created as the identity of the volume changes which will likely cause the host to experience issues.

User response

Remove the host mapping for the auxiliary volume to avoid the host issues.

CMMVC9716E **The command failed because the secondary volume was not mapped to an online host.**

Explanation

The switch is only allowed if the secondary volume has mappings, and these mappings have online hosts.

User response

Ensure that the host configuration for the secondary volume is correct and has online hosts.

CMMVC9717E **The command failed because the specified provisioning policy has incompatible properties with the specified properties for the pool.**

Explanation

The properties of the provisioning policy must match with the properties of the pool.

User response

Verify the properties of the pool and assign a provisioning policy that matches the properties of the pool.

CMMVC9718E **The source and the target volumes in a FlashCopy mapping must have the same preferred node.**

Explanation:
A Safeguarded backup must have the same preferred node as its Safeguarded source. The FlashCopy function creates copies of content on a source volume to a target volume in a mapping.

User response

Create a Safeguarded backup in the Safeguarded backup location with the same preferred node as the Safeguarded source. Use the Safeguarded source as the target volume to run **mkfcmap** command to create a new FlashCopy mapping.

CMMVC9719E **The command failed because the source and the target volumes in a FlashCopy mapping are not in the same parent storage pool.**

Explanation:
A Safeguarded backup must have the same parent storage pool as its Safeguarded source. The FlashCopy function creates copies of content on a source volume to a target volume in a mapping.

User response

Create a Safeguarded backup in the Safeguarded backup location associated with the Safeguarded source. Use the Safeguarded source as the target

volume to run **mkfcmap** command to create a new FlashCopy mapping.

CMMVC9720E **The command failed because a Safeguarded Policy is not assigned to the source volume.**

Explanation

A Safeguarded policy needs to be assigned to the volume group that has the source volume.

User response

Assign a Safeguarded policy to the volume group where the source volumes are defined.

CMMVC9721E **The command failed because the volume is a target volume of an existing FlashCopy mapping.**

Explanation

A volume can only be a target of a single FlashCopy mapping.

User response

Create a new target volume , and resubmit the **mkfcmap** command.

CMMVC9722E **The command failed because the `-autodelete` parameter cannot be specified when the target volume is a Safeguarded backup.**

Explanation:
When the target volume is a Safeguarded backup, the **-autodelete** parameter cannot be specified.

User response

Resubmit **mkfcmap** command without specifying the **-autodelete** parameter.

CMMVC9724E **The command failed because the `-iogrp` parameter cannot be specified when the target volume is a Safeguarded backup.**

Explanation:
When the target volume is a Safeguarded backup, the **-iogrp** parameter cannot be specified.

User response

Resubmit **mkfcmap** command without specifying the **-iogrp** parameter.

CMMVC9725E **Cannot create a distributed RAID array with this number of member drives for this drive class.**

Explanation:
The drive class does not support the given number of member drives for the specified distributed RAID array level.

User response

Select following steps to resolve the issue:

- Retry the command, specifying a lower value for the **-totaldrivecount** parameter.
- Use a different drive class if available.
- Select a different array level. For example, select RAID-6 instead of RAID-5.

CMMVC9726E **The command failed because migration relationships can not be deleted. Use `rmvolume` to delete the secondary volume instead.**

Explanation

Migration relationships cannot be deleted directly.

User response

To delete the secondary volume, use **rmvolume**

CMMVC9727E **The command failed because the volume is the primary in a migration relationship. The secondary volume must be deleted first.**

Explanation

The primary volume in a migration relationship cannot be deleted while the relationship exists.

Note: The secondary volume must be deleted first.

User response

To delete the secondary volume first, use **rmvolume**.

CMMVC9728E **The command failed because the volume is in a migration relationship and should be deleted using `rmvolume`.**

Explanation

The volumes in a migration relationship cannot be deleted using **rmvdisk**.

User response

To delete the secondary volume instead, use **rmvolume**

CMMVC9729E **The command failed because the managed disk (MDisk) is not compatible with the provisioning policy that is associated with the storage pool.**

Explanation

The provisioning policy has properties that restrict the MDisk to be added to the storage pool. For example, if you are adding an MDisk to a pool that uses a provisioning policy, the capabilities of the MDisk must match the capacity savings method that the provisioning policy defines. An MDisk that uses self-compressing drive technology cannot be added to a data reduction pool that defines a provisioning policy that uses compression and deduplication.

User response

Use a different MDisk or review the provisioning policy that is associated with the storage pool.

CMMVC9730E **The command failed because the specified pools use different provisioning policies.**

Explanation

Using **mkvolume**, you can specify multiple pools. The specified pools must configure the same provisioning policy. For example, if you are adding an MDisk to a pool that uses a provisioning policy, the capabilities of the MDisk must match the capacity savings method that the provisioning policy defines.

User response:
Match the pools or the provisioning policies and run **mkvolume** again.

CMMVC9731E **The migration relationship could not be created because the UUID already exists on the remote cluster.**

Explanation

The UUID of the volume is already present on the aux system which can lead to data corruption if not correctly resolved.

User response

Immediately contact support.

CMMVC9733E **The command failed because a multi-factor authentication service is already enabled.**

Explanation

The user attempted to enable a new multi-factor authentication service, but a service is already enabled on the system. Only one multi-factor authentication service can be enabled at any time.

User response

Disable the current multi-factor authentication service and rerun the command.

CMMVC9732E **The command failed because the portset must be empty and unmapped from host, partnership, or storage objects before deletion.**

Explanation

A portset cannot be deleted if it has any IP addresses, hosts or replication that are associated with it.

User response

Select the following steps to resolve the issue:

1. Use **lsportset** command to find that how many IP addresses and hosts are associated with the portset.
2. Use **lshost -filtervalue portset=<portset>** to list hosts associated with the portset.
3. Use **lspartnership -filtervalue -link1=<portset>** to list partnership associated with the portset.
4. Use **lspartnership -filtervalue -link2=<portset>** to list partnership associated with the portset.
5. Use **lsip -filtervalue portset=<portset>** to view IP addresses associated with the portset.

CMMVC9734E **The command failed because there are users or user groups configured to use the multi-factor authentication service.**

Explanation

This error occurs when an attempt was made to disable the multi-factor authentication service. The multi-factor authentication service cannot be disabled while user groups are configured for multi-factor authentication, or if the superuser is configured for multi-factor authentication.

User response

Disable multi-factor authentication for superuser or any user groups that are configured on the system, then rerun the command.

CMMVC9735E The command failed because the multi-factor authentication credentials have not been set.

Explanation

The error occurred because an attempt was made to enable the multi-factor authentication service, without specifying all of the required credentials for multi-factor authentication.

User response

Set the OpenID Connect client ID, OpenID Connect client secret, CLI client ID, and CLI client secret credentials and rerun the command.

CMMVC9736E The command failed because there was a problem configuring multi-factor authentication on the system.

Explanation

The error occurred when an attempt was made to configure the multi-factor authentication service, but there was a problem applying the new configuration on the system.

User response

Contact IBM support.

CMMVC9737E The command failed because the portset parameter is applicable only for the host with an iSCSI name.

Explanation

A portset can only be associated with an iSCSI host.

User response:
Specify an iSCSI host for the portset and resubmit the command.

CMMVC9738E The command failed because the maximum number of allowed partnerships are exceeded.

Explanation:
Only three partnerships of type Fibre Channel or IP address is permitted in a system.

User response:
Remove the existing partnership to establish the new partnership.

CMMVC9739E Portset type mismatch.

Explanation:
Only a replication type portset can be mapped to a partnership.

User response:
Select replication type portset.

CMMVC9740E The auto-delete option cannot be used in a HyperSwap scenario.

Explanation

An auto-delete option cannot be used to create a volume copy in a HyperSwap Volume.

User response

Try the following actions:

1. Reissue the **addvolumecopy** command using **auto-delete** option.
2. Remove the source volume manually when the command is completed.
3. Contact to the system administrator for further assistance.

CMMVC9741E The command failed because the specified provisioning policy has incompatible properties with the specified properties for the MDisks.

Explanation:
The properties of the provisioning policy must match with the properties of the assigned MDisks.

User response:
Verify the properties of the MDisk and assign a provisioning policy that matches the properties for the MDisk. If the error is not fixed, select an MDisk that matches the properties of provisioning policy.

CMMVC9742E Empty portset is not allowed.

Explanation

Each portset in the partnership must be configured with only one unique IP address (gateway specified). The portset that is specified for *link1* and *link2* must contain at least one IP address.

User response

Map the portset that contain at least one identical IP address types to *link1* and *link2* .

CMMVC9743E **IP addresses cannot be shared across multiple replication type portsets.**

Explanation:
Replication portsets cannot share IP addresses to avoid the same IP address from being used for both link1 and link2.

User response:
Select unique IP addresses for IP replication type portsets.

CMMVC9744E **Mixing IP address types for *link1* and *link2* is not allowed.**

Explanation

All IP addresses in the portsets must be either IPv4 or IPv6 addresses. IP protocol versions cannot be mix on the portsets. The protocol for both the IP partnership and the portset must be the same.

User response

Map portsets that contain identical IP address types to *link1* and *link2* .

CMMVC9745E **The command failed because the multi-factor authentication hostname is not specified.**

Explanation

The error occurred when an attempt was made to enable multi-factor authentication services without specifying the hostname on the system.

User response

Specify the hostname and rerun the command.

CMMVC9746E **The command failed because the volume has a copy in a standard pool and a copy in a data reduction pool.**

Explanation:
Replication can only be enabled on a mirrored volume when both copies are in the same type of storage pool. Either both copies must be in a data reduction pool or in a standard pool.

User response:
Remove or migrate one of the volume copies so that all copies are in the same type of storage pool, then retry the command.

CMMVC9748E **The command failed because a multi-factor authentication service has not been enabled.**

Explanation

The error occurred when an attempt was made to create or modify a user group with multi-factor authentication set. A user group can only use multi-factor authentication if there is a multi-factor authentication service enabled on the system.

