Hewlett Packard Enterprise

HPE 3PAR Command Line Interface reference guide

Abstract

This reference describes the HPE 3PAR Command Line Interface (CLI) commands used to administer and maintain the HPE StoreServ storage system.

Part Number: P29135-002
Published: March 2021
Edition: 1

© Copyright 2002, 2021 Hewlett Packard Enterprise Development LP

Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

Acknowledgments

Intel®, Itanium®, Optane™, Pentium®, Xeon®, Intel Inside®, and the Intel Inside logo are trademarks of Intel Corporation in the U.S. and other countries.

Microsoft® and Windows® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

Java® and Oracle® are registered trademarks of Oracle and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

All third-party marks are property of their respective owners.

Revision history

Part Number	Publication date	Edition	Summary of changes
TBD	March 2021	1	First release

Contents

Command Quick Reference 16

Banner commands 17

Certificate commands 18

CIM server commands 19

Disk enclosure management commands 20

Drive cage management 20
Encryption commands 20
Physical disk management commands 21

Domain management commands 22

File persona commands 23

File persona antivirus service 24
File persona archiving 24
File Persona group accounts 25
File Persona NDMP 25
File Persona routes 26
File Persona snapshots 26
File Persona user accounts 27

Health and alert management commands 28

Alerts 28
Alert security levels 28
Events 29
System health 29

Help and utility commands 30

LDAP management commands 31

License management commands 32

Maintenance mode commands 33

Node subsystem management 34
Array and controller node 34
Battery management 34
Controller node EEPROM log 35
Controller node properties 35
Firmware versions 35
Network interface configuration 35
Node date information 36
Node rescue 36
Port information 36
System manager 37

Performance management commands 38
Chunklet statistics 38
Data cache memory statistics 38
Dynamic optimization 39
Link statistics 39
Logical disk statistics 39
Node CPU statistics 39
Physical disk statistics 39
Port statistics 40
Preserved data commands 40
Remote Copy volume statistics 40
System tuner 40
Virtual LUN (export) statistics 41
Virtual volume statistics 41

Replication commands 42
Physical copy commands 42
Remote copy commands 42
Virtual copy commands 45

Security hardening commands 46
Service Processor credential commands 46
Support recovery account commands 46

Service commands 47
Disk enclosure commands 47
Node commands 48
General system maintenance commands 48
QoS commands 48
Software version commands 49

SNMP agent commands 50

Sparing commands 51

SSH access commands 52

Task management commands 53

Task schedule commands 54

User management commands 55

Volume management commands 56

Common Provisioning Group management commands 56
Host management commands 57
Logical disk management commands 58
Space and storage management 58
Template management commands 59
Virtual volume management commands 59
Virtual LUN (export) statistics 61

Web Services API commands 62

Storage management with the HPE HPE 3PAR CLI 63

Backslash characters in filenames for Microsoft Windows environments 63
Command syntax 63
Syntax rules for the CLI 64
Documentation layout for command descriptions 65
Exit status codes for CLI commands 66
Glob-style patterns 66
Global options and environment variables 67
Help commands and topics 67
Numbering formats and values 68
Object sets and domain sets 68
Regular expressions and patterns 69
Supported destination VV name patterns 70
User space provisioning 71

Add Command 73

`addsnmpmgr` 73

Admit Commands 76

`admithw` 76
`admitpd` 77
`admitrcopylink` 78
`admitrcopytarget` 79
`admitrcopyvv` 80
`admitvv` 83

Backup Command 84
backupfsconf 84

Cancel Command 86
canceltask 86

Change Command 87
changedomain 87

Check Commands 88
checkhealth 88
checkld 89
checkpassword 90
checkpd 91
checkport 93
checkrclink 94
checksnmp 96
checkvv 97

CLI Command 99
cli 99

Compact Commands 103
compactcpg 103
compactld 104

Control Commands 106
controlencryption 106
controliscsiport 108
controlmag 110
controlpd 111
controlport 112
controlrecoveryauth 120
controlsecurity 121
controlsr 123

Create Commands 128
createaocfg 128
createcert 129
createcpg 131
createdomain 135
createdomainset 136
createflashcache 137
createfpg 137
createfsgroup 139

createfshare 140
createfsip 146
createfsnap 147
createfspn 148
createfsroute 149
createfstore 150
createfsuser 151
creategroupsv 152
creategroupvvcopy 154
createhost 157
createhostset 159
createmaint 160
creatercopygroup 161
creatercopytarget 163
createsched 164
createsnmpuser 166
createspare 167
createsralertcrit 168
createsv 179
createtemplate 181
createuser 186
createvfs 187
createvlun 190
createvv 192
createvvcopy 196
createvvset 200

Dismiss Commands 202
dismisspd 202
dismissrcopylink 202
dismissrcopytarget 203
dismissrcopyvv 204

Free Command 206
freespace 206

Grow Commands 207
growfpg 207
growvv 208

Hist Commands 209
histch 209
histld 211
histpd 214
histport 219
histrcvv 222
histvlun 225
histvv 228

Import Commands 232
importcert 232
importvv 233

Locate Commands 236
locatecage 236
locatenode 237
locatesys 238

Move Commands 240
movech 240
movechtospare 241
movepd 242
movepdtospare 243
moverelocpd 244
movetodomain 247

Promote Commands 249
promotegroupsv 249
promotesv 250
promotevvcopy 251

Remove Commands 253
removealert 253
removeaocfg 254
removebanner 254
removecert 255
removecorequest 256
removecpg 257
removedomain 258
removedomainset 259
removeflashcache 260
removefpg 260
removefsarchive 261
removefsgroup 263
removefshare 264
removefsip 265
removefsnap 266
removefspn 267
removefsroute 268
removefstore 269
removefsuser 269
removehost 270
removehostset 271
removeld 272
removercopygroup 273
removercopytarget 275

removesched 276
removesnmpmgr 277
removesnmppw 277
removesnmpuser 278
removespare 279
removespcredential 280
removesralertcrit 281
removesshkey 282
removetask 282
removetemplate 283
removeuser 284
removeuserconn 285
removevfs 286
removevlun 287
removevv 288
removevvset 290
removewsapisession 291

Restore Command 293
restorefsconf 293

Service Commands 294
servicecage 294
servicehost 295
servicemag 297
servicenode 300

Set Commands 302
setalert 302
setaocfg 302
setauthparam 304
setbanner 309
setbattery 309
setcage 310
setcim 312
setclienv 314
setcorequest 316
setcpg 317
setdate 321
setdomain 323
setdomainset 324
setflashcache 324
setfpg 326
setfs 327
setfsarchive 337
setfsaudit 340
setfsav 345
setfsgroup 348
setfshare 349
setfsip 357

setfsndmp 358
setfsquota 360
setfsroute 362
setfstore 363
setfsuser 364
sethost 365
sethostset 368
setlicense 368
setmaint 370
setnet 370
setnode 374
setpassword 375
setpd 376
setqos 377
setrcopygroup 378
setrcopytarget 386
setsched 390
setsnmpmgr 391
setsnmppw 393
setsnmpuserpw 394
setsralertcrit 395
setsshkey 403
setstatch 404
setstatpdch 405
setsys 405
setsysmgr 411
settask 412
settemplate 413
setuser 414
setuseracl 415
setvasa 416
setvfs 418
setvv 419
setvvolsc 422
setvvset 423
setwsapi 423

Show Commands 426

showalert 426
showaocfg 427
showauthparam 428
showbanner 428
showbattery 429
showblock 431
showcage 433
showcert 435
showcim 436
showclienv 437
showcorequest 438
showcpg 438
showdate 441
showdomain 442

showdomainset 443
showeeprom 444
showencryption 445
showeventlog 445
showfirmwaredb 447
showflashcache 448
showfpg 450
showfs 451
showfsarchive 455
showfsaudit 458
showfsav 460
showfsgroup 461
showfshare 462
showfsip 464
showfsnap 464
showfsnapclean 466
showfsndmp 467
showfspn 468
showfsquota 469
showfsroute 470
showfstore 471
showfsuser 472
showhost 472
showhostset 474
showinventory 475
showiscsisession 476
showld 477
showldch 480
showldmap 481
showlicense 482
showmaint 482
shownet 483
shownode 484
shownodeenv 487
showpatch 488
showpd 489
showpdata 493
showpdch 493
showpdvv 497
showport 499
showportarp 504
showportdev 504
showportisns 507
showportlesb 508
showportpel 511
showqos 513
showrcopy 514
showrctransport 516
showrole 516
showrsv 518
showsched 519
showsnmpmgr 520
showsnmppw 521

showsnmpuser 521
showspace 522
showspare 525
showsr 526
showsralertcrit 527
showsshkey 528
showsys 529
showsysmgr 531
showtarget 532
showtask 534
showtemplate 535
showtoc 536
showtocgen 536
showuser 537
showuseracl 538
showuserconn 539
showvasa 540
showversion 542
showvfs 543
showvlun 543
showvv 546
showvvcpg 551
showvvmap 552
showvvolsc 553
showvvolvm 554
showvvpd 557
showvvset 559
showwsapi 560
showwsapisession 561

Shutdown Commands 563
shutdownnode 563
shutdownsys 564

System Reporter (SR) Commands 567
sraomoves 567
srcpgspace 568
srhistld 571
srhistpd 574
srhistport 577
srhistvlun 580
srldspace 584
srpdspace 587
srrgiodensity 591
srstatcache 592
srstatcmp 595
srstatcpu 598
srstatfsav 601
srstatfsblock 603
srstatfscpu 606
srstatfsfpg 609

srstatfsmem 611
srstatfsnet 614
srstatfsnfs 617
srstatfssmb 620
srstatfssnapshot 622
srstatiscsi 625
srstatiscsisession 628
srstatld 631
srstatlink 633
srstatpd 636
srstatport 639
srstatqos 642
srstatrcopy 645
srstatrcvv 648
srstatvlun 652
srstatvv 656
srsysspace 659
srvvspace 662

Start Commands 667

startao 667
startcim 669
startfs 670
startfsarchive 671
startfsav 672
startfsnapclean 673
startfsndmp 674
startld 675
startnoderescue 675
startrcopy 676
startrcopygroup 677
starttask 678
startvasa 679
startvv 680
startwsapi 680

Stat Commands 682

statcache 682
statch 684
statcmp 687
statcpu 689
statfcoe 690
statfs 691
statiscsi 692
statiscsisession 694
statld 695
statlink 698
statpd 698
statport 702
statqos 705
statrcopy 708

statrcvv 709
statvlun 712
statvv 715
statwsapi 718

Stop Commands 720
stopcim 720
stopfs 720
stopfsarchive 721
stopfsav 722
stopfsnapclean 723
stopfsndmp 724
stoprcopy 725
stoprcopygroup 726
stopvasa 727
stopwsapi 728

Sync Command 730
syncrcopy 730

Tune Commands 732
tunepd 732
tunesys 734
tunevv 739

Update Commands 742
updatesnapspace 742
updatevv 743

Upgrade Commands 745
upgradecage 745
upgradepd 746

Wait Command 748
waittask 748

Mapping roles and rights 749
3PAR_AO Role 749
3PAR_RM Role 750
Audit Role 751
Basic_edit Role 752
Browse Role 754
CO Role 755
Create Role 755
Edit Role 757
Security Role 762
Service Role 763

Super Role 766

Command Authorities 777

Websites 791

Support and other resources 792

Accessing Hewlett Packard Enterprise Support 792
Accessing updates 792
Remote support 793
Warranty information 793
Regulatory information 793
Documentation feedback 794

Command Quick Reference

The HPE 3PAR CLI provides commands for managing your 3PAR storage systems.

Banner commands

Command	Description	Authority
removebanner	Remove the banners.	Super Any role granted the `sshbanner_set` right.
setbanner	Set the banner that is displayed before the user logs in.	Super Any role granted the `sshbanner_set` right.
showbanner	Show the banners.	Any role in the system.

Certificate commands

Command	Description	Authority
createcert	Create self-signed SSL certificate or a certificate signing request (CSR) for the 3PAR storage system SSL services.	Super
showcert	Show information about SSL certificates of the 3PAR storage system.	Any role in the system.
importcert	Import certificates for a given service.	Super
removecert	Remove certificates that are no longer trusted.	Super

CIM server commands

Command	Description	Authority
setcim	Set the properties of the CIM server, including options to enable or disable the SLP, HTTP, and HTTPS ports for the CIM server.	Super, Service [1] Any role granted the `cim_set` right.
showcim	Display the CIM server setting information and status.	Any role in the system.
startcim	Start the CIM server to service CIM requests.	Super, Service[1] Any role granted the `cim_start` right.
stopcim	Stop the CIM server from servicing CIM requests.	Super, Service[1] Any role granted the `cim_stop` right.

[1] Requires access to all domains.

Disk enclosure management commands

The 3PAR CLI provides a number of commands you can use to manage disks and disk enclosures.

More information

Drive cage management
Encryption commands
Physical disk management commands

Drive cage management

Command	Description	Authority
locatecage	Locate a particular drive cage.	Super, Service[1] Any role granted the `cage_locate` right.
setcage	Set parameters for a drive cage.	Super, Service[1] Any role granted the `cage_set` right.
showcage	Display drive cage information.	Any role in the system.

1 Requires access to all domains.

Encryption commands

Command	Description	Authority
controlencryption	Control data encryption.	Super Any role granted the `encryption_control` right. Super, Service (for `status` subcommand only). Any role granted the `encryption_status` right.
showencryption	Show data encryption.	Super, Service Any role granted the `control_encryption_status` right.

Physical disk management commands

Command	Description	Authority
admitpd	Admit one or all physical disks to enable their use.	Super, Service[1] Any role granted the `pd_admit` right.
checkpd	Execute surface scans on physical disks.	Super, Service[1] Any role granted the `pd_check` right.
controlpd	Spin physical disks up or down.	Super, Service[1] Any role granted the `pd_control` right.
dismisspd	Dismiss one or more physical disks from use.	Super, Service[1] Any role granted the `pd_dismiss` right.
movepd	Move data from specified physical disks to a temporary location selected by the system.	Super, Service, Edit Any role granted the `pd_move` right.
setpd	Mark physical disks as allocatable for logical disks.	Super, Service[1]
showpd	Display physical disks in the system.	Any role in the system.

[1] You must have access to all domains to run this command.

Domain management commands

Command	Description	Authority
changedomain	Change the current domain CLI environment parameter.	Any role in the system.
createdomain	Show a list of domains on the system.	Super [1] Any role granted the `domain_create` right.
createdomainset	Define a new set of domains and provide the option of assigning one or more domains to that set.	Super Any role granted the `domainset_create` right can create a domain set. Any role granted the `domainset_set` right can add domains to a domain set.
movetodomain	Move objects from one domain to another.	Super, Edit[1] Any role granted the `domain_moveto` right.
removedomain	Remove an existing domain from the system.	Super[1] Any role granted the `domain_remove` right.
removedomainset	Remove a domain set or remove domains from an existing set.	Super, Edit Any role granted the `domainset_remove` right can remove a domain set or domains from a domain set.
setdomain	Set the parameters and modify the properties of a domain.	Super[1] Any role granted the `domain_set` right.
setdomainset	Set the parameters and modify the properties of a domain set.	Super, Edit Any role granted the `domainset_set` right.
showdomain	Display a list of domains on a system.	Any role in the system.
showdomainset	Display domain sets defined on the 3PAR Storage System and their members.	Any role in the system.

[1] Requires access to all domains,

File persona commands

Command	Description	Authority
createfspn	Create SPN for a given VFS to enable Kerberos authentication for File Persona.	Super, Edit Any role granted the `fspn_create` right.
removefspn	Remove SPN associated with a VFS to disable Kerberos authentication for File Persona.	Super, Edit Any role granted the `fspn_delete` right.
setfs	Update global File Persona settings.	Super Any role granted the `fs_set` right.
showfs	Show information of a File Persona cluster.	Any role in the system.
showfspn	Display information about SPN and VFS association.	Any role in the system.
startfs	Initialize and start File Persona on the system.	Super, Service Any role granted the `fs_start` right.
statfs	Show statistics for File Persona.	Any role in the system.
stopfs	Stop or remove File Persona.	Super, Service

More information

File persona antivirus service
File persona archiving
File persona commands
File Persona group accounts
File Persona NDMP
File Persona routes
File Persona snapshots
File Persona user accounts

File persona antivirus service

Command	Description	Authority
setfsav	Set antivirus properties for File Persona.	Super, Edit Any role granted the `fsav_set` right.
showfsav	Show antivirus properties for File Persona.	Any role in the system.
startfsav	Start antivirus service or scan for File Persona.	Super, Edit Any role granted the `fsav_start` right.
stopfsav	Stop the antivirus service or stop/pause a scan.	Super, Edit Any role granted the `fsav_stop` right.

File persona archiving

Command	Description	Authority
removefsarchive	Delete the WORM/WORM-retained files, remove the retention period, or delete retention validation scan on the fstore.	Super, Edit [1] Any role granted the `fsarchive_remove` right. Super for `-importfile <source_path>`.
setfsarchive	Set or modify archiving properties for File Persona.	Super, Edit[1] Any role granted the `fsarchive_set` right. Super for `-importfile <source_path>`.
showfsarchive	Display policy setting, retention setting, and status of validation scans.	Any role in the system[1]. The `-export` subcommand, as well as the `-importfile <source_path>` and `-export <target_path>` options can be accessed only by users with Super role.

Table Continued

Command	Description	Authority
startfsarchive	Start or resume data validation jobs.	Super, Service[1] Any role granted the `fsarchive_start` right.
stopfsarchive	Stop or pause data validation jobs.	Super, Service[1] Any role granted the `fsarchive_stop` right.

[1] Requires access to all domans.

File Persona group accounts

Command	Description	Authority
createfsgroup	Create a local group account associated with File Persona.	Super Any role granted the `fsgroup_create` right.
showfsgroup	Show local group information associated with File Persona.	Any role in the system.
removefsgroup	Remove a local group account associated with File Persona.	Super Any role granted the `fsgroup_remove` right.
setfsgroup	Modify a local group account associated with File Persona.	Super Any role granted the `fsgroup_set` right.

File Persona NDMP

Command	Description	Authority
setfsndmp	Set NDMP properties for File Persona.	Super, Edit Any role granted the `fsndmp_set` right.
showfsndmp	Show NDMP properties for File Persona.	Any role in the system.

Table Continued

Command	Description	Authority
startfsndmp	Start NDMP and ISCSI service.	Super, Edit Any role granted the `fsndmp_start` right.
stopfsndmp	Stop NDMP and ISCSI service.	Super, Edit Any role granted the `fsndmp_stop` right.

File Persona routes

Command	Description	Authority
createfsroute	Create a route for a target address with a gateway.	Edit Any role granted the `fsroute_set` right.
removefsroute	Remove an existing route for a target address.	Edit Any role granted the `fsroute_set` right.
setfsroute	Modify an existing route for a target address.	Edit Any role granted the `fsroute_set` right.
showfsroute	Display routes for target addresses.	Any role in the system.

File Persona snapshots

Command	Description	Authority
createfsnap	Create a snapshot for File Persona.	Super, Edit Any role granted the `fsnap_create` right.
removefsnap	Remove file store snapshots from File Persona.	Super, Edit Any role granted the `fsnap_remove` right.
showfsnap	Show snapshot information for File Persona.	Any role in the system.
showfsnapclean	Show details of an on-demand snapshot reclamation task.	Any role in the system.

Table Continued

Command	Description	Authority
startfsnapclean	Start or resume an on-demand snapshot reclamation task.	Super, Service Any role granted the `fsnapclean_start` right.
stopfsnapclean	Stop or pause an on-demand snapshot reclamation task.	Super, Service Any role granted the `fsnapclean_stop` right.

File Persona user accounts

Command	Description	Authority
createfsuser	Create a local user account associated with File Persona.	Super Any role granted the `fsuser_create` right.
removefsuser	Remove a local user account associated with File Persona.	Super Any role granted the `fsuser_remove` right.
setfsuser	Modify a local user account associated with File Persona.	Super Any role granted the `fsuser_set` right.
showfsuser	Show local user information associated with File Persona.	Any role in the system.

Health and alert management commands

Use the 3PAR CLI to review alerts, events, and system health.

More information

Alerts
Events
System health

Alerts

Command	Description	Authority
removealert	Remove one or more alerts.	Super, Service [1] Any role granted the `alert_remove` right.
setalert	Set the status of system alerts.	Super, Service[1] Any role granted the `alert_set` right.
showalert	Display system alerts.	Any role in the system.