User response

Enable a multi-factor authentication service and try again, or create or modify the user group with no multi-factor authentication set.

CMMVC9749E **The command failed because the provisioning policy that is attached to a pool cannot be removed.**

Explanation:
Only the provisioning policy that is not attached to a pool can be removed.

User response

Use the **chmdiskgrp -noprovisioningpolicy** command to remove the provisioning policy from the pool.

CMMVC9750E **Cannot create a secured partnership with remote system running software version prior to 8.5.2.0.**

Explanation:
As older systems do not have IPsec capability, a secure partnership cannot be built with a system with a software version previous to 8.5.2.0.

User response:
To construct a secured partnership, upgrade the peer system to a supported version 8.5.2.0, or set the parameter that is **-secured** no to form an unsecured relationship.

CMMVC9751E **The command failed because multi-factor authentication cannot be enabled for the superuser unless Remote Support Assistance has been enabled.**

Explanation

The user attempted to enable multi-factor authentication for the superuser on IBM Storage Virtualize for Public Cloud, which is only allowed once the Remote Support Assistance feature has been enabled on the system.

User response

Enable the Remote Support Assistance service and rerun the command.

CMMVC9752E **The command failed because Remote Support Assistance cannot be disabled while multi-factor authentication is enabled for the superuser.**

Explanation

The error occurred when an attempt was made to disable Remote Support Assistance on IBM Storage Virtualize for Public Cloud, while multi-factor authentication is enabled for the superuser. Remote Support Assistance must be enabled on IBM Storage Virtualize for Public Cloud when you use multi-factor authentication because it provides a recovery mechanism in case the superuser is locked out.

User response

Disable multi-factor authentication for the superuser and rerun the command.

CMMVC9753E **The command failed because the specified VDisk is used by the system to support volume snapshots.**

Explanation

To add a copy to a volume that is being used to maintain a snapshot, you must use the **addvdiskcopy** command.

User response

Run the **addvdiskcopy** command.

CMMVC9754E **The command failed because a duplicate iSCSI qualified name (IQN) was specified. iSCSI names are case-insensitive.**

Explanation:
The command provided an IQN that is already in use.

User response:
Specify a different iSCSI qualified name (IQN).

CMMVC9755E **The command failed because the snapshot specified does not exist.**

Explanation

The parentuid and snapshot combination that has been entered is invalid or the volumegroup and snapshot combination is invalid or the snapshotid is invalid.

User response

Correct the parameters and resubmit the command.

CMMVC9756E **The command failed because the snapshot specified is not in the ready state.**

Explanation

The snapshot specified is not in the ready state.

User response

Retry the command specifying the correct snapshot

CMMVC9757E **The command failed because there are legacy FlashCopy mappings using one or more of the volumes as a source VDisk.**

Explanation

Adding snapshots to volumes that are source volumes of legacy FlashCopy mappings will result in restricted usability of the legacy FlashCopy mappings.

User response

Retry the command specifying the correct set of volumes or override this check to accepted restricted legacy FlashCopy functionality.

CMMVC9758E **The command failed because there are legacy FlashCopy mappings using one or more of the volumes as a target VDisk.**

Explanation

Snapshots cannot be added to volumes that are target volumes of legacy FlashCopy mappings.

User response

Retry the command specifying the correct set of volumes or remove the legacy FlashCopy mappings before retrying.

CMMVC9759E **The command failed because there are not enough FlashCopy mappings available to create the snapshot objects.**

Explanation

A snapshot requires a FlashCopy mapping object, there will need to be a FlashCopy mapping for each volume in the specified volume group or list of volumes.

User response

Retry the command once FlashCopy mapping objects have been freed for use.

CMMVC9760E **The command failed because the specified pool is not a child pool of the pool containing the volumes.**

Explanation

If the pool option is used the pool must be a child pool of all the volumes in the list or the volume group.

User response

Retry the command specifying the correct pool or without specifying a pool.

CMMVC9761E **The command failed because the specified pool and volume group ownership do not match.**

Explanation

If a pool and a volume group are specified then both must be in the same ownership group, or none must be in an ownership group.

User response

Retry the command after correcting the ownership of the objects.

CMMVC9762E **The command failed because the ownership of the pool of one or more volumes does not match the ownership of the volume group.**

Explanation

The pools of the volumes, as well as the volume group must be in the same ownership group, or none must be in an ownership group.

User response

Retry the command after correcting the ownership of the objects.

CMMVC9763E **The command failed because the volume group specified is empty.**

Explanation

The user has entered a volume group that contains no volumes.

User response

Correct the volume group and rerun the command.

CMMVC9765E **The command failed because there are not enough VDisks available to create the snapshot objects.**

Explanation

A snapshot requires a VDisk object and need to have a VDisk for each volume in the specified volume group or list of volumes.

User response

Retry the command once VDisk objects have been freed for use.

CMMVC9766E **The command failed because either there are no online nodes in the I/O group or because there are unrecovered FlashCopy mappings in the I/O group.**

Explanation

All the nodes in the IO group of the volumes in the volume group are unavailable whilst trying to create snapshot objects.

User response

Retry the command once IO group nodes have recovered.

CMMVC9767E **The command failed because there is not enough FlashCopy bitmap space available to create the snapshots.**

Explanation

A snapshot requires memory for the FlashCopy bitmap, need to have a bitmap associated with each volume in the specified volume group or list of volumes.

User response

Retry the command once memory has been freed for FlashCopy use.

CMMVC9768E The command failed because one or more volume snapshots could not be triggered.

Explanation

One or more of the volume snapshots could not be triggered. This can be due to the parent volumes being offline during the snapshot trigger or the IO group nodes being interrupted during the snapshot trigger.

User response

Retry the command once all the parent volumes and IO group nodes have recovered.

CMMVC9769E The action failed because the specified portset is not valid.

Explanation

The error occurred because the specified portset does not exist.

User response

Run **lsportset** command and check the list of available portsets. Specify the available portset.

CMMVC9770E The action failed because the Fibre Channel port is already added to the portset.

Explanation

The Fibre Channel port is already mapped with the portset. Same Fibre Channel port cannot be mapped multiple times with the same portset.

User response

No user action is required.

CMMVC9771E The action failed because the specified Fibre Channel port is not part of the specified portset.

Explanation

The error occurred because the FC port specified is not added to specified portset. Correct Fibre Channel port should be specified that is added to specified portset.

User response

Check the correct Fibre Channel port association with the specified port by using **lsfcportsetmember** command and rerun the command.

CMMVC9772E The action failed because limit of maximum Fibre Channel port types per node for a portset is reached.

Explanation

A portset limits maximum number of Fibre Channel ports that can be associated with it. This error indicates that the limit has been reached.

User response

Select another portset that has not reached the port limit.

CMMVC9773E The command failed because it could not configure a change volume for replication as the volume does not have enough FlashCopy mappings available.

Explanation:
A change volume for replication requires two FlashCopy mapping objects. The volume is already a member of other FlashCopy mappings and needs to have two FlashCopy mapping objects available for change volume mappings.

User response:
Retry the command once the volume has sufficient FlashCopy mapping objects available.

CMMVC9774E The command failed because there are not enough FlashCopy mappings available to configure a change volume for replication.

Explanation:
Two FlashCopy mapping objects are needed for a change volume for replication. For each volume in the volume group, there must be two FlashCopy mappings available.

User response:
After sufficient FlashCopy mapping objects are available, then retry the command.

CMMVC9775E The command failed because there is not enough FlashCopy bitmap space available to configure a change volume for replication.

Explanation:
To create internal FlashCopy mappings that are used for change recording, FlashCopy bitmap space in needed for each volume in the volume group.

User response:

Update the FlashCopy bitmap space to a sufficient amount and retry the command. For more information, see Bitmap space configuration.

CMMVC9776E The command failed because the backend storage or the replication function is not supported on this Ethernet port.

Explanation

The storage or the replication function is not supported on the specified Ethernet port.

User response

Retry the command with a valid Ethernet port.

CMMVC9778E The command failed because default Fibre Channel portset cannot be removed.

Explanation:
The error occurred because an attempt was made to delete the default Fibre Channel portset. Default portset cannot be removed.

User response

Not applicable.

CMMVC9779E The command failed because portset and Fibre Channel port mapping exist.

Explanation:
The error occurred because an attempt was made to delete a portset that is associated with an Fibre Channel portset.

User response

Run **rmfcportsetmember** command to delete an association between Fibre Channel ports under the portset, then delete the portset.

CMMVC9780E The command failed because the maximum limit of a VLAN on an Ethernet port_id is reached. Use another port_id for a new VLAN allocation.

Explanation:
The error occurred because the maximum VLAN limit per port is reached.

User response:
Select a different port to assign a new VLAN.

CMMVC9781E The command failed due to portset type mismatch.

Explanation

The error occurred because the portset specified does not have compatible port_type.

User response

Run **lsportset** command to identify portsets with the required port_type.

CMMVC9782E The action failed because it may result in loss of VDisk access.

Explanation

The error occurred because the cluster level vdiskprotection is enabled. There are some active hosts that might lose VDisk access if this mapping is removed.

User response

Contact Admin to disable the vdiskprotection using **chsystem -vdiskprotectionenabled no** command.

CMMVC9783E The action failed because neither storage or replication functionality is supported for Fibre Channel portset.

Explanation

Fibre Channel portset cannot be created for storage or replication. Fibre Channel portset can only be created with type host. No other portset type is supported.

User response

Not applicable.

CMMVC9786E The command failed because the single sign-on authentication service is already enabled. .

Explanation

The error occurred when an attempt was made to enable a new single sign-on multi-factor authentication service that is already enabled on the system. Only one single sign-on authentication service can be enabled at any time.

User response

Disable the current single sign-on multifactor authentication and try again.