[1] Requires access to all domains.

Alert security levels

Table 1: Alert Severity Levels

Severity	Description
Fatal	A fatal event has occurred. It is no longer possible to take remedial action.
Critical	The event is critical and requires immediate action.
Major	The event requires immediate action.
Minor	An event has occurred that requires action, but the situation is not yet serious.
Degraded	An aspect of performance or availability may have become degraded. You must decide if action is necessary.
Informational	The event is informational. No action is required other than acknowledging or removing the alert.

Events

Command	Description	Authority
showeventlog	Display event logs.	Any role in the system.

System health

Command	Description	Authority
checkhealth	Display the status of system hardware and software components.	Super, Service Any role granted the `health_check` right.

Help and utility commands

Command	Description	Authority
cli	Provide a means to set up your CLI session or to enter directly into a CLI shell.	Super, Service, Edit, Browse
clihelp	List all commands or details for a specified command.	Super, Service, Edit, Browse
cmore	Page the output of commands.	Super, Service, Edit, Browse
help	List all commands or details for a specified command.	Super, Service, Edit, Browse
setclienv	Set the CLI environment parameters.	Super, Service, Edit, Browse Any role granted the `clienv_set` right.
showclienv	Display the CLI environment parameters.	Any role in the system.

LDAP management commands

Command	Description	Authority
setauthparam	Set the authentication parameters.	Super Any role granted the `authparam_set` right.
showauthparam	Show authentication parameters and integrates the authentication and authorization features using LDAP.	Any role in the system.
checkpassword	Support authentication and authorization using LDAP.	Super, Service, Edit, Browse Any role granted the `password_checkany` right can check the password of any user. Any role granted the `password_checkown` right can only check their own password.

License management commands

Command	Description	Authority
setlicense	Set the license key.	Super, Service [1] Any role granted the `license_set` right.
showlicense	Display the installed license info or key.	Any role in the system.

[1] Requires access to all domains.

Maintenance mode commands

Command	Description	Authority
createmaint	Creates a maintenance window record with the specified options and maintenance type.	Super, Service, maint_set
setmaint	Allows modification of the maintenance window record with the specified options for the maintenance type.	Super, Service, maint_set
showmaint	Displays maintenance window records.	Any role in the system.

Node subsystem management

Use the 3PAR CLI to manage subsystems such as nodes and ports, and view and verify network information.

More information

Array and controller node
Battery management
Controller node EEPROM log
Controller node properties
Firmware versions

Array and controller node

Command	Description	Authority
locatenode	Locate a particular node component by blinking LEDs on the node.	Super, Service [1] Any role granted the `node_locate` right.
locatesys	Locate a system by blinking its LEDs.	Super, Service[1] Any role granted the `sys_locate` right.
setsys	Set system-wide parameters such as the raw space alert.	Super, Service[1] Any role granted the `sys_set` right.
showsys	Display the 3PAR storage system properties, including system name, model, serial number, and system capacity.	Any role in the system.

[1] Requires access to all domains.

Battery management

Command	Description	Authority
setbattery	Set battery properties.	Super, Service [1] Any role granted the `battery_set` right.
showbattery	Display battery status information.	Any role in the system.

[1] Requires access to all domains.

Controller node EEPROM log

Command	Description	Authority
showeeprom	Display node EEPROM information.	Any role in the system.

Controller node properties

Command	Description	Authority
setnode	Set the properties of the node components such as the serial number of the power supply.	Super, Service [1] Any role granted the `node_set` right.
shownode	Display an overview of node-specific properties.	Any role in the system.
shownodeenv	Display the node's environmental status.	Any role in the system.

[1] Requires access to all domains.

Firmware versions

Command	Description	Authority
showfirmwaredb	Display a current database of firmware levels.	Any role in the system.

Network interface configuration

Command	Description	Authority
setnet	Set the administration network interface configuration.	Super, Service [1] Any role granted the `net_set` right.
shownet	Display the network configuration and status.	Any role in the system.

[1] Requires access to all domains.

Node date information

Command	Description	Authority
setdate	Set the system time and date on all nodes.	Super, Service [1] Any role granted the `date_set` right.
showdate	Display the date and time on all system nodes.	Any role in the system.

[1] Requires access to all domains.

Node rescue

Command	Description	Authority
startnoderescue	Initiate a node rescue, which initializes the internal node disk of the specified node to match the contents of the other node disks.	Super, Service

Port information

Command	Description	Authority
checkport	Perform a loopback test on Fibre Channel ports.	Super, Service [1] Any role granted the `port_check` right.
controlport	Control Fibre Channel or Remote Copy ports.	Super, Service[1] Any role granted the `port_control` right.
controliscsiport	Set up parameters and characteristics of an iSCSI port.	Super, Service[1] Any role granted the `port_control` right.
showiscsisession	Show iSCSI active sessions per port.	Any role in the system.
showport	Display system port information.	Any role in the system.
showportarp	Show ARP table for iSCSI ports in the system.	Any role in the system.

Table Continued

Command	Description	Authority
showportdev	Displays detailed information about devices on a Fibre Channel port.	Any role in the system.
showportisns	Show iSNS host information for iSCSI ports in the system.	Any role in the system.
showportlesb	Displays Link Error Status Block information about devices on a Fibre Channel port.	Any role in the system.
showtarget	Displays unrecognized targets.	Any role in the system.
statiscsi	Displays the iSCSI statistics.	Any role in the system.
statiscsisession	Displays the iSCSI session statistics.	Any role in the system.

[1] Requires access to all domains.

System manager

Command	Description	Authority
setsysmgr	Set the system manager startup state.	Super [1] Any role granted the `sysmgr_set` right.
showsysmgr	Display the system manager startup state.	Any role in the system.
showtoc	Display the system table of contents summary.	Any role in the system.
showtocgen	Display the system table of contents generation number.	Any role in the system.

[1] Requires access to all domains.

Performance management commands

Use the HPE 3PAR CLI to check performance statistics and manage virtual volumes.

More information

Chunklet statistics
Data cache memory statistics
Dynamic optimization
Link statistics
Logical disk statistics
Node CPU statistics
Physical disk statistics
Port statistics
Remote Copy volume statistics
System tuner
Virtual LUN (export) statistics
Virtual volume statistics

Chunklet statistics

Command	Description	Authority
histch	Display histogram data for individual chunklets.	Any role in the system.
setstatch	Set statistics collection mode on chunklets.	Super, Edit Any role granted the `statch_set` right.
setstatpdch	Set statistics collection mode on physical disk chunklets.	Super, Edit Any role granted the `statpdch_set` right.
statch	Display statistics for individual chunklets.	Any role in the system.

Data cache memory statistics

Command	Description	Authority
statcmp	Display statistics for cache memory pages.	Any role in the system.

Dynamic optimization

Command	Description	Authority
tunevv	Change the layout of a virtual volume.	Super, Edit Any role granted the `vv_tune` right.

Link statistics

Command	Description	Authority
statlink	Display statistics for links.	Any role in the system.

Logical disk statistics

Command	Description	Authority
histld	Display histogram data for logical disks.	Any role in the system.
statld	Display statistics for logical disks.	Any role in the system.

Node CPU statistics

Command	Description	Authority
statcpu	Display statistics for CPU use.	Any role in the system.

Physical disk statistics

Command	Description	Authority
histpd	Display histogram data for physical disks.	Any role in the system.
statpd	Display statistics for physical disks.	Any role in the system.

Port statistics

Command	Description	Authority
histport	Display histogram data for Fibre Channel ports.	Any role in the system.
statport	Display statistics for Fibre Channel ports.	Any role in the system.

Preserved data commands

Command	Description	Authority
showpdata	Display preserved data status.	Any role in the system.

Remote Copy volume statistics

Command	Description	Authority
histrcvv	Display histogram data for Remote Copy volumes.	Any role in the system.
statrcvv	Display statistics for Remote Copy volumes.	Any role in the system.

System tuner

Command	Description	Authority
tunepd	Display physical disks with high service times, and optionally perform load balancing.	Super, Edit Any role granted the `pd_tune` right.
tunesys	Analyze disk usage and adjust resources.	Super, Edit Any role granted the `sys_tune` right.

Virtual LUN (export) statistics

Command	Description	Authority
histvlun	Display histogram data for VLUNs.	Any role in the system.
statvlun	Display statistics for VLUNs.	Any role in the system.

Virtual volume statistics

Command	Description	Authority
histvv	Display histogram data for virtual volumes.	Any role in the system.
statvv	Display statistics for virtual volumes.	Any role in the system.

Replication commands

Replication commands available within the CLI include physical copy, remote copy, virtual copy commands.

More information

Physical copy commands
Remote copy commands
Virtual copy commands

Physical copy commands

Command	Description	Authority
creategroupvvcopy	Create consistent group physical copies of a list of virtual volumes.	Super, Edit Any role granted the `groupvvcopy_create` right.
createvvcopy	Copy a virtual volume.	Super, Edit Any role granted the `vvcopy_create` right.
promotevvcopy	Promote a physical copy back to a base volume.	Super, Edit Any role granted the `vvcopy_promote` right.

Remote copy commands

HPE 3PAR Remote Copy functionality requires a 3PAR Remote Copy license. See the *HPE 3PAR StoreServ Storage Concepts Guide* for additional information.

Command	Description	Authority
admitrcopylink	Admit a network link for Remote Copy use.	Super, Edit Any role granted the `rcopylink_admit` right.
admitrcopytarget	Add a target to a Remote Copy volume group	Super, Edit Any role granted the `rcopytarget_admit` right.

Table Continued

Command	Description	Authority
admitrcopyvv	Admit a virtual volume to a Remote Copy volume group.	Super, Edit Any role granted the `rcopyvv_admit` right.
checkrclink	Perform a latency and throughput test on a remote copy link.	Super, Edit Any role granted the `rclink_check` right.
creatercopygroup	Create a group for Remote Copy.	Super, Edit Any role granted the `rcopygroup_create` right.
creatercopytarget	Create a target for Remote Copy.	Super, Edit [1] Any role granted the `rcopytarget_create` right.
dismissrcopylink	Dismiss a network link from Remote Copy use.	Super, Edit[1] Any role granted the `rcopylink_dismiss` right.
dismissrcopytarget	Dismiss a Remote Copy target from a Remote Copy volume group.	Super, Edit Any role granted the `rcopytarget_dismiss` right.
dismissrcopyvv	Dismiss a virtual volume from a Remote Copy volume group.	Super, Edit Any role granted the `rcopyvv_dismiss` right.
removercopygroup	Remove a group used for Remote Copy.	Super, Edit Any role granted the `rcopygroup_remove` right.
removercopytarget	Remove a target used for Remote Copy.	Super, Edit[1] Any role granted the `rcopytarget_remove` right.
setrcopygroup	Set a volume group's policy for dealing with I/O failure and error handling, or switch the direction of a volume group.	Super, Edit Any role granted the `rcopygroup_set` right.

Table Continued

Command	Description	Authority
setrcopytarget	Set the Remote Copy target state.	Super, Edit[1] Any role granted the `rcopytarget_set` right.
showrcopy	Display the details of a Remote Copy configuration.	Any role in the system.
showrctransport	Show status and info about end-to-end transport for Remote Copy in the system.	Any role in the system.
startrcopy	Start a Remote Copy subsystem.	Super, Edit[1] Any role granted the `rcopy_start` right.
startrcopygroup	Start a Remote Copy volume group.	Super, Edit Any role granted the `rcopygroup_start` right.
statrcopy	Display Remote Copy statistics.	Any role in the system.
stoprcopy	Stop a Remote Copy subsystem.	Super, Edit[1] Any role granted the `rcopy_stop` right.
stoprcopygroup	Stop a Remote Copy volume group.	Super, Edit Any role granted the `rcopygroup_stop` right.
syncrcopy	Synchronize Remote Copy volume groups.	Super, Edit Any role granted the `rcopy_sync` right.

1 Requires access to all domains.

Virtual copy commands

Command	Description	Authority
createsv	Create snapshot volumes.	Super, Edit Any role granted the `sv_create` right.
creategroupsv	Create consistent group snapshots of a list of virtual volumes.	Super, Edit Any role granted the `groupsv_create` right.
promotesv	Copy the differences of a virtual copy back to its base volume.	Super, Edit Any role granted the `sv_promote` right.
promotegroupsv	Copy the differences of snapshots back to their base volumes, to allow to revert the base volumes to an earlier point in time.	Super, Edit Any role granted the `svgrp_promote` right.
updatevv	Update a snapshot virtual volume with a new snapshot.	Super, Edit Any role granted the `vv_update` right.

Security hardening commands

Authenticate recovery accounts and manage SP credentials.

More information

Service Processor credential commands
Support recovery account commands

Service Processor credential commands

Command	Description	Authority
removespcredential	Remove all Service Processor credentials on the array.	Super Any role granted the `spcred_remove` right.

Support recovery account commands

Command	Description	Authority
controlrecoveryauth	Control the method used to authenticate recovery accounts.	Super Any role granted the `recoveryauth_control` right. Super, Service (for `status` subcommand only). Any role granted the `recoveryauth_status` right.

Service commands

Use the 3PAR CLI to for tasks such as general system maintenance, managing nodes and disk enclosures, or checking software versions.

More information

Disk enclosure commands
Node rescue
General system maintenance commands
QoS commands
Software version commands

Disk enclosure commands

Command	Description	Authority
admithw	Admit new hardware into the system.	Super, Service [1]
controlmag	Take drives or magazines on or off loop.	Super, Service[1] Any role granted the `mag_control` right.
servicecage	Prepare a drive cage for service.	Super, Service[1]
servicehost	Prepare a port for host attachment.	Super, Service[1] Any role granted the `host_service` right.
servicemag	Prepare a drive magazine for service.	Super, Service[1] Any role granted the `mag_service` right.
upgradecage	Upgrade drive cage firmware.	Super, Service[1] Any role granted the `cage_upgrade` right.
upgradepd	Upgrade disk firmware.	Super, Service[1] Any role granted the `pd_upgrade` right.

[1] Requires access to all domains.

Node commands

Command	Description	Authority
servicenode	Prepare a node for service.	Super, Service Any role granted the node_service right. Requires access to all domains.

General system maintenance commands

Command	Description	Authority
shutdownnode	Shut down an individual system node.	Super, Service [1] Any role granted the node_shutdown right.
shutdownsys	Shut down the entire system.	Super, Service[1] Any role granted the sys_shutdown right.

[1] Requires access to all domains.

QoS commands

Command	Description	Authority
setqos	Create and update QoS rules in a system.	Super, Edit Any role granted the qos_set right can set QoS configurations.
showqos	List the QoS rules configured in a system.	Any role in the system.
statqos	Display historical performance data reports for QoS rules.	Any role in the system.

Software version commands

Command	Description	Authority
`showpatch`	Display patches applied to a system.	Any role in the system.
`showversion`	Display software versions.	Any role in the system.

SNMP agent commands

Command	Description	Authority
addsnmpmgr	Add an SNMP manager to receive trap notifications.	Super [1] Any role granted the `snmpmgr_add` right.
checksnmp	Allow a user to send an SNMPv2 test trap to the list of managers	Super, Service Any role granted the `snmp_check` right.
removesnmpmgr	Remove an SNMP trap manager.	Super[1] Any role granted the `snmpmgr_remove` right.
removesnmppw	Remove an SNMP password.	Super[1] Any role granted the `snmppw_remove` right.
removesnmpuser	Remove an SNMP user.	Super, Service, Edit, Browse Any role granted the `snmpuser_remove` right.
setsnmpmgr	Change an SNMP manager's properties.	Super[1] Any role granted the `snmpmgr_set` right.
setsnmpuserpw	Update SNMP v3 user authentication and privacy passwords.	Super[1] Any role granted the `snmppw_set` right.
setsnmppw	Allow users to update SNMP passwords.	Super[1] Any role granted the `snmppw_set` right.
showsnmpmgr	Display SNMP trap managers.	Any role in the system.
showsnmppw	Display SNMP access passwords.	Any role in the system.
showsnmpuser	Display information about SNMP users.	Any role in the system.

1 Requires access to all domains.

Sparing commands

Command	Description	Authority
createspare	Create spare chunklets.	Super, Service [1] Any role granted the `spare_create` right.
movech	Move specified chunklets.	Super, Service, Edit[1] Any role granted the `ch_move` right.
movechtospare	Move specified chunklets to spare.	Super, Service, Edit[1] Any role granted the `ch_movetospare` right.
movepd	Move data from specified physical disks to a temporary location selected by the system.	Super, Service, Edit Any role granted the `pd_move` right.
movepdtospare	Move specified physical disks to spare.	Super, Service, Edit[1] Any role granted the `pd_movetospare` right.
moverelocpd	Move chunklets relocated from a physical disk to another physical disk.	Super, Service[1] Any role granted the `pd_movereloc` right.
removespare	Remove spare chunklets.	Service, Edit[1] Any role granted the `spare_remove` right.
showspare	Display information about spare and relocated chunklets.	Any role in the system.

[1] Requires access to all domains.

SSH access commands

Command	Description	Authority
setsshkey	Set the SSH public key for users, enabling login without a password.	Super, Service, Edit, Browse Any role granted the `sshkey_set` right.
showsshkey	Display all SSH public keys that have been set with `setshhkey`.	Any role in the system.
removesshkey	Remove a user's SSH public key.	Super, Service, Edit, Browse Any role granted the `sshkey_remove` right.

Task management commands

Command	Description	Authority
canceltask	Cancel one or more tasks.	Super, Edit Any role granted the `task_cancel` right.
removetask	Remove information about one or more tasks and their details.	Super, Edit Any role granted the `task_remove` right.
settask	Set the priority on a specified task.	Super, Edit Any role granted the `task_set` right.
showtask	Display information about tasks.	Any role in the system.
starttask	Execute commands with long running times.	Super, Service, Edit Any role granted the `task_start` right.
waittask	Ask the CLI to wait for a task to complete before proceeding.	Super, Service, Edit, Browse Any role granted the `task_wait` right.

Task schedule commands

Command	Description	Authority
createsched	Allow users to schedule tasks that are periodically run by the scheduler.	Super, Service Any role granted the `sched_create` right.
removesched	Remove a scheduled task from the system.	Super, Service, Edit Any role granted the `sched_remove` right can remove any user scheduled tasks. Any role granted the `sched_setown` right can only remove their own scheduled tasks.
setsched	Allow users to suspend, pause, change the schedule, change the parameters, and change the name of currently scheduled tasks.	Super, Service Any role granted the `sched_set` right.
showsched	Display the state of tasks currently scheduled on the system.	Super, Service Any extended role in the system.

User management commands

Command	Description	Authority
createuser	Create user accounts.	Super [1] Any role granted the `user_create` right.
removeuser	Remove user accounts.	Super[1] Any role granted the `user_remove` right.
removeuserconn	Remove user connections.	Super[1] Any role granted the `userconn_remove` right.
setpassword	Change a password.	Super, Service, Edit, Browse [2] Any role granted the `password_setany` right can set any user password. Any role granted the `password_setown` right can only set their own password.
setuser	Set user properties.	Super, Edit, Browse Any role granted the `user_setany` right can add, remove, and set default domains for any user. Any role granted the `user_setown` right can only set their own default domain.
setuseracl	Set an Access Control List (ACL).	Super, Edit Any role granted the `useracl_set` right.
showrole	Display information about rights assigned to roles in the system.	Any role in the system.
showuser	Display user accounts.	Any role in the system.
showuseracl	Display an access control list (ACL).	Any role in the system.
showuserconn	Display user connections.	Any role in the system.

1 Requires access to all domains

2 Only the Super user can edit another user's password.

Volume management commands

Use the 3PAR CLI to manage volumes, including CPGs, hosts, virtual volumes.

More information

Common Provisioning Group management commands
Host management commands
Logical disk management commands
Space and storage management
Template management commands
Virtual volume management commands
Virtual LUN (export) statistics

Common Provisioning Group management commands

Command	Description	Authority
compactcpg	Consolidate logical disk space in a CPG into as few logical disks as possible, allowing unused logical disks to be removed.	Super, Edit Any role granted the `cpg_compact` right.
createcpg	Create a Common Provisioning Group (CPG).	Super, Edit [1] Any role granted the `cpg_create` right.
removecpg	Remove CPGs or remove specified logical disks from CPGs.	Super, Edit[1] Any role granted the `cpg_remove` right.
setcpg	Set or change CPG properties.	Super, Edit[1] Any role granted the `cpg_set` right.
showcpg	Display CPGs in the system.	Any role in the system.