CMMVC9788E **The command failed because the single sign-on authentication credentials have not been set.**

Explanation

The error occurred because an attempt was made to enable the single sign-on authentication service, but either the OpenID Connect client ID or the OpenID Connect client secret has not been set. The single sign-on authentication service cannot be enabled without specifying both required credentials.

User response

Set the OpenID Connect client ID, and the OpenID Connect client secret credentials and rerun the command.

CMMVC9789E **The command failed because the OpenID Connect configuration endpoint has not been set.**

Explanation

The error occurred because an attempt was made to enable the single sign-on authentication service without specifying the OpenID Connect configuration endpoint. The single sign-on authentication service cannot be enabled without specifying the .well-known/openid-configuration URL for the single sign-on provider.

User response

Set the OpenID Connect configuration endpoint and try again.

CMMVC9790E **The command failed because there was a problem configuring single sign-on authentication on the system.**

Explanation

The error occurred when an attempt was made to configure the single sign-on authentication service on the system, but there was a problem applying the new configuration on the system.

User response

Contact IBM support.

CMMVC9792E **The command failed because the volume has prepared snapshots.**

Explanation

The vdisk cannot be shrunk because it has prepared snapshots.

User response

Remove the snapshots and then retry.

CMMVC9795E **The command failed because the peer system disagrees with the security status of the partnership**

Explanation:
A secure partnership is attempted with a remote system that has an unsecured partnership with the local system, or vice versa. Ensure that both systems in the collaboration seek to construct a secured or unsecured relationship.

User response:
If the user wants to build a protected IP partnership, the **-secured yes** option must be explicitly supplied on both peer systems.

CMMVC9797E **The command failed because the user group contains users who do not have both a password and an SSH key set.**

Explanation

The error occurred when an attempt was made to modify a user group with password and key required set. When the password and key required feature has been set for a user group, all users in the group require both an SSH key and a password. The user group contains users that do not meet this requirement.

User response

Ensure that all the users in the user group have both an SSH key and a password set, then rerun the command.

CMMVC9798E **The command failed because the user does not have both a password and an SSH key set.**

Explanation

The user attempted to enable the password and key required feature for a user that does not have both a password and an SSH key set.

User response

Ensure that the user has both an SSH key and a password set, then rerun the command.

CMMVC9799E **The command failed because the local system must be specified as one of the systems in the replication policy.**

Explanation:
The system where the command is being submitted must be included in the replication policy.

User response:
Ensure that the local system is included in the replication policy, and retry the command.

CMMVC9801E **The command failed because the specified I/O group on the local system contains no nodes.**

Explanation:
A replication policy can be created only by using I/O groups that contain nodes.

User response:
Identify the I/O groups that contain nodes on the local system, and retry the command.

CMMVC9802E **The command failed because the maximum number of replication policies has been reached on the local system.**

Explanation:
The maximum number of replication policies are defined on the local system.

User response:
Check and remove an existing replication policy from the local system and resubmit the command.

CMMVC9803E **The command failed because a replication policy with the specified name is already in use on the local system.**

Explanation:
A replication policy with the specified name already exists on the local system.

User response:
Ensure that the name for the new policy is not used by an existing policy, and retry the command.

CMMVC9804E **The command failed because the system does not support replication policies.**

Explanation:
One or more nodes in the system does not support replication policies.

User response:
Check the hardware compatibility matrix in the documentation. Remove the nodes that do not support the replication policies from the system, and retry the command.

CMMVC9805E **The command failed because there is insufficient memory available for replication in the specified I/O group on the local system.**

Explanation:
One or more nodes in the specified I/O group does not have enough memory installed to use replication policies.

User response:
Check the hardware compatibility matrix in the documentation and install more memory or replace the nodes in the I/O group.

CMMVC9806E **The command failed because the port does not support the specified connection types.**

Explanation

The error occurred when an attempt was made to run the command on the specified Management Ethernet port.

User response

Retry the command with a valid Ethernet port.

CMMVC9807E **The command failed because there is not enough FlashCopy bitmap space available to create the volume group.**

Explanation

In order to create the internal FlashCopy mappings FlashCopy bitmap space will be needed for each snapshot.

User response

Retry the command once FlashCopy bitmap space has been freed for use.

CMMVC9808E **The command failed because the specified volume group does not exist.**

Explanation:
The user has specified a volumegroup id that does not exist.

User response:
Specify a volume group that exists and rerun the command.

CMMVC9809E **The command failed because there are not enough VDisks**

available to create the volume group.

Explanation

There are not enough VDisk ID available to create this volume group.

User response

Re-try the command after freeing up VDisk objects.

CMMVC9810E **The command failed either because there are no online nodes in the I/O group or because there are unrecovered FlashCopy mappings in the I/O group.**

Explanation

All the nodes in the IO group of the source snapshots are unavailable whilst trying to create FlashCopy objects.

User response

Retry the command once IO group nodes have recovered.

CMMVC9811E **The command failed either because there are no online nodes in the I/O group or because there are unrecovered FlashCopy mappings in the I/O group.**

Explanation:
All the nodes in the IO group of the source snapshots are unavailable whilst trying to create FlashCopy objects.

User response:
Retry the command once IO group nodes have recovered.

CMMVC9812E **The command failed because one or more FlashCopy mappings could not be started.**

Explanation:
One or more of the FlashCopy mappings could not be triggered. This can be due to the parent snapshots being offline or the IO group nodes being interrupted.

User response:
Retry the command once all the snapshots and IO group nodes have recovered.

CMMVC9816E **The command failed because the policy is assigned to a volume group.**

Explanation:
A replication policy can only be removed when it is not assigned to any volume groups.

User response:
Remove the replication policy from the volume groups it is currently assigned to, or assign an alternative replication policy to those groups, and retry the command.

CMMVC9818E **The command failed because superuser password key required cannot be enabled while superuser password reset is disabled.**

Explanation

The error occurred when an attempt was made to enable the password and key required feature when superuser password reset is disabled.

User response

Enable the superuser password reset feature and rerun the command.

CMMVC9819E **The command failed because superuser password reset cannot be disabled while superuser password key required is enabled.**

Explanation

The error occurred when an attempt was made to disable the superuser password reset feature when the superuser password and key required feature is enabled.

User response

Disable the "password and key required" feature for superuser and rerun the command.

CMMVC9820E **The command failed because the proxy server cannot be removed while proxy clients are configured.**

Explanation:
The user attempted to remove the proxy server while MFA or SSO is configured. Proxy server details cannot be removed while MFA or SSO is configured.

User response

Disable the MFA or SSO configuration and try again.

To disable the MFA or SSO use the following commands:

- Use the **chauthmultifactorverify** to disable MFA (IBM security verify).

- Use the **chauthmultifactorduo** to disable MFA (Duo Security).
- Use the **chauthsinglesignon** to disable SSO.
- Use the **chproxy** to change proxy server details if user does not want to disable MFA or SSO configuration.

CMMVC9822E The command failed because the specified replication policy does not exist.

Explanation:
The specified replication policy does not exist.

User response:
Use the **lsreplicationpolicy** command to list the existing replication policies. Retry the command with a valid replication policy.

CMMVC9823E The command failed because the user group cannot have all the login methods disabled.

Explanation

The error occurred when an attempt was made to create or modify a user group with all the login access methods disabled.

User response

Enable at least one login access method to create or modify a user group.

CMMVC9825E The command failed because the local volume group is not a production copy.

Explanation:
The user has attempted a configuration change to a volume group that is operating as a target for replication. Configuration changes are possible only on a production copy of a volume group.

User response:
Identify the system with a production copy of the volume group and apply the configuration change from that system, or enable access to the local volume group, then retry the command.

CMMVC9826E The command failed because the specified remote system is used in a replication policy.

Explanation:
This action is only permitted when there are no replication policies associated with the partnership.

User response:
Remove the replication policies that include the remote system, and retry the command.

CMMVC9827E The command failed because the new replication policy does not have the same topology and locations as the existing policy assigned to the volume group.

Explanation:
The user has attempted to apply a new replication policy to a volume group that already has a replication policy assigned. The differences between the two policies prevent a direct transition from the current policy to the new policy. A replication policy can only be replaced with a new policy if the topology and locations specified in each policy are the same.

User response:
Remove the current replication policy from the volume group before assigning the new policy.

CMMVC9828E The command failed because the shared disk feature is not enabled on the local managed disks.

Explanation

The error occurred when an attempt is made to add local managed disks to a storage pool or an image volume without enabling the shared disk feature on the managed disks.

User response

- Select a local managed disk on which shared disk feature is enabled and then retry the command.
- Enable the shared disk feature on Azure disk by using Azure CLI or Azure web console.

CMMVC9829E The command failed because the volume has replication enabled and the command specified is prohibited for these volumes.

Explanation:
The volume is replicated, and the command that is specified is not allowed on these volumes.

User response:
Disable replication on the volume, and retry the command.

CMMVC9830E The command failed because the volume is not a production or independent copy.

Explanation:
The volume is in a volume group that is configured as a recovery copy. The command that is specified is only permitted for a volume in a production or independent copy of a volume group.

User response:

Complete the action on the system with the production copy of the volume, or enable independent access to this volume group to suspend replication, and permit configuration changes.

CMMVC9831E **The command failed because the specified volume group has an associated replication policy.**

Explanation:
The command specified is prohibited for a volume group with a replication policy.

User response:
Retry the command without specifying a volume group, or specifying a volume group that does not have a replication policy.

CMMVC9832E **The command failed because the specified pool contains volumes in a volume group with an associated replication policy, or volumes used by the system for replication.**

Explanation:
A storage pool cannot be removed if it would result in volumes being deleted from a volume group with a replication policy, or volumes used by the system for replication.

User response:
Migrate the affected volumes to another pool or remove the replication policy from the corresponding volume groups, then retry the command.