[1] Requires access to all domains.

Host management commands

Command	Description	Authority
createhost	Create host and assign paths to the host.	Super, Edit [1] Any role granted the `host_create` right.
createhostset	Create a new set of hosts, and assign existing hosts to the host set.	Super, Edit Any role granted the `hostset_create` right can create a host set. Any role granted the `hostset_set` right can add hosts to a host set.
removehost	Remove host definitions from the system.	Super, Edit[1] Any role granted the `host_remove` right.
removehostset	Remove a host set or remove hosts from an existing host set.	Super, Edit Any role granted the `hostset_remove` right can remove a host set or hosts from a host set.
showhost	Display defined hosts and host paths in the system.	Any role in the system.
showhostset	List the defined host sets and their members.	Any role in the system.
sethost	Set properties on existing system hosts, annotate a host with descriptor information, configure/remove iSCSI CHAP information, or reset a host.	Super, Edit[1] Any role granted the `host_set` right.
sethostset	Set parameters and modify the properties of a host set.	Super, Edit Any role granted the `hostset_annotate` right can only annotate the comment of a host set. Any role granted the `hostset_set` right can set any host set property.

[1] Requires access to all domains.

Logical disk management commands

Command	Description	Authority
checkld	Perform validity checks of data on logical disks (LD) and optionally repairs inconsistent LDs.	Super, Service [1] Any role granted the `ld_check` right.
compactld	Consolidate space on LDs.	Super, Edit Any role granted the `ld_compact` right.
removeld	Remove LDs from system service group.	Super, Service, Edit Any role granted the `ld_remove` right.
showld	Display configuration of LDs.	Any role in the system.
startld	Start data service on LDs.	Super, Service[1] Any role granted the `ld_start` right.

[1] Requires access to all domains.

Space and storage management

Command	Description	Authority
showblock	Display block mapping information for virtual volumes, logical disks, and physical disks.	Any role in the system.
showldch	Display logical disk to physical disk chunklet mapping.	Any role in the system.
showldmap	Display logical disk to virtual volume mapping.	Any role in the system.
showpdch	Display the status of selected chunklets of physical disks.	Any role in the system.
showpdvv	Display physical disk to virtual volume mapping.	Any role in the system.
showspace	Display estimated free space for logical disk creation.	Any role in the system.

Table Continued

Command	Description	Authority
showvvmap	Display virtual volume to logical disk mapping.	Any role in the system.
showvvolvm	Display information about all virtual machines (VVol-based) or a specific virtual machine in a system.	Any role in the system.
showvvpd	Display virtual volume distribution across physical disks.	Any role in the system.

Template management commands

Command	Description	Authority
createtemplate	Create templates for the creation of logical disks, virtual volumes, thinly provisioned virtual volumes, and common provisioning groups.	Super [1] Any role granted the `template_create` right.
removetemplate	Remove one or more templates.	Super[1] Any role granted the `template_remove` right.
settemplate	Modify template properties.	Super[1] Any role granted the `template_set` right.
showtemplate	Display existing templates.	Any role in the system.

[1] Requires access to all domains.

Virtual volume management commands

Command	Description	Authority
admitvv	Create and admit remotely exported virtual volume definitions to enable the migration of these volumes.	Super Any role granted the `vv_admit` right.
checkvv	Perform validity checks of virtual volume administrative information.	Super, Service, Edit[1] Any role granted the `vv_check` right.

Table Continued

Command	Description	Authority
createvv	Create a virtual volume from logical disks.	Super, Edit[1] Any role granted the `vv_create` right.
createvvset	Define a new set of virtual volumes and provide the option of assigning one or more existing virtual volumes to that set.	Super, Edit Any role granted the `vvset_create` right can create a volume set. Any role granted the `vvset_set` right can add volumes to a volume set.
freespace	Free SA and SD spaces from a virtual volume if they are not in use.	Super, Edit Any role granted the `space_free` right.
growvv	Increase the size of a virtual volume by adding logical disks.	Super, Edit[1] Any role granted the `vv_grow` right.
importvv	Migrate data from a remote LUN to the local 3PAR storage system.	Super, Edit Any role granted the `vv_import` right.
removevv	Remove virtual volumes or logical disks from common provisioning groups.	Super, Edit Any role granted the `basevv_remove` right can remove base volumes. Any role granted the `vvcopy_remove` right can remove physical volumes. Any role granted the `sv_remove` right can remove virtual copies.
removevvset	Remove a virtual volume set or virtual volumes from an existing set.	Super, Edit Any role granted the `vvset_set` right can remove a volume set or volumes from a volume set.
setvv	Modify properties associated with a virtual volume.	Super, Edit Any role granted the `vv_set` right.
setvvolsc	Create, remove, and set properties of Storage Containers for virtual volumes.	Super, Edit Any role granted the `vvset_set` right.

Table Continued

Command	Description	Authority
setvvset	Set parameters and modify properties of a virtual volume set.	Super, Edit Any role granted the `vvset_set` right.
showrsv	Display information about reservation and registration of VLUNs connected on a Fibre Channel port.	Any role in the system.
showvv	Display virtual volumes in the system.	Any role in the system.
showvvolsc	Display information about VVol storage containers in the system.	Any role in the system.
showvvcpg	Display virtual volume sets defined on the 3PAR storage system and their associated members.	Any role in the system.
showvvset	Display virtual volume sets defined on the 3PAR storage system and their members.	Any role in the system.
startvv	Start virtual volumes.	Super, Service[1] Any role granted the `vv_start` right.
updatesnapspace	Start a task to update the actual snapshot space used by a virtual volume.	Super, Edit Any role granted the `snapspace_update` right.

[1] Requires access to all domains.

Virtual LUN (export) statistics

Command	Description	Authority
histvlun	Display histogram data for VLUNs.	Any role in the system.
statvlun	Display statistics for VLUNs.	Any role in the system.

Web Services API commands

Command	Description	Authority
removewsapisession	Removes the WSAPI user connections.	Super Any role granted the `wsapisession_remove` right.
setwsapi	Sets properties of the Web Services API server.	Super, Service Any role granted the `wsapi_set` right.
showwsapi	Displays the WSAPI server service configuration state.	Any role in the system.
showwsapisession	Displays the WSAPI server sessions connection information.	Any role in the system.
startwsapi	Starts the WSAPI server.	Super, Service Any role granted the `wsapi_start` right.
stopwsapi	Stops the WSAPI server.	Super, Service Any role granted the `wsapi_stop` right.

Storage management with the HPE HPE 3PAR CLI

The HPE 3PAR CLI is a powerful and simple way to manage your HPE 3PAR storage arrays. You can use the CLI to create scripts that automate repetitive storage administrative tasks, such as taking snapshots and auditing logs. While simple, the right switches and options can produce powerful output.

Backslash characters in filenames for Microsoft Windows environments

Use of backslash (\ ) is okay when using a filename with an absolute or relative path at the command prompt:

```
cli% importcert wsapi -ca C:\certs\ca-bundle.pem
```

For Windows® shell, use:

```
C:\Users\w_user> importcert wsapi -ca C:\certs\ca-bundle.pem
```

Double each backslash when using the CLI shell or a Tcl application to enter a filename:

```
cli% importcert wsapi -ca C:\\certs\\ca-bundle.pem
```

Or, surround the filename with braces:

```
cli% importcert wsapi -ca {C:\certs\ca-bundle.pem}
```

Command syntax

Most HPE 3PAR CLI commands use the following syntax:

```
cmd [subcmd] [option [arg]]... [<spec>...][<pattern>...]
```

The following table lists the syntax elements and their meanings:

Table 2: CLI command syntax

Element	Meaning
`cmd`	Specifies the operation to be executed, such as `createcpg`, `movepdtospare`, or `showcage`.
`subcmd`	Specifies action that a command will execute, such as `controlencryption enable` (where `enable` is a subcommand of the `controlencryption` command). Not all commands contain subcommands.
`option`	Indicates an optional command line element, such as `-rw`, `-f`, or `-waittask`. Options begin with the dash (–) symbol.

Table Continued

Element	Meaning
`arg`	Indicates a specific variable (argument) for an option. Often used to identify a specific node, volume, or disk. For example, in the command `controlport rst -m target 6:3:1`, `target` is an `arg`.
`spec`	Indicates a specifier used with a required command line element such as a command or option. For example, `backupfsconf <vfs>` (where *<vfs>* indicates a value that should be replaced with the actual name of the VFS server). Specifiers are enclosed in `< >` characters.
`|`	Requires choosing one `option` or `spec` when separated by this character, such as `-y|-n` or *<VV_name|pattern|VV_Set>*.
`{ }`	Indicates grouped elements. Do not enter the braces; enter only the information inside the braces.
`[ ]`	Indicates optional elements. Do not enter the brackets; enter only the information inside the brackets.
`< >`	Indicates user-supplied input. See, `spec`, above.
`...`	Indicates that a specifier or an option can be used more than once in a command.

More information

Syntax rules for the CLI

Syntax rules for the CLI

The 3PAR CLI command syntax uses the following rules:

- All command line elements are in lowercase letters, except where indicated in this reference.
- Subcommands immediately follow the commands they modify.
- Options, as indicated in this guide, are indicated by one or more letters, are preceded by a hyphen, and are enclosed in brackets (for example: `removealert [-a]`).
- Options follow subcommands on the command line and precede any specifier.
- An argument directly follows the option or subcommand it is modifying and is required on the command line, unless otherwise specified (for example: `removealert -i <alert_ID>`).
- Multiple options and arguments on a command line are separated with a space.
- Specifiers follow options.
- User supplied input is identified by angled brackets (< >).
- Unless noted otherwise, valid character formats include alphanumeric characters, periods, dashes, and underscores.

The 3PAR CLI enforces the following length limits:

- Virtual volume name ≤ 31 characters
- Thinly Provisioned Virtual Volume (TPVV) name ≤ 31 characters
- Virtual copy name ≤ 31 characters
- Logical disk name ≤ 31 characters
- Host name ≤ 31 characters
- Common Provisioning Group (CPG) name ≤ 31 characters
- Template name ≤ 31 characters
- Domain name ≤ 31 characters
- Snapshot name ≤ 31 characters
- User name ≤ 31 characters
- System name ≤ 31 characters
- File Provisioning Group (FPG) name ≤ 21 characters

More information

Command syntax

Documentation layout for command descriptions

The CLI documentation displays command information using the following format:

COMMAND NAME

The name of the CLI command.

SYNTAX

Lists syntax patterns for the command (see, **Command syntax**).

DESCRIPTION

Briefly explains the use or purpose of the command.

OPTIONS

Lists and describes options you can use within the command line for various results. Some commands require certain options and their specifiers. Unless otherwise stated, you can choose to use options or not. Use of options can increase command functionality and increase the level of detail provided in the output.

SPECIFIERS

Lists and describes specifiers that control aspects of how the command is interpreted.

SUBCOMMANDS

Itemizes any subcommands available to the command. Appears only when applicable.

AUTHORITY

Defines the user access and rights required to use the command.

USAGE

Provides other pertinent information that you might need to know about the command.

EXAMPLES

Provides sample commands with results similar to what you can expect when running the command.

Exit status codes for CLI commands

The 3PAR CLI uses exit status codes to indicate the execution status of a given command. To determine the overall results of individual CLI commands, use the command output.

Except where noted, the CLI returns the following codes to indicate success or failure for an individual command.

- `0` indicates that the command executed successfully.
- `1` indicates that the command failed.
- `10` indicates that the command failed, but it can be retried.

  The CLI returns exit status 10 when using the `-enherror` option or TPDCLIENHERROR environment variable (see the **`cli`** command for details).

Glob-style patterns

The following `clihelp` command returns an explanation about the use of glob-style (shell-style) pattern matching:

```
clihelp sub,globpat
```

Several CLI commands allow you to specify a pattern that refines the command output. The command uses the patterns provided to match against a string. Strings include names, such as host or virtual volume names. For example, in the `showhost host_name|pattern...` command, you can specify a pattern to refine the command output for a string of host names that match the specified pattern.

The CLI commands use the `pattern` specifier in the form of a glob-style (shell-style) pattern. Tcl implements Glob-style matching in the CLI.

The following table lists the symbols used in glob-style patterns. Use a backslash (\) or enclosed in braces ({ }) to escape brackets ([ ]) in Tcl. Escape or quote other characters, such as star (*), when running CLI commonds from another shell.

Symbol	Explanation of Action
`*`	Matches any sequence of characters in a string, including a null string.
`?`	Matches any single character in a string.
`[chars]`	Matches any character in the set given by `chars`. Specify a range of `chars` using a dash (–). For example, `[a-z]` represents all the ASCII characters from a through z.
`\x`	Matches the single character x.

The following table lists practical examples of the glob-style pattern in use:

Example	Explanation of Action
*	Shows all names.
[a]*	Matches all names beginning with the letter a.
[a–z]x*	Matches any character a–z in the first character position and x in the second position, followed by any character (*).

Global options and environment variables

Several options and environment variables are available at the global level.

Global options are provided to help configure the CLI and control the operation of the system. Environment variables are provided to customize the CLI.

- You can use the global options with the CLI program and individual CLI commands.
- When used with the CLI program, global options remain in effect until you exit the Tcl shell. The following example shows the –sys option used with the CLI program for a system named betasystem1:

```
$ cli –sys betasystem1
```

- When used with individual commands, global options are in effect only for that command. The following example shows the –nohdtot option used with the showsysmgr CLI command:

```
$ showsysmgr -nohdtot
```

- Global options, when used, override environment variables.

NOTE: Global options cannot be specified on the command line for commands issued from a Tcl shell. Global options must be specified before starting a Tcl shell.

You can use environment variables when configuring the CLI.

NOTE: Environment variables can only be used if the CLI is set up remotely. If you are accessing the CLI through SSH, environment variables cannot be used.

Help commands and topics

Use the commands listed in the following table to find an online version of help topics related to HPE 3PAR-specific CLI commands.

See the tclsh and Tcl documentation for help on the Tcl language (**https://www.tcl.tk/**).

Help command	Result
`help`	Displays a list of CLI command categories and help usage.
`help <command>` or `<commandname> -h`	Provides help text for a specific CLI `<command>`. For example, `help addsnmpmgr` or `admithw -h`.
`help <category>`	Displays help for a specific `<category>` of commands. For example, `help tune` displays help on commands in the Tune category.
`help sub,<subject>` or `help -sub <subject>`	Displays help on subjects other than commands. For example, `help sub,number` displays help on the use of number format in commands.
`help -col <command>` or `help col,<command>`	Provides an explanation of column headers for a CLI `<command>` output.
`help -search <regex>`	Searches help text for a regular expression `<regex>` that you supply.
`help opensource`	Display opensource usage and attribution information

Numbering formats and values

The output of the following command explains the use of numbering formats in CLI commands:

```
clihelp sub,number
```

Specify integer values in following formats:

- **Decimal** – Base-10 number system. Uses the digits 0 through 9.
- **Octal** – Base-8 number system. Uses the digits 0 through 7. The first character of the operand is 0 (for example: 010).
- **Hexadecimal** – Base-16 number system. Uses the digits 0 through 9 to represent values zero to nine, and the alpha characters a, A, b, B, c, C, d, D, e, E, f, F to represent values ten to fifteen. The first two characters of the operand are 0x or 0X (for example: 0x100).

Object sets and domain sets

The output of the following command explains the use of object sets with CLI commands:

```
clihelp sub,objectsets
```

Object sets provide a convenient method of grouping logically connected objects together. For example, an `esxhosts` host set might group all host members of an ESX cluster, while the common VVs exported to them are in an `esxexports` VV set.

Commands refer to sets using the `set:` prefix. For example, the following command exports the `esxexports` VV set to the `esxhosts` host set:

```
createvlun set:esxexports 1 set:esxhosts
```

The command exports the VVs sequentially, from LUN 1 onward, to the hosts in the `esxhosts` set. Any hosts added to the host set at a later time can see the exported VVs.

The CLI also supports domain sets.

Only a host can be a member of a domain set. This membership allows the host to include VVs from all domains that are members of the domain set exported to the host.

Using domain sets allows you to back up VVs centrally, even though the VVs are under separate administrative control.

The system orders set members according to when the set was added. This affects VV set exports and when they are used for `createvvcopy`.

Not all commands support object sets in their parameters. Commands that do support object sets indicate a set parameter in the command description. Commands with set support include `createhost, createvlun, createsv, createvvcopy, movetodomain, removevlun`.

Regular expressions and patterns

The output of the following command describes regular expressions in further detail:

```
clihelp sub,regexpat
```

In the CLI, specify patterns as regular expressions (`regex`). The CLI uses regular expression patterns to match against strings. The CLI strings include descriptive text such as help text or event message strings.

See the Tcl re_syntax documentation for detailed help (**https://www.tcl.tk/**).

Regular expressions are one or more branches separated by a pipe (|), and match any string that uses the same pattern as any of the branches.

A branch is zero or more quantified atoms (an element of a regular expression) concatenated. A quantified atom is an atom followed by an optional quantifier. A branch makes a match of the first quantified atom, followed by a match of the second quantified atom, and so on. An empty branch matches an empty string.

Quantifiers are:

`* — 0 or more matches of the atom.`

`+ — 1 or more matches of the atom.`

`? — exactly one match of the atom.`

`{m} — exactly m matches of the atom.`

`{m,} — m or more matches of the atom.`

`{m,n} — m through n (inclusive) matches of the atom.`

`*? +? ?? {m}? {m,}? {m,n}? — Match the same possibilities as above but prefer the smallest number instead of the largest number of matches.`

`m and n are unsigned decimal integers from 0 through 255 inclusive.`

Atoms are:

`(re) — Where re is any regexp, matches a match for re.`

`() — Matches an empty string.`

`[chars] — A bracket expression, usually matches a single character in the list, or if the list begins with ^ matches a single character not in the`

```
list. The chars may be a range specified, for example, as a-z, or 0-9. See
detailed Tcl re_syntax help on bracket expressions.

. — Matches any single character.

\k — Matches the non-alphanumeric character k.

\c — If c is alphanumeric, matches an escape (see Tcl re_syntax help).

{ — When followed by a character other than a digit, matches {, otherwise it is a quantifier as described above.

x — Without any other significance, matches the character x.
```

Be sure to escape square brackets (`[]`) in Tcl using forward slash (`\`) or enclosing the argument in curly braces (`{}`). Other characters, such as the asterisk (`*`), are also significant in most shells. Use the same methods to escape them, or enclose them in quotations (`" "`) when running CLI commands from another shell.

Supported destination VV name patterns

The output of this command explains the supported patterns for VV names:

```
clihelp sub,vvnamepat
```

Some commands allow the use of VV name patterns to indicate the destination virtual volume names for the command. These patterns can expand based upon information such as the source virtual volume, allowing a single pattern string to represent a group of different VVs.

The CLI supports the following patterns:

- `@count@` – The current count of VVs created.
- `@vvname@` – The name of the current volume being snapshot.

The CLI supports the following date/time related patterns where it makes sense. For example, not when the destination exists, but when creating a volume.

- `@a@` – Abbreviated weekday name (Mon, Tue, and so on.).
- `@A@` – Full weekday name (Monday, Tuesday, and so on.).
- `@b@` – Abbreviated month name (Jan, Feb, and so on.).
- `@B@` – Full month name.
- `@C@` – First two digits of the four-digit year (19 or 20).
- `@d@` – Day of month (01 - 31).
- `@e@` – Day of month (1 - 31), no leading zeros.
- `@g@`—The ISO8601 year number corresponding to the ISO8601 week (%V), expressed as a two-digit year-of-the-century, with leading zero if necessary.
- `@G@` – The ISO8601 year number corresponding to the ISO8601 week (%V), expressed as a four-digit number.
- `@h@` – Abbreviated month name.
- `@H@` – Hour in 24-hour format (00 - 23).
- `@I@` – Hour in 12-hour format (01 - 12).
- `@j@` – Day of year (001 - 366).
- `@k@` – Hour in 24-hour format, without leading zeros (0 - 23).