CMMVC9833E **The command failed because the volume is in a volume group with a replication policy. Use the rmvolume command to remove the volume.**

Explanation:
The rmvdisk and rmvdiskcopy commands cannot be used to remove a volume from a volume group with a replication policy.

User response:
Use the rmvolume command to remove the volume.

CMMVC9834E **The command failed because either the current or specified volume group has a replication policy assigned.**

Explanation:
The user has attempted to move a volume directly between volume groups when a replication policy is assigned to either or both of the volume groups.

User response:
Move the volume out of the current volume group, then move the volume into the new volume group.

CMMVC9837E **The command failed because the snapshot supplied is not active.**

Explanation:
The snapshot has been prepared, but not triggered.

User response:
Retry once the snapshots has been triggered.

CMMVC9838E **The command failed because maximum number of snapshot objects in the ready state has already been reached.**

Explanation

It is not possible to create any additional snapshot objects with the -prepare option.

User response

Trigger or delete one of the existing snapshots before retrying the command.

CMMVC9841E **The command failed because the specified pool contains snapshot objects or volumes associated with snapshot objects.**

Explanation

All snapshot objects must be deleted before the pool can be deleted.

User response

Remove the snapshots and then retry.

CMMVC9842E **The command failed because the I/O group of one or more volumes in the volume group does not match the I/O group of the replication policy.**

Explanation:
The user has attempted to assign a replication policy to a volume group that contains volumes from I/O groups that do not match the I/O group of the replication policy. For a replication policy to be assigned to a volume group, all the volumes in the group must be from the I/O group defined for the local system in the policy.

User response:
Move all volumes in the volume group into the I/O group specified in the replication policy, or use different volume groups and replication policies for the volumes in each I/O group.

CMMVC9843E **The command failed because the specified volume group contains**

volumes with Transparent Cloud Tiering enabled.

Explanation:
The user attempted to assign a replication policy to a volume group that contains volumes with Transparent Cloud Tiering enabled. If a replication policy need to be assign to a volume group, Transparent Cloud Tearing cannot be enabled for the volumes in that group.

User response:
Disable Transparent Cloud Tiering on all volumes in the volume group, then retry the command.

CMMVC9844E **The command failed because the specified volume group contains volumes that are part of a remote copy relationship.**

Explanation:
For a replication policy to be assigned to a volume group, the volumes in the group cannot be part of a remote copy relationship.

User response:
Remove the remote copy relationships on all volumes in the volume group, then retry the command.

CMMVC9845E **The command failed because the specified volume group contains volumes with caching disabled.**

Explanation:
The user attempted to assign a replication policy to a volume group that contains volumes that have write caching disabled, the volumes in the group must have caching enabled whenever a replication policy to be assigned to a volume group.

User response:
Enable caching on all volumes in the volume group, then retry the command.

CMMVC9846E **The command failed because the specified volume group contains volumes with image mode copies.**

Explanation:
The user attempted to assign a replication policy to a volume group that contains volumes with image mode copies. To assign a replication policy to a volume group, the group cannot contain volumes with image mode copies.

User response:
Convert the image mode volume copies to non-image mode copies, then retry the command.

CMMVC9847E **The command failed because the replication settings could not be applied.**

Explanation:
The system encountered an error while configuring replication.

User response:
Contact the support center.

CMMVC9848E **The command failed because there is not enough bitmap space available in the I/O group to enable replication on the volume.**

Explanation:
The system attempted to enable replication on a volume, it failed as sufficient bitmap space for the replication feature is not available in the I/O group.

User response

To replicate this volume, you need to reduce the amount of replicated capacity within the I/O group or use a different I/O group. You can reduce the amount of replicated capacity by:

- Removing volumes in volume groups that have a replication policy.
- Moving volumes out of volume groups that have a replication policy, so they are no longer replicated.
- Removing the replication policy from volume groups, so they are no longer replicated.

CMMVC9849E **The command failed because a cache disabled volume cannot be created in a volume group with a replication policy.**

Explanation:
A cache disabled volume cannot be created in a volume group with a replication policy.

User response:
Create a cache disabled volume without specifying a volume group, specify a volume group that does not have a replication policy, or create the volume with caching enabled.

CMMVC9850E **The command failed because a HyperSwap volume cannot be created in a volume group with a replication policy.**

Explanation:
A HyperSwap volume cannot be created in a volume group with a replication policy.

User response:
Create the volume without specifying a volume group, specify a volume group that does not have a replication policy, or create a basic volume.

CMMVC9851E **The command failed because there is already a distributed RAID array with compressing drives in the storage pool.**

Explanation

The error occurred because it is not permitted to have more than one distributed RAID array with compressing drives in the same storage pool.

User response

Use the specified drives to expand an existing distributed RAID array or to create a new array in another storage pool.

CMMVC9852E **The command failed because sufficient pool links have not been established to the remote system. The first failing pool is id [%1].**

Explanation:
The user has attempted to assign a replication policy to a volume group, or a storage partition with an associated volumegroup, which contains volumes with copies in pools that are not linked to pools on the remote systems in the replication policy. For a replication policy to be assigned to a volume group or storage partition, the pools on the local system that contain volumes to be replicated must be linked to pools in each of the remote systems in the replication policy.

User response:
Configure pool links between the local and remote pools, then retry the command.

CMMVC9853E **The command failed because there is insufficient space to store the certificate, available space is n KiB.**

Explanation

The error occurred when the user attempted to create a Store with a file of size > 12KiB.

User response

User need to create the Store with fewer certificates, adding the remaining to another Store.

CMMVC9854E **The command failed as the specified parent uid is not valid.**

Explanation

The snapshot object with the specified parent uid does not exist.

User response

Retry the command specifying the correct parent uid.

CMMVC9855E **The command failed because one of more of the specified volumes does not exist.**

Explanation:
The user has specified a volume id that does not exist.

User response:
Specify a volume that exists and rerun the command.

CMMVC9856E **The command failed because a duplicate volume was entered in the list.**

Explanation

The user has entered duplicate volumes in the **-volumes** list.

User response

Remove the duplicates and rerun the command.

CMMVC9857E **The command failed because the volume has replication enabled and the command would change the I/O group**

Explanation:
A volume in a volume group with a replication policy cannot be moved to a different I/O group. The preferred node for the volume can be changed if the new preferred node is in the same I/O group.

User response:
Retry the command by specifying a preferred node in the current I/O group.

CMMVC9858E **The command failed because snapshots cannot be added to mirrored volumes.**

Explanation

At least one of the specified volumes or one volume in the volume group is mirrored. It is not possible to create a snapshot of mirrored volumes.

User response

Remove one copy of the mirrored volumes and retry.

CMMVC9859E **The command failed because a copy cannot be added to a VDisk with snapshots without autodelete.**

Explanation

VDisks with snapshots cannot be mirrored except for migration.

User response

Add the -autodelete flag and retry.

CMMVC9860E	**The command failed because a volume with snapshots cannot be part of a legacy FlashCopy mapping as source VDisk.**

Explanation

Specified volumes that have snapshot cannot be included in legacy FlashCopy mapping as source VDisk.

User response

Remove snapshots from volume, and retry the command.

CMMVC9861E	**The command failed because a volume with snapshots cannot be part of a legacy FlashCopy mapping as target VDisk.**

Explanation

Specified volumes that have snapshot cannot be included in legacy FlashCopy mapping as target VDisk.

User response

Remove snapshots from volume, and retry the command.

CMMVC9862E	**The legacy FlashCopy mapping was not started or prepared as the source VDisk has snapshots.**

Explanation

Mappings that have snapshots cannot be started as FlashCopy mappings.

User response

Remove snapshots from mappings, and retry the command.

CMMVC9863E	**The source or target volume has a type set.**

Explanation

A volume of type thinclone or clone cannot be part of a legacy FlashCopy mapping.

User response

Locate a volume without a type and retry the command.

CMMVC9865E	**The command failed because a management port can only be configured with a management IP address.**

Explanation

Only management IP addresses are supported on the specified port.

User response

Select a different port.

CMMVC9866E	**The command failed as the clone is dependent on a snapshot or volume in a different site**

Explanation

The clone could not be made HyperSwap as it is dependent on a snapshot or parent volume that is not site aligned.

User response

Either wait for the clone to complete copying or attempt to correct alignment, then retry the command.

CMMVC9867E	**The command failed because the topology is HyperSwap, one of the parent volumes is not site aligned and in a data reduction pool**

Explanation

It is not possible to create a non-aligned snapshot vdisk in a data reduction pool if the topology is HyperSwap.

User response

Align or remove the parent volume and retry.

CMMVC9868E	**The command failed because the topology is hyperswap the target iogroup and pool are not site aligned and the target pool is data reduction.**

Explanation

It's not possible to create a non site aligned data reduction volume if the topology is HyperSwap.

User response

Retry with parameters such that the sites are aligned.

CMMVC9869E **The command failed as there is a dependent HyperSwap clone and site alignment would be broken.**

Explanation

There is a dependent HyperSwap clone and so cannot add a vdisk copy to this volume in a different storage site.

User response

Wait for the HyperSwap clone to complete and then retry.

CMMVC9870E **The command failed as the snapshot is still part of a volume group..**

Explanation

The snapshot is still part of a volume group and therefore it's ownership will follow the volume groups.

User response

Modify the volume groups ownership group and the snapshots will automatically change.

CMMVC9871E **The command failed because distributed RAID-5 arrays are not supported for this IO group.**

Explanation

The hardware for this IO group does not support creation of distributed RAID-5 arrays.

User response

Create a distributed RAID-6 array.

CMMVC9872E **The command failed because a copy cannot be added to a VDisk that is used to support a snapshot without autodelete.**

Explanation

VDisks that are supporting snapshots cannot be mirrored except for migration.

User response

Add the -autodelete flag and retry.

CMMVC9873E **The command failed because a snapshot parent volume is not in the pool specified.**

Explanation

A snapshot vdisk copy can only be added to a pool if the snapshot parent volume also has a copy in that pool (or a parent of that pool).

User response

Add a vdisk copy of the snapshot parent volume to the pool specified and then retry.

CMMVC9874E **The command failed because the local volume group is not a recovery copy.**

Explanation:
The user has attempted to enable independent access or start a disaster recovery test on a volume group that is not a recovery copy.

User response:
Identify the system with a recovery copy of the volume group and apply the configuration change from that system.

CMMVC9875E **3-site consistency group has only one relationship, use `convertrcconsistgrp` command.**

Explanation

Only one relationship exists in 3-site consistency group else the command fails.

User response:
Run **`convertrcconsistgrp`** command to convert back the consistency group in only one relationship.

CMMVC9876E **3-site RC relationship was partially converted.**

Explanation

3-site RC relationship was partially converted on only 1 or 2 sites in 3-site configuration due to connection issue or SVC site internal error.

User response

Verify the SSH connection to 3-sites and rerun the command to convert back the relationship.

CMMVC9877E **Operation failed, run `stop3siterconsistgrp` command and retry.**

Explanation

The error occurred because the orientation is not set to default.

User response

Run **stop3siterccconsistgrp** command to set the orientation to default.

CMMVC9878E **Another relationship is being removed from 3-site consistency group.**

Explanation

The error occurred when an attempt was made to remove the specified relationship that cannot be removed from 3-site consistency group because another relationship is already being removed and is in partial state.

User response

Rerun the command after previous operation to remove relationship.

CMMVC9879E **The command failed because the volume group has a replication policy assigned.**

Explanation:
The user has attempted to change the name of a volume group that has a replication policy assigned. This operation is prohibited.

User response:
Remove the replication policy from the volume group, or from the storage partition associated with the volume group, and retry the command.

CMMVC9880E **The command failed because the specified volume is used by the system for replication**

Explanation:
The volume is a system-owned volume for replication and the command specified is prohibited for system-owned volumes.

User response:
Some operations for managing the system's replication volumes are prohibited. Review the documentation to determine which commands and parameters are supported.

CMMVC9887E **The volume group cannot be deleted because it has snapshots.**

Explanation

The volume group specified has snapshots associated with it, these need to be removed before the group can be deleted.

User response

Re-try the command after removing the all of snapshots.

CMMVC9889E **The command failed because snapshots cannot be added to zero size volumes.**

Explanation

At least one of the specified volumes or one volume in the volume group is zero size. It is not possible to create a snapshot of zero size volumes.

User response

Remove any zero size volumes and retry the command.

CMMVC9890E **The command failed because there are legacy FlashCopy mappings using one or more of the volumes as a source VDisk.**

Explanation

Scheduled snapshots cannot be added to volumes that are the source volumes of legacy FlashCopy mappings.

User response

Retry the command specifying the correct set of volumes or remove the legacy FlashCopy mappings before retrying.

CMMVC9891E **The command failed because there are not enough VDisks available to create the snapshots.**

Explanation

Each snapshot requires a VDisk object for every volume in the specified volume group. Multiple concurrent snapshots may be required depending on the policy interval and retention period.

User response

Retry the command once VDisk objects have been freed for use.

CMMVC9892E **The command failed because there is not enough FlashCopy**

bitmap memory available to create the snapshots.

Explanation

The memory required depends on the size of the volumes in the specified volume group. Multiple concurrent snapshots may be required depending on the policy interval and retention period.

User response

Retry the command once memory has been freed for FlashCopy use.

CMMVC9893E **The command failed because there are not enough FlashCopy mappings available to create the snapshot objects.**

Explanation

Each snapshot requires a FlashCopy mapping object for every volume in the specified volume group. Multiple concurrent snapshots may be required depending on the policy interval and retention period.

User response

Retry the command once FlashCopy mapping objects have been freed for use.

CMMVC9894E **The command failed because VDisk is a clone or thin clone volume or is a snapshot VDisk.**

Explanation

The volume is a thinclone or clone volume or a VDisk being used for a snapshot. Therefore the split process would result in one of the copies being corrupt.

User response

Either remove the copy instead of splitting or for a clone wait until the population process is complete before retrying the command.

CMMVC9895E **The command failed because snapshots cannot be added to mirrored volumes.**

Explanation

The volume is mirrored and is being created in a volume group with prepared snapshots. It is not possible to create a snapshot of mirrored volumes.

User response

Create the volume with a single copy or create the volume in a different volume group.

CMMVC9896E **The command failed because the volume is too small to use with the features configured on the specified volume group.**

Explanation:
The user attempted to move a zero-sized volume to a volume group with a replication policy. When a replication policy is assigned to a volume group, all volumes in the group must have a non-zero size.

User response:
Expand the size of the volume, and then retry the command.

CMMVC9897E **The command failed because certuid and ssh keys are mutually exclusive.**

Explanation

This error occurs when a user tries to set both a ssh key and a certuid as both values cannot be set at the same time.

User response

Either you need to create another user or remove the existing ssh key assigned to the user.

CMMVC9898E **The command failed because one or more volumes in the volume group are offline.**

Explanation:
The system cannot complete I/O operations to volumes in a volume group because some or all of the volumes are offline.

User response:
Investigate the cause of offline volumes and fix the problem. You can also delete offline volumes from the volume group.

CMMVC9900E **The command failed because the local volume group is not an independent copy.**

Explanation:
The user has attempted to restart replication with the local volume group as the production copy but the volume group is not independent.

User response:
Enable independent access from the system that has a recovery copy of the local volume group before running the command again on the system that will have the production copy.

CMMVC9901E **The command failed because the partnerships to systems in the replication policy are not fully configured, or the systems are not connected.**

Explanation:
The user attempted to change the replication settings, but either the remote systems cannot be accessed or the partnerships to the systems in the replication policy are not fully configured.

User response:
Check and ensure that the partnerships to all systems in the replication policy are fully configured and connected, then retry the command.

CMMVC9902E **The command failed because the volume group in one or more systems is unhealthy.**

Explanation:
The user has attempted to restart replication, but the copy of the volume group in one or more systems is unhealthy.

User response:
Identify and fix the problem that is causing the volume group to be unhealthy, and retry the command.

CMMVC9903E **The command failed because the I/O group of the volume does not match the I/O group of the replication policy.**

Explanation:
The user attempted to add a new volume to a volume group with a replication policy, or move an existing volume into one. When a replication policy is assigned to a volume group, all volumes in the group must belong to the I/O group indicated in the replication policy for the local system.

User response:
Create a new volume in the I/O group provided in the replication policy, or relocate an existing volume there, or use a different volume group and replication policy for each I/O group.

CMMVC9904E **The command failed because the volume has Transparent Cloud Tiering enabled.**

Explanation:
One or more of the volumes to be refreshed has Transparent Cloud Tiering enabled, this prevents the volume from being refreshed.

User response:
Disable Transparent Cloud Tiering for those volumes that need to be refreshed and retry the command.

CMMVC9905E **The command failed because the volume is part of a Remote Copy relationship.**

Explanation:
The user has attempted to move a volume that is part of a Remote Copy relationship into a volume group with a replication policy. When a replication policy is assigned to a volume group, the volumes in the group cannot be part of a Remote Copy relationship.

User response:
Remove the use of Remote Copy on the volume, then retry the command.

CMMVC9906E **The command failed because the volume has caching disabled.**

Explanation:
The user has attempted to move a volume with write caching disabled into a volume group with a replication policy. When a replication policy is assigned to a volume group the volumes in the group must have write caching enabled.

User response:
Enable the write caching on the volume, then retry the command.

CMMVC9907E **The command failed because the volume has image mode copies.**

Explanation:
The user has attempted to move a volume with image mode copies into a volume group with a replication policy. When a replication policy is assigned to a volume group the volumes in the group cannot have image mode copies.

User response:
Either remove the image mode volume copies or convert the image mode copies to non-image mode copies, then retry the command.

CMMVC9909E **The command failed because a certuid cannot be set for superuser.**

Explanation

A certuid cannot be set for superuser.

User response

Create a new or change an existing local user to set a certuid.

CMMVC9910E **The command failed because snapshots cannot be added to a volume that is used by the system to support an existing snapshot.**

Explanation

At least one of the specified volumes is being used by the system to support an existing snapshot. It is not possible to create a snapshot of internal snapshot volumes.

User response

Remove any internal snapshot volumes from the list and retry the command.

CMMVC9911E The command failed because a snapshots parent volume failed the trigger.

Explanation

The addsnapshot command could not trigger the snapshots because one of the parent volumes could not be prepared for the trigger. This can be because

1. A parent volume is/was offline.
2. A parent volumes IO performance is experiencing significant delay.

User response

Rectify any issues with the parent volumes and retry the command.

CMMVC9912E The command failed because the specified Ethernet port can only be used for system management.

Explanation:
This error occurs when the specified port to create an IP address is a management port. Management ports are only used for system management. Functions such as IP connectivity to hosts and partner systems, iSCSI virtualization, and IP replication are provided by the optional adapters that are installed in the nodes.

User response

Complete the following steps to fix this error:

- Use the **`svcinfo lsportethernet`** command to identify the non-management Ethernet ports in the system.
- Retry the command by specifying a non-management Ethernet port that is suitable for the intended usage.

CMMVC9915E The command failed because the maximum number of I/O group connections that can be configured with replication policies has been reached on the local system.

Explanation:
The local I/O group already has the maximum number of local I/O group to remote I/O group connections. For asynchronous replication, replication policies can be configured so that each I/O group can replicate to a maximum of six I/O groups in partnered systems. For high availability, replication policies can be configured for a single I/O group in the local system to a single I/O group in a partnered system.

User response:
Review the policy-based replication configuration limits. Use an existing replication policy, configure the replication policy between I/O groups that already have a replication policy configured or remove existing replication policies before retrying the command. Retry with one of the I/O groups that already have replication policies defined; remove unused replication policies, or use an existing replication policy instead.