- `@l@` – Hour in 12-hour format, without leading zeros (1 - 12).
- `@m@` – Month number (01 - 12).
- `@M@` – Minute (00 - 59).
- `@p@` – AM/PM indicator.
- `@s@` – Count of seconds since the epoch, expressed as a decimal integer.
- `@S@` – Seconds (00 - 59).
- `@u@` – Weekday number (Monday = 1, Sunday = 7).
- `@U@` – Week of year (00 - 52), Sunday is the first day of the week.
- `@V@` – Week of year according to ISO-8601 rules. Week 1 of a given year is the week containing 4 January.
- `@w@` – Weekday number (Sunday = 0, Saturday = 6).
- `@W@` – Week of year (00 - 52), Monday is the first day of the week.
- `@y@` – Year without century (00 - 99).
- `@Y@` – Year with century (such as, 1990).
- `@Z@` – Time zone name.

User space provisioning

The output of the following command explains the provisioning of a virtual volume:

```
clihelp sub,provisioning
```

Each virtual volume includes user space and snapshot space that can be provisioned or allocated (respectively) by the CPG. The output of the `showvv -cpgalloc` command indicates whether the user space of the virtual volume is CPG-provisioned.

Use the **`createvv`**, **`growvv`**, and **`tunevv`** commands to create, grow, and tune the virtual volume with CPG-provisioned user space.

The following table displays the virtual volume `type` and the corresponding CLI commands for creating, growing and tuning. The `type` definition used matches the Type column of **`showvv`** output.

Snapshot space	**No**	**Not from a CPG**	**From a CPG**

User space			
Full Provisioned, from a CPG	Base - createvv - growvv - tunevv	Base - createvv - growvv - tunevv	Base, cpvv - createvv - growvv - tunevv
Thin Provisioned	N/A	N/A	Base, tpvv - createvv - growvv - tunevv

Add Command

addsnmpmgr

Syntax

```
addsnmpmgr [options <arg>] <manager_IP>
```

Description

Adds an SNMP manager to receive trap notifications.

Options

-alertclear {all|nodup|standard}

all

Enable both kinds of alert clear trap:

1. The original alert which state has changed, but with a severity of "clear" if the new state is acknowledged, fixed or autofixed.
2. The legacy alert state change `alertNotify` trap, with a `detailedMessage` of "Alert <id> changed from state <prev state> to <new state>"

Note that if an alert is removed through `removealert`, only type 2 trap above is generated.

nodup

Enable only the type 1 trap above.

Note that if an alert is removed through `removealert`, the type 2 trap above is generated instead.

standard

Enable only the type 2 trap above. This is the default.

-notify {all|nodup|standard}

all

Enables all trap notification types defined by the HPE 3PAR MIB.

nodup

Enables all trap notification types defined by the HPE 3PAR MIB. Send `alertNotify` trap only if no other trap type is sent for an event.

standard

Sends `alertNotify` type traps only. Default.

-p <port_number>

Specifies the port number where the SNMP manager receives traps.

Used together, the port number and IP address uniquely identify the SNMP manager. Use this option if the port number differs from the default of `162`.

-pw <password>

Specifies the SNMP manager access community string (password), using up to 32 alphanumeric characters. If unspecified, defaults to `public`. This option applies to SNMPv2 only.

-r <retry>

Specifies the number of times to send a trap (retry) if the SNMP manager is not available. Use an integer from 1 through 15. With none specified, defaults to 2 retries.

This option is no longer used and will be removed in a subsequent release.

-snmpuser <username>

Specifies the SNMPv3 user name using up to 31 characters. To enable for SNMPv3, use the **createsnmpuser** command. Configure the user name with the same credential on both the SNMPv3 manager and the HPE HPE 3PAR agent.

-t <timeout>

Specifies the number of seconds to wait before sending a trap (timeout). Use an integer from 1 through 300. With none specifies, the time defaults to 200 seconds. This option is no longer used and will be removed in a subsequent release.

-version <version>

Specifies the SNMP version supported. Use an integer value of `2` for SNMPv2 or `3` for SNMPv3. Defaults to 2.

Specifiers

`<manager_IP>`

Specifies the IP address of the host from which the manager runs. Use a valid IPv4 or IPv6 address. The IPv6 address is hexadecimal, case insensitive, and separated using colons, similar to the following example:

```
5def:2008:0:0:abcd:0:0:161a
```

Use a double colon (::) one time in an address to replace multiple fields of zeros. For example:

```
5def:2008:0:0:abcd::161a
```

Authority

- Super
- Any role granted the `snmpmgr_add` right

Usage

- Requires access to all domains.
- Does not support any form of name resolution. Specify IP addresses directly.
- Display the list of registered SNMP managers using the **showsnmpmanager** command.
- Change the SNMP user or version using the **setsnmpmanager** command.
- Change the SNMP3 passwords using the **setsnmpuserpw** command.
- Change the SNMP passwords using the **setsnmppw** command.
- Remove SNMP passwords using the **removesnmppw** command.
- Remove SNMP managers using the **removesnmpmgr** command.

Examples

Adding SNMPv2 manager IPv4 address 123.45.67.89 with the assigned password of *alpha1*:

```
cli% addsnmpmgr -pw alpha1 123.45.67.89
```

Adding SNMPv2 manager IPv6 address with port number and `nodup` notification:

```
cli% addsnmpmgr -notify nodup -p 9162 5def:2008:abcd::161a
```

Adding SNMPv3 manager IPv6 address with user:

```
cli% addsnmpmgr -version 3 -snmpuser joe 5def:2008:abcd::161b
```

Admit Commands

admithw

Syntax

```
admithw [options]
```

Description

Admits new hardware into the system. On discovery of any two-node HPE 3PAR system, **tunesys** starts automatically and redistributes existing volumes to use the new capacity. You can disable this facility using either the –notune option or by setting the AutoAdmitTune system parameter to no. On systems with more than two nodes, always run tunesys manually after disk installation.

Options

-checkonly

Performs passive checks only. Does not make any changes.

-f

Ignores errors and continue. Error messages still display.

-nopatch

Suppress the check for drive table update packages for new hardware enablement.

-tune

Always run tunesys to rebalance the system after discovering new disks.

-notune

Do not automatically run tunesys to rebalance the system after discovering new disks.

Specifiers

None.

Authority

Super, Service

Usage

- Requires access to all domains.
- Handles any nodes, disks, or cages added into the system.
- Verifies the presence of all expected hardware and handles all checks, including valid states, cabling, and firmware revisions.
- Handles creating system logical disks while adding and rebalancing spare chunklets.
- Allocates spares according to the algorithm specified by the SparingAlgorithm system parameter.
- If new disks are discovered, the set size for existing CPGs is recalculated. Changes to the CPG occur prior to any tunesys operation so that the affected LDs are automatically tuned.
- Checks for drive table patch updates unless you specify the –nopatch option.

- In addition, discovery of new disks in any combination can cause `tunesys` to start automatically and rebalance the system after the `admithw` command has completed.
- Automatic `tunesys` occurs under the following conditions:
  - With `admithw -tune`, rebalancing occurs on all systems regardless of the number of controller nodes.
  - On systems with two controller nodes, `tunesys` runs automatically. To suppress this behavior, use the system variable `AutoAdmitTune` with the following command structure: `cli% setsys AutoAdmitTune no`. `AutoAdmitTune` defaults to `yes`.
  - With `admithw -notune`, rebalancing does not occur after new discovery of new disks. In all circumstances, run `tunesys` as soon as possible after discovery of new disks.

Examples

Displaying a passive check of the system:

```
cli% admithw -checkonly
Checking nodes...
Checking volumes...
Checking system LDs...
Checking ports...
Checking state of disks...
Checking cabling...
Check complete.
```

admitpd

Syntax

`admitpd [option] [<world_wide_name>...]`

Description

Creates and admits PD definitions to enable the use of those disks.

Options

`-nold`

Do not use the PD (as identified by the `<world_wide_name>` specifier) for logical disk allocation.

`-nopatch`

Suppress the check for drive table update packages for new hardware enablement.

Specifiers

`<world_wide_name>`

Specifies the World Wide Name (WWN) of the PD to be admitted. With WWNs specified, admits only the specified PD. Otherwise, admits all available PDs.

Authority

- Super, Service
- Any role granted the `pd_admit` right

Usage

- Requires access to all domains.
- The HPE 3PAR OS cannot use PDs for storage until you admit the PDs into the system.

- To replace a failed disk that has had chunklets moved to spare space, use the `-nold` option when adding the replacement physical disk. Specifying `-nold` prevents the allocation of the newly added physical disk, allowing you to move the chunklets back to the new disk. After moving the chunklets back to the new disk, use the **setpd** command to enable logical disk allocation.
- To verify the admittance of physical disk definitions, enter the **showpd** command.
- With no WWN specified, or with all specified WWNs admitted, the command succeeded. If you specify multiple WWNs and not all can be admitted, the `admitpd` command fails.
- This command checks for drive table patch updates unless you specify the `-nopatch` option.

Examples

Admitting PDs in a 20-disk system:

```
cli% admitpd
20 disks admitted
```

admitrcopylink

Syntax

For remote copy over IP (RCIP), the syntax for the `admitrcopylink` command is as follows:

```
admitrcopylink <target_name> <N:S:P:IP_address>...
```

For remote copy over fibre channel (RCFC), the syntax for the `admitrcopylink` command is as follows:

```
admitrcopylink <target_name> <N:S:P:WWN>...
```

Description

Adds one or more links (connections) to a remote copy target system.

Options

None.

Specifiers

<target_name>

Specifies the **creatercopytarget** command.

<node:slot:port:IP_address>...

Specifies the node, slot, and port of the Ethernet port on the local system and an IP address of the peer port on the target system.

<node:slot:port:WWN>...

Specifies the node, slot, and port of the Fibre Channel port on the local system and World Wide Name (WWN) of the peer port on the target system.

Authority

- Super, Edit
- Any role granted the `rcopylink_admit` right

Usage

- IP targets are made up of pairs composed of the node, slot, and port of the ethernet port on the local system and an IP address of the peer port on the target system.
- FC targets are made up of sets with the node, slot, and port of the FC port on the local system and WWN of the peer port on the target system.
- This command concludes by returning a list of one or more links to be admitted.
  - For IP targets, the list consists of pairs composed of the node containing the Ethernet port on the local system and an IP address of the peer port on the target system.
  - For FC targets, the list consists of sets with the node, slot, and port of the FC port on the local system and the WWN of the peer port on the target system.
- This command requires the HPE 3PAR Remote Copy Software license. Contact your local service provider for more information.
- Access to all domains is required to run this command.
- The transport layer (RCIP or RCFC) is set using the **creatercopytarget** command.
- The specifier `<node:IP_address>` has been deprecated.