CMMVC9916E The command failed because the volume group is not a production or independent copy.

Explanation:
The user has attempted to create a new volume in a volume group that is configured as a recovery copy. New volumes can only be created in production or independent copies of a volume group.

User response:
Identify the system with the production copy of the volume group and create the volume on that system, or enable independent access to the local volume group, then retry the command.

CMMVC9917E The command failed because the password and key required feature cannot be enabled for superuser on this platform

Explanation

On a platform that does not allow it, the user attempted to enable the "password and key required" option for the superuser. Since password authentication for the superuser is prohibited, this feature is not enabled on systems like IBM Storage Virtualize for Public Cloud.

User response

No action required.

CMMVC9918E The command failed because the maximum number of hyperswap volumes have already been created.

Explanation

There is a limit on the number of hyperswap volumes you can create on a system (refer to the product documentation for that limit). The user has attempted to create more than the permitted amount of hyperswap volumes

User response

Check if there are any existing hyperswap volumes that are no longer being used and can be removed in order to allow for the creation of new hyperswap volumes.

CMMVC9921E **The command failed because there is an existing FlashCopy mapping for the volume using bitmap memory from a different I/O group.**

Explanation:
The I/O group of a change volume that is used for replication must be the same as the I/O group used for any existing FlashCopy mappings.

User response:
Remove and then re-create the FlashCopy mappings in the same I/O group as the volume or relocate the volume to the I/O group that contains the existing FlashCopy mappings.

CMMVC9922E **The command failed because the I/O group of one or more FlashCopy mappings for volumes in the volume group does not match the I/O group of the replication policy.**

Explanation:
The user has attempted to assign a replication policy to a volume group that contains volumes with FlashCopy mappings using bitmap memory from I/O groups that do not match the I/O group of the replication policy. For a replication policy to be assigned to a volume group, any existing FlashCopy mappings for volumes in the group must be from the same I/O group defined for the local system in the policy.

User response:
Remove and then recreate any FlashCopy mappings in the I/O group specified in the replication policy, or move the volumes with FlashCopy mappings that use a different I/O group to a different volume group and assign a replication policy using that I/O group.

CMMVC9924E **The command failed because the ownership of the pool specified does not match the ownership of the specified snapshot.**

Explanation:
If a pool is specified then its ownership must match the ownership of the specified snapshot or both must have no ownership.

User response:
Retry with parameters such that the pool ownership matches.

CMMVC9927E **The command failed because the specified volume group contains volumes of zero size.**

Explanation:
The user attempted to assign a replication policy to a volume group that contains volumes with no capacity, or zero-sized volumes. Replication policies cannot be assigned to a volume group if any of the volumes in the group are zero-sized volumes.

User response:
Remove the volumes with zero size from the volume group, and retry the command.

CMMVC9928E **The command failed because the maximum number of replication policies has been reached on the remote system.**

Explanation

This error occurred because the maximum number of replication policies have been reached on the remote system.

User response

Reuse an existing replication policy or reduce the number of replication policies in use on the remote system and resubmit the command.

CMMVC9929E **The command failed because a replication policy with the specified name is already in use on the remote system.**

Explanation:
The specified name or the default name for the replication policy on the local system exists on the remote system. Names of replication policies must be unique.

User response:
Verify that the name of the new policy is not used by an existing policy on the remote system. Resubmit the command, specifying a unique name for the replication policy.

CMMVC9930E **The command failed because the maximum number of I/O group connections that can be configured with replication**

policies has been reached on the remote system.

Explanation:
The remote I/O group already has the maximum number of I/O group to I/O group connections.

User response:
Review the policy-based replication configuration limits. Use an existing replication policy, configure the replication policy between I/O groups that already have a replication policy configured or remove existing replication policies before retrying the command.

CMMVC9931E The command failed because the I/O group that manages the volume needs to be recovered.

Explanation:
The command cannot be completed on a volume in an I/O group which is either offline or unable to bring the volumes online.

User response:
Contact IBM support center.

CMMVC9932E The command failed because the volume group on the remote system is not an independent copy.

Explanation:
The user attempted to restart replication with the remote volume group as the recovery copy, which does not have independent access enabled.

User response:
Before running the command again on the system that has the production copy, review and fix any unfixed errors on the remote system and ensure that the volume group is independent and the volumes are online.

CMMVC9934E The command failed because one or more volumes have had their size changed since the snapshot was taken.

Explanation:
When restoring an entire volume group from a snapshot the composition of that volume group must match what it was at the time that the snapshot was taken.

User response:
Reverse any changes (volume resizes, volumes added to / moved from the volume group) that were made to the volume group since the snapshot was taken or use the **-volumes** parameter to restore a subset of the volumes in the volume group.

CMMVC9935E The command failed because volumes have been moved out of the volume group since the snapshot was taken.

Explanation:
When restoring an entire volume group from a snapshot the composition of that volume group must match what it was at the time that the snapshot was taken.

User response:
Reverse any changes (volume resizes, volumes added to / moved from the volume group) that were made to the volume group since the snapshot was taken or use the **-volumes** parameter to restore a subset of the volumes in the volume group.

CMMVC9936E The command failed because volumes have been added to the volume group since when the snapshot was taken.

Explanation:
When restoring an entire volume group from a snapshot the composition of that volume group must match what it was at the time that the snapshot was taken.

User response:
Reverse any changes (volume resizes, volumes added to / moved from the volume group) that were made to the volume group since the snapshot was taken or use the **-volumes** parameter to restore a subset of the volumes in the volume group.

CMMVC9937E The command failed because there is an error event against the volume group.

Explanation:
Configuration changes to the volume group are disabled because the specified volume group has an unfixed error event.

User response:
Identify and fix the error event, and retry the command.

CMMVC9938E The command failed because there is an error event against the partnership.

Explanation:
Configuration changes to the partnership are disabled because the specified partnership has an unfixed error event.

User response:
Identify and fix the error event, and retry the command.

CMMVC9939E The command failed because statistics interval time was not 1 or 5 minutes.

Explanation

For Storage Insights Integration support, startstats doesnot allow to set the statistics other than 1 or 5 minute interval while sending data.

User response

Always set the statistics interval to 1 or 5 minutes while sending data to Storage Insights.

CMMVC9940E The command failed because cloud backup of thin cloned volume is not allowed.

Explanation:
A thin cloned volume cannot be enabled for cloud backup.

User response:
No action possible. Enabling cloud backup for thin cloned volumes is not allowed.

CMMVC9941E The command failed because restore target has Snapshots.

Explanation:
Restoring a cloud backup to a volume other than the backup source volume is not allowed if the target volume has snapshots.

User response:
Specify a volume without snapshots or remove all snapshots from the volume specified and attempt a restore.

CMMVC9942E The FlashCopy mapping could not be prepared or started because the target volume has replication enabled.

Explanation:
FlashCopy mappings cannot be initiated if the target volume is in a volume group with a replication policy assigned.

User response

Make any of the following changes, and retry the command.

1. Change the target volume of the mapping.
2. Remove the replication policy from the volume group.
3. Move the volume out of the volume group.

CMMVC9943E The FlashCopy consistency group could not be prepared or started because one or more of the mappings in the group have a target volume with replication enabled.

Explanation:
FlashCopy consistency groups cannot be initiated if one or more of the mappings in the group have a target volume with a replication policy assigned.

User response

Use the following command to find the FlashCopy mappings in the consistency group that have a target volume with replication enabled.

```
lsfcmap
```

Make any of the following changes and retry the command:

1. Change the target volume of mapping.
2. Remove the replication policy from the volume group.
3. Move the volume out of the volume group.

CMMVC9944E The command failed because snapshot cannot be created if restore to the volume is in progress.

Explanation:
Snapshot operation on a volume is not allowed if the volume is being restored from the cloud backup. This restriction prevents the creation of snapshots with inconsistent data.

User response:
Create snapshot after the restore operation is completed on the volume.

CMMVC9946E The command failed because the specified I/O group on the remote system contains no nodes.

Explanation:
The specified I/O group on the remote system does not contain nodes. A replication policy can be created only in an I/O group that contains nodes.

User response:
Specify an I/O group on the remote system that contains nodes in the replication policy, and retry the command. Alternatively you can add nodes to the specified I/O group, and retry the command.

CMMVC9947E The command failed because there is insufficient memory available for replication in the specified I/O group on the remote system.

Explanation:
The remote system does not have enough memory installed for replication policies on at least one node in the designated I/O group.

User response:

Check the hardware compatibility matrix in the documentation. Install additional memory or replace the nodes in the I/O group.

CMMVC9948E The command failed because the requested action took too long to complete.

Explanation:
The command requires that the system makes changes to the replication configuration on one or more partnered systems. These changes took too long to complete, and the command failed.

User response:
Verify and fix any issues that might interrupt IP communication between systems.

CMMVC9949E The command failed because the local system is unable to communicate with the remote system via the Ethernet management network.

Explanation:
The Ethernet management network no longer allows communication between the local system and the remote system.

User response:
Make sure that the local and remote systems are connected over Ethernet. Check the system event log, identify, and solve the unfixed errors. When there are no unfixed errors, retry the command.

CMMVC9950E The command failed because of container quota limitation on the cloud service provider account.

Explanation

This error occurs when the cloud service provider account reaches its maximum container limits. The following container limits are supported by cloud service providers:

- For Amazon S3, a container is called a bucket. By default you can create up to 100 buckets. It can be increased to 1000 buckets by submitting a request to Amazon S3 for a service limit increase.
- For Microsoft Azure, an unlimited number of containers are available according to the current policy.

User response:
Request to increase the container limits or delete the containers that are not required, and then retry the command.

CMMVC9951E The command failed because one or more volumes in the volume group are corrupt.

Explanation:
The user has attempted to resume replication or start a disaster recovery test on a volume group that contains corrupt volumes. The corrupt volumes must be recovered before this action is permitted.

User response

Use the **recovervdiskbyiogrp** or **recovervdiskbysystem** commands to acknowledge the data loss and recover the affected volumes, then retry the command

CMMVC9952E The command failed because automatic renewal is only supported for system-signed certificates.

Explanation:
Automatic renewal is allowed only when a certificate is signed by the system certificate authority.

User response:
Verify and confirm that the current system certificate is signed by the system certificate authority, and retry the command.

CMMVC9953E The command failed because the volume has a system-owned volume used for replication.

Explanation:
Volumes that are associated with existing system-owned volumes cannot be used for replication. This can occur if a replication policy has recently been removed and the associated resources are still being removed by the system.

User response:
Wait for the system to complete removing the previous resources used for replication, then retry.

CMMVC9954E The command failed because one or more volumes in the volume group have a system-owned volume used for replication.

Explanation:
This error occurs when a replication policy is added to a volume group that contains volumes linked with system-owned volumes that are used for replication. This can occur if a replication policy has recently been removed from the volume group and the associated system-owned resources are still being removed.

User response:
Wait for the system to finish removing the previous resources used for replication, then retry.

CMMVC9955E The command failed because a conflicting conditional operation is

currently in progress against the cloud storage resource.

Explanation

The consistency model used by cloud service providers for container or bucket configurations can lead to conflicting processes for a certain amount of time. The following examples list such conflicts:

1. If you delete a bucket or container and immediately list all the buckets or containers, the deleted bucket or container can still appear for a certain amount of time in the list. It takes time to replicate the latest configuration with data transparently stored in multiple locations.
2. Deleting and creating an object that contains same name cannot be processed at the same time. This conflict can also occur if an administrator changes the state of the existing bucket or container on the cloud service provider console. For example, a container can be disabled or archived by the administrator on Microsoft Azure. A bucket can be locked to have write-once-read-many model in place on Amazon S3.

User response

The following action can prevent conflicts:

- Wait for some time after you delete a bucket or container, and then retry the command.
- Verify bucket or container state on cloud service provider console if operation fails again after some reasonable time for using existing bucket or container.
- Verify the resource state for reading and writing permission on buckets and containers. Undo any task made on the cloud service provider console that prevents reading and writing to an existing bucket or container. You can find the bucket or container resource on the console by using the following naming convention:

  For Metadata resource: <prefix>.m1 and Data resource: <prefix>.d1

CMMVC9956E **The command failed because the request to the cloud service provider contains invalid inputs.**

Explanation

The system's attempt to read or write a certain resource failed during communication with the cloud service provider. This error can occur due to a problem on the system or cloud service provider. On the system, this error can occur when an invalid character is specified for the -storageaccountname parameter or when an excessive number of characters are specified for the -containerprefix parameter. On the cloud service provider, this error can occur for following reasons:

1. When a mismatch occurs in the request that is received or the HTTP header is malformed. These conditions can lead to a cyberattack.
2. When an Amazon S3 bucket is moved from one region to another or deleting in one region and creating in another region with the same bucket name.

User response:
Retry the command. If the error persists, examine the event logs on the resources of the cloud service providers and take appropriate action. For more information about enabling and analyzing event logs, see Amazon S3 or Microsoft Azure documentation.

CMMVC9957E **The command failed because a restore operation is active for this volume from the change volume.**

Explanation:
Independent access has previously been enabled to the volume and a restore operation is still active. The volume cannot be split until the restore operation has completed

User response:
Retry the command when the restore operation has completed.

CMMVC9958E **The command failed because of an internal server error at the cloud service provider.**

Explanation:
This error occurs when cloud service provider encounters any unexpected conditions that pause the system to process the request for some time.

User response:
Wait for 120 seconds and then retry the command. If the error still exists, check for outage information on the status page for the cloud system provider.

CMMVC9959E **The command failed because the specified partnership does not support high availability.**

Explanation

The user has attempted to generate a quorum app to be used for high availability with another system but the specified partnership does not support high availability. This could be because the remote system has not been upgraded to a code version that supports this feature or the hardware does not support the feature.

User response

Ensure that both this system and the specified partnered system support policy-based replication high availability.

CMMVC9960E **The command failed because there is no truststore configured for the partner system.**

Explanation

The user has asked to generate a policy-based replication high availability IP quorum app when the system has no truststore configured for that type of partnership.

User response

Configure a truststore with valid certificates from the target system.

CMMVC9961E **The command failed because the snapshotpreserveparent property cannot be turned off while some snapshots do not satisfy the snapshot mirroring rules.**

Explanation

The internal volume mirroring properties of the snapshots do not match their parent volume.

User response

Fix the inconsistent snapshot mirroring configuration and run the command again.

CMMVC9962E **The command failed because there are not enough FlashCopy mappings available to create the volume.**

Explanation:
Creating a volume of type thinclone or clone requires a FlashCopy mapping. There are no FlashCopy mappings available in the system.

User response:
Re-try the command after freeing up FlashCopy mapping objects.

CMMVC9963E **The command failed because there is not enough FlashCopy bitmap space available.**

Explanation:
In order to create the internal FlashCopy mapping FlashCopy bitmap space will be needed.

User response:
Retry the command once FlashCopy bitmap space has been freed for use.

CMMVC9964E **The command failed because the source volume is not present in the specified snapshot.**

Explanation:
The snapshot specified does not include an image of the specified source volume.

User response:
Retry the command specifying the correct snapshot and source volume.

CMMVC9965E **Failed to complete request.**

Explanation:
Failed to run command, as some operation failed to complete.

User response:
Check connectivity with sites and check the orchestrator health, and then retry the command.

CMMVC9966E **Operation cannot be performed with specified standby orchestrator configuration.**

Explanation:
Failed to run command, as standby orchestrator state is not valid.

User response:
Run valid operation based on standby and active orchestrator, and then retry the command.

CMMVC9967E **Operation cannot be performed while standby orchestrator is functionality is running.**

Explanation:
This operation cannot be completed while standby orchestrator is running.

User response:
Stop the standby orchestrator, and then retry the command.

CMMVC9968E **The command failed to generate the certificate because an invalid value was supplied for the Subject Alternative Name field.**

Explanation:
This error occurs when you enter invalid data for the Subject Alternative Name field while generating a new certificate. Enter an IP address, URI, email address, DNS name, or a combination, in the Subject Alternative Name field of the certificate. Each entry must be in the form of a key:value pair, where the key and value are separated using a colon character. The valid keys are IP, DNS, URI, email. Each pair must be separated using a valid delimiter such as a space or comma character.

User response:
Add valid data in subject alternate name field, and then retry **chsystemcert** command.

CMMVC9969E **The cloud service provider is currently unable to process requests due to a temporary overload situation, Please try after some time.**

Explanation
This error occurs when cloud service provider account exceeds the limit that is offered to receive requests. Cloud service providers have the following request limits:

- Amazon S3 supports 3500 request per second to add data and 5500 request per second to retrieve data. This error can also occur when it gets two or more PUT calls on the same object key.
- Microsoft Azure uses the terminology of ingress and egress limit. The default maximum ingress per general-purpose v2 account is 25 Gbps; the default maximum egress for general-purpose v2 account is 120 Gbps. It also supports a limit of 500 requests per second on a single blob. By default, system software does not allow more than 500 requests per second on a single blob.

User response:
Wait for some time and then retry the command. You can also request to increase the limit with your cloud service provider. If the error persists, check your environment to see if two or more entities are accessing the cloud storage at the same time. This situation can result in two PUT requests to the same object key.

CMMVC9970E **Operation cannot be performed as active orchestrator is not configured.**

Explanation:
Standby orchestrator is not initiated on active orchestrator, which failed to run the command on standby orchestrator.

User response:
Start standby orchestrator on active orchestrator, and then retry the command.

CMMVC9971E **The current user is not allowed to change users that are members of SecurityAdmin user groups when two person integrity is enabled.**

Explanation:
This error occurs because you don't have the required elevated role to change users that are members of SecurityAdmin user groups.

User response
To resolve this error, complete these steps:

1. Request a role elevation.
2. Check with a user having Restricted Security Administrator or Security Administrator roles and provide them with the user name and ID of the request created by step 1, and get that request approved.
3. After the request is approved, retry before the temporary access expires.

Related reference
“mktwopersonintegrityrequest” on page 635
Use the **mktwopersonintegrityrequest** command to request a role elevation for yourself. This command requests on behalf of the current user invoking this command.

CMMVC9972E **The current user is not allowed to move the specified user to a SecurityAdmin user group when two person integrity is enabled.**

Explanation:
This error occurs because you don't have the required elevated role to move users to Security Administrator user groups.

User response
To resolve this error, complete these steps:

1. Request a role elevation.
2. Check with a user having Restricted Security Administrator or Security Administrator roles and provide them with the user name and ID of the request created by step 1, and get that request approved.
3. After the request is approved, retry before the temporary access expires.

Related reference
“mktwopersonintegrityrequest” on page 635
Use the **mktwopersonintegrityrequest** command to request a role elevation for yourself. This command requests on behalf of the current user invoking this command.

CMMVC9973E **The command failed because the specified user currently holds a two person integrity request.**

Explanation:
This error occurs when you attempt to change any attribute of the user specified in the command (for example password or user group) and the specified user holds a two person integrity request.

User response

To change any attribute of a user that holds a two person integrity request, you can wait until the role elevation (two person integrity request) expires or cancel it by completing the following steps:

1. Locate the two person integrity request in the view that shows all active two person integrity requests.
2. Cancel the request.
3. After the two person integrity request is canceled or expired, retry the command.

Related reference

"chtwopersonintegrityrequest" on page 633
Use the **chtwopersonintegrityrequest** command to approve or cancel an existing role elevation request. You must have the Restricted Security Administrator role or the Security Administrator role to use this command.

CMMVC9974E The current user is not allowed to change user groups with SecurityAdmin role while two person integrity is enabled.

Explanation:

You must have approved two person integrity request to gain elevated access to make any changes to the user groups with Security Administrator role.

User response

To resolve this error, complete these steps:

1. Request a role elevation.
2. Check with a user having Restricted Security Administrator or Security Administrator roles and provide them with the user name and ID of the request created by step 1, and get that request approved.
3. After the request is approved, retry before the temporary access expires.

Related reference

"mktwopersonintegrityrequest" on page 635
Use the **mktwopersonintegrityrequest** command to request a role elevation for yourself. This command requests on behalf of the current user invoking this command.

CMMVC9975E The current user is not allowed to upgrade the role of the specified user group to SecurityAdmin role while two person integrity is enabled.

Explanation:

This error occurs when you try to upgrade the role of a user group to the Security Administrator role while two person integrity is enabled and you don't have the required elevated role.

User response

To resolve this error, complete these steps:

1. Request a role elevation.
2. Check with a user having Restricted Security Administrator or Security Administrator roles and provide them with the user name and ID of the request created by step 1, and get that request approved.
3. After the request is approved, retry before the temporary access expires.

Related reference

"mktwopersonintegrityrequest" on page 635
Use the **mktwopersonintegrityrequest** command to request a role elevation for yourself. This command requests on behalf of the current user invoking this command.

CMMVC9976E The command failed because the specified role change affects a user group that has at least a user member currently holding a two person integrity request.

Explanation

This error occurs when two person integrity is enabled, and:

- you attempt to change the role of a local user group from Security Administrator role to a lesser role than Security Administrator while the local user group being changed has one or more user members holding two person integrity requests.
- or you attempt to remove a local user group that has one or more user members holding two person integrity requests.

User response

Follow these steps to resolve this error:

1. Note down the user group ID that you tried to change or remove.
2. Identify the user members of the user group.
3. Locate if any of the user member names is found in the view of two person integrity requests. If a user member name shows in a request, complete one of the following actions:
   - Wait until the request that is associated with the user name expires and the user name no longer shows in that view.
   - Or cancel the request.

4. When no user member names are present in the two person integrity requests view, retry the command..

Related reference

“chtwopersonintegrityrequest” on page 633
Use the **chtwopersonintegrityrequest** command to approve or cancel an existing role elevation request. You must have the Restricted Security Administrator role or the Security Administrator role to use this command.

CMMVC9977E The current user is not allowed to remove any user group of SecurityAdmin role while two person integrity is enabled.

Explanation:
This error occurs when two person integrity is enabled and you try to remove a user group of Security Administrator role but you don't have an approved role elevation to complete the action.

User response

To resolve this error, complete these steps:

1. Request a role elevation.
2. Check with a user having Restricted Security Administrator or Security Administrator roles and provide them with the user name and ID of the request created by step 1, and get that request approved.
3. After the request is approved, retry before the temporary access expires.

Related reference

“mktwopersonintegrityrequest” on page 635
Use the **mktwopersonintegrityrequest** command to request a role elevation for yourself. This command requests on behalf of the current user invoking this command.

CMMVC9978E The command failed because there are enabled services that do not support the security level.

Explanation

This error occurs when the specific security level is not supported by some of the enabled services.

For example, this error might occur when the security level specified allows only TLS 1.3, but the services currently enabled does not support TLS 1.3. For more information about the services that do not support TLS 1.3, see Changing security protocol levels .

User response:
Ensure that you limit the maximum security protocol level to 5 to keep compatibility with the protocol level that is supported by the enabled services.

CMMVC9979E The command failed to enable the service because it does not support the current security level.

Explanation:
This error occurs when the service that you are attempting to enable does not support the current protocol security level. For example, this error might occur when the current security level specified allows only TLS 1.3, but the service that you are trying to enable does not support TLS 1.3. For more information about the services that do not support TLS 1.3, see Changing security protocol levels .

User response:
Ensure that you limit the maximum security protocol level to 5 to keep compatibility with the protocol level that is supported by the enabled services.

CMMVC9981E The command failed because the replication policy is a HA policy.

Explanation

The volume group has a HA replication policy assigned. Changing the mode is prohibited for a volume group with a replication policy of this type. It is only permitted for volume groups with an asynchronous replication policy.

User response

Run the command on a volume group with an asynchronous replication policy.

CMMVC9982E The command failed because the specified storage partition does not exist.

Explanation

The specified storage partition does not exist.

User response

Use the **lspartition** command to list the storage partitions that exist. Run the command again with a valid storage partition.

CMMVC9983E The Remote Copy relationship or consistency group cannot replicate in this direction because it makes a volume that is configured with policy-based replication a Remote Copy secondary volume.

Explanation:

You can configure remote copy and policy-based replication on the same volume only if the volume is operating as the production volume for both functions.

User response:
Remove the replication policy that is associated with the volume group and retry the action.

CMMVC9984E **The command failed because the volume is being replicated using policy-based replication.**

Explanation:
You cannot assign a remote-copy change volume when remote copy and policy-based replication are configured on the same volume.

User response:
No action needed.

CMMVC9985E **This volume cannot be moved because it is part of a Remote Copy relationship that does not meet the prerequisites to be moved.**

Explanation

To move volumes in a remote-copy relationship to a different I/O group, the following conditions must be met:

- The relationship must be either a Metro Mirror or Global Mirror relationship.
- The relationship must not be in a consistency group.
- The relationship must be consistent_synchronized.
- The relationship must not have a change volume that is associated with the primary volume.
- The volume that is moved must be the primary of the relationship.

User response:
Change the configuration so that the volume meets all the prerequisites.

CMMVC9986E **The command failed because the volume is part of a Remote Copy relationship that does not meet the prerequisites to use policy-based replication.**

Explanation

To use policy-based replication and remote copy on the same volume, the following prerequisites must be met:

- The volume must be the primary in the remote-copy relationship.
- The remote-copy relationship must be Global Mirror or Metro Mirror.
- The remote-copy relationship must not have a change volume that is configured on the primary volume.

User response:
For more information, see Converting remote copy replication to policy-based replication.

CMMVC9987E **This action is not permitted because the volume group contains volumes using Remote Copy that do not meet the prerequisites to use policy-based replication.**

Explanation

To use policy-based replication and remote copy on the same volume, the following prerequisites must be met:

- The volume must be the primary in the remote-copy relationship.
- The remote-copy relationship must be Global Mirror or Metro Mirror.
- The remote-copy relationship must not have a change volume that is configured on the primary volume.

User response:
For more information, see Converting remote copy replication to policy-based replication. .

CMMVC9988E **This action is not permitted because of the state of Remote Copy relationships on volumes in the volume group on one of the systems.**

Explanation

You can configure remote copy and policy-based replication on the same volume only if the volume is operating as the production volume for both functions.

User response:
On each system, identify the volumes in the volume group that have remote-copy relationship. Remove the relationships and retry the command.

CMMVC9990E **The command failed because one or more of the snapshots of the specified VDisk do not have a copy with a matching parent pool.**

Explanation

To add a volume copy to the volume that contains snapshots, consider the following scenarios:

- The snapshots of the volume must be mirrored to the correct pool.
- The parent pool of at least one volume copy that is used to maintain the snapshot is the same as the parent pool of the copy that is added.

User response:
Ensure that the snapshots are mirrored to the correct pool or run the command again with the correct pool.

CMMVC9991E **This VDisk maintains a snapshot for a parent volume. A copy cannot be removed from a snapshot VDisk if the snapshot would then have fewer copies than the parent volume.**

Explanation:
A volume copy can be removed from a snapshot only when the parent volume has a single copy.

User response:
Remove a copy of the parent volume and run the command again.

CMMVC9992E **This VDisk maintains a snapshot for a parent volume. A copy cannot be removed from a snapshot VDisk if the snapshot would no longer have a copy in the same pool as the parent volume's copy.**

Explanation:
A volume copy can be removed from a snapshot only if the pools of the remaining copies are the same as the pools of the parent volume copies.

User response:
Ensure that the parent volume has a copy in the same pool as the remaining copy of the volume.

CMMVC9993E **The command failed because the host is already in a storage partition.**

Explanation

The host is already in a storage partition.

User response

Remove the host from the storage partition first.

CMMVC9994E **The command failed because the host is mapped to one or more volumes.**

Explanation:
The host is mapped to one or more volumes.

User response:
Remove all host mappings .

CMMVC9995E **The command failed because the maximum number of storage partitions has been reached on the system.**

Explanation

The maximum number of storage partitions are already defined on the system.

User response

Remove an existing storage partition from the system and resubmit the command.

CMMVC9996E **The command failed because the replication policy is not a HA policy.**

Explanation

A storage partition replication policy must be a HA replication policy.

User response

If assigning a replication policy to a storage partition, use only a HA replication policy and resubmit the command.

CMMVC9997E **The command failed because the system does not support storage partitions.**

Explanation

At least one node in the system does not support storage partitions.

User response

Check the hardware compatibility matrix in the documentation. Remove the nodes that do not support the function from the system and try again.

Appendix A. Accessibility features for the system

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

Accessibility features

These are the major accessibility features for the system:

- This product uses standard Windows navigation keys.
- Interfaces are commonly used by screen readers.
- Industry-standard devices, ports, and connectors.

For information about the accessibility features of the online information, see IBM Documentation (https://ibm.biz/BdqxdY).

Keyboard navigation

You can use keys or key combinations for operations and to initiate menu actions that can also be done through mouse actions. You can go to the system online documentation from the keyboard by using the keyboard shortcuts for your browser or screen-reader software. See your browser or screen-reader software Help for a list of keyboard shortcuts that it supports.

IBM and accessibility

See the IBM Human Ability and Accessibility Center for more information about the commitment that IBM has to accessibility.

© Copyright IBM Corp. 2003, 2025

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785

© Copyright IBM Corp. 2003, 2025

US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information at www.ibm.com/legal/copytrade.shtml.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux and the Linux logo is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other product and service names might be trademarks of IBM or other companies.