Examples

The following example adds a link on System2 (targetname), node 1. The IP address 193.1.2.11 specifies the address on the target system:

```
cli% admitrcopylink System2 1:2:1:193.1.2.11
```

The following WWN creates an RCFC link to target System2, which connects to the local 5:3:2 (N:S:P) in the target system:

```
cli% admitrcopylink System2 5:3:2:20010002AC000060
```

admitrcopytarget

Syntax

```
admitrcopytarget <target_name> <mode> <group_name>
[<pri_vv_name:sec_vv_name> ...]
```

Description

The `admitrcopytarget` command adds a target to a remote-copy volume group.

Options

None.

Specifiers

<target_name>

Specifies the name of the target that was previously created with the **creatercopytarget** command.

<mode>

Specifies the mode of the target as either synchronous (sync), asynchronous periodic (periodic), or asynchronous streaming (async).

<group_name>

Specifies the name of the existing remote copy volume group created with the **creatercopygroup** command to which the target will be added.

<pri_vv_name:sec_vv_name>

Specifies the mapping between the names of a volume in the primary group and the corresponding volume on the added target. This specifier must be included for every volume in the primary volume group. This specifier is required only if the group contains volumes.

Authority

- Super, Edit
- Any role granted the `rcopytarget_admit` right

Usage

- Functionality of this command requires the HPE 3PAR Remote Copy Software license.
- Contact your local service provider for more information.
- Access to all domains is required to run this command.
- A primary<->secondary volume mapping must be provided for each volume currently in the group.

Examples

In the following example, the `admitrcopytarget` command adds the target target1 in synchronous mode to volume group Group1.

```
cli% admitrcopytarget target1 sync Group1
```

admitrcopyvv

Syntax

```
admitrcopyvv [options] <VV_name>[:<snapname>][,<psnapname>] <group_name>
<target_name>:<sec_VV_name>...
```

Description

The `admitrcopyvv` command adds an existing virtual volume to an existing Remote Copy volume group.

Options

-pat

Specifies that the `<VV_name>` is treated as a glob-style pattern, and that all Remote Copy volumes matching the specified pattern are admitted to the Remote Copy group. When this option is used, the `<sec_VV_name>` and `<snapname>` (if specified) are also treated as patterns. It is required that the secondary volume names and snapshot names can be derived from the local volume name by adding a prefix, suffix or both. `<snapname>` and `<sec_VV_name>` must take the form prefix@vvname@suffix, where @vvname@ resolves to the name of each volume that matches the `<VV_name>` pattern.

-createvv

Specifies that the secondary volumes must be created automatically. This specifier cannot be used when starting snapshots (`<VV_name>:<snapname>`) are specified.

If primary volumes are of DECO (dedup and compr) type, and the target CPG is not capable of DECO, then TPVV volumes are created in the secondary system.

-nowwn

When used with `-createvv`, it ensures that a different WWN is used on the secondary volume. Without this option, `-createvv` will use the same WWN for both primary and secondary volumes.

-nosync

Specifies that the volume must skip the initial sync. This is for the admission of volumes that have been presynced with the target volume.

This specifier cannot be used when starting snapshots (<VV_name>:<snapname>) are specified.

Specifiers

<VV_name>

Specifies the name of the existing virtual volume to be admitted to an existing Remote Copy volume group that was created with the **createcopygroup** command.

<psnapname>

Valid only for the SLD/3DCPP configuration. An optional read-only snapshot <psnapname> for periodic target can be specified with <snapname> separated by ",". <snapname> is considered the snapshot for the sync target. Only automatically created snapshots are allowed.

<snapname>

An optional read-only snapshot <snapname> can be specified along with the virtual volume name <VV_name>. This snapshot is a starting snapshot. When the group is started, a full sync is not performed. Instead, for synchronous groups, the volume will synchronize deltas between <snapname> and the base volume. For asynchronous periodic groups, the volume will synchronize deltas between <snapname> and a snapshot of the base. For asynchronous streaming groups, the volume will synchronize deltas between <snapname> and the base volume.

<group_name>

Specifies the name of the existing Remote Copy volume group created with the **createcopygroup** command, to which the volume will be added.

<sec_vv_pattern>

Specifies the pattern that is used to create the secondary virtual volume on the target system.

<target_name>:<sec_VV_name>

The target name associated with this group, as set with the **createcopygroup** command. The target is created with the **createcopytarget** command. <sec_VV_name> specifies the name of the secondary volume on the target system. One <target_name>:<sec_VV_name> must be specified for each target of the group.

<pattern>

Specifies a glob-style pattern. If this specifier is not used, the <VV_name> specifier must be used. (For more information, enter clihelp sub,globpat or see **Glob-style patterns**).

Authority

- Super, Edit
- Any role granted the rcopyvv_admit right .

Usage

- Functionality of this command requires the HPE 3PAR Remote Copy Software license. Contact your local service provider for more information.
- A secondary volume mapping must be provided for each target in the group.
- The virtual volume and the Remote Copy group must be in the same domain or both in no domain.
- The -nosync option will skip the full synchronization of this volume once the Remote Copy group is started. On the primary system, the showrcopy SyncStatus (see the **showrcopy** command) will appear as "New-PreSynced". On the secondary, the SyncStatus will appear as "New".

- When a starting snapshot is specified, the snapshot will not be automatically deleted after the group is started and resynced.
- When a starting snapshot is specified, any changes between the specified snapshot and the point when the group is started will be resynced. This is done instead of a full synchronization. On the primary system, the `showrcopy SyncStatus` (see the **showrcopy** command) will appear as "New-SyncFromSnap". On the secondary, the SyncStatus will appear as "New".
- If the `-nosync` option is not used and a starting snapshot is not specified, the volume will undergo a full synchronization when the group is started. The `showrcopy SyncStatus` (see the **showrcopy** command) will appear as "New" on both primary and secondary systems.
- If a group's target has the mirror_config policy set and the group is a primary group, then this command is mirrored to that target and the volume is added to the secondary volume group. If the policy is set and the group is a secondary, this command fails.
- For multiple target Remote Copy groups, and adding the primary/secondary mapping on the target systems, the relevant mapping between the target systems will be added to the backup groups.
- With the `-createvv` option, if the VV has an associated VLUN (host is exported), VLUN is also created on the secondary. If the id is available, the VLUN will have the same id as the primary.

Examples

In the following example, the `admitrcopyvv` command adds the volume vv1 to the primary volume group Group1. At the same time, it adds the secondary volume vv1_remote on the target system System2 to the corresponding secondary volume group (Group1.r<sys_ID>), which was previously created when the **creatercopygroup** command was issued:

```
cli% admitrcopyvv vv1 Group1 System2:vv1_remote
```

In the following example, the `admitrcopyvv` command adds all volumes beginning with tpvv to the primary volume group (Group1). At the same time, it adds the secondary volumes, which are named the same as the primary volumes (as specified by @vvname@), on the target system System2 to the corresponding secondary volume group (Group1.r<sys_ID>):

```
cli% admitrcopyvv -pat tpvv* Group1 System2:@vvname@
```

In the following example, the `admitrcopyvv` command adds the volume vv1 to the primary volume group (Group1). At the same time, it creates and adds the volume vv1.r on the target system InServ1_in to the corresponding secondary volume group that was previously created when the **creatercopygroup** command was issued:

```
cli% admitrcopyvv -createvv vv1 Group1 InServ1_in:vv1.r
```

In the following example, the `admitrcopyvv` command adds the volume vv1 to the primary volume group (Group1). At the same time, it adds the volume vv1.r on the target system InServ1_in to the corresponding secondary volume group that was previously created when the **creatercopygroup** command was issued. It also specifies that the volume must not undergo an initial synchronization when the group is started:

```
cli% admitrcopyvv -nosync vv1 Group1 InServ1_in:vv1.r
```

In the following example, the `admitrcopyvv` command adds the volume vv1 to the primary volume group (Group1). At the same time, it adds the volume vv1.r on the target system InServ1_in to the corresponding secondary volume group that was previously created when the **creatercopygroup** command was issued. It also specifies that the volume must be synced using rosv_vv1 as the starting snapshot:

```
cli% admitrcopyvv vv1:rosv_vv1 Group1 InServ1_in:vv1.r
```

In the following example, the `admitrcopyvv` command adds all the volumes that start with the name testvv to the primary volume group (Group1). At the same time, it adds all the volume testvv*.r on the target system InServ1_in to the

corresponding secondary volume group that was previously created when the **createcopygroup** command was issued. It also specifies that each volume must be synchronized using rosv_testvv* as the starting snapshot:

```
cli% admitrcopyvv -pat testvv*:rosv_@vvname@ Group1 InServ1_in:@vvname@.r
```

admitvv

Syntax

```
admitvv [-domain <domain>] <vvname>:<WWN> [<vvname>:<WWN> ...]
```

```
admitvv [-domain <domain>] <vvname:<WWN>:<New WWN> [<vvname:<WWN>:<New
WWN>] ...
```

Description

The admitvv command creates and admits remotely exported virtual volume definitions to enable the migration of these volumes. The newly created volume will have the WWN of the underlying remote volume.

Options

-domain <domain>

Creates the admitted volume in the specified domain. The default is to create it in the current domain, or no domain if the current domain is not set.

Specifiers

<vvname>

Specifies the local name that must be given to the volume being admitted (up to 31 characters in length).

<WWN>

Specifies the World Wide Name (WWN) of the remote volumes to be admitted.

<New WWN>

Specifies the World Wide Name (WWN) for the local copy of the remote volume. If the keyword "auto" is specified the system automatically generates a WWN for the virtual volume.

Authority

- Super
- Any role granted the vv_admit right

Usage

- You need access to all domains to run this command.
- The volume admitted using this command can be exported to hosts. However, all the data access to the volume will be directed to the external LUN.
- The volume admitted using this command can be imported into local storage system by using the **importvv** command.

Examples

The following example admits two remote VVs that have been exported to the storage system in preparation for migration:

```
cli% admitvv migvv.0:50002AC00037001A migvv.1:50002AC00047001A
2 VVs admitted
```

Backup Command

backupfsconf

Syntax

```
backupfsconf [-fpg <fpgname>] <vfs>
```

Description

The `backupfsconf` command creates a configuration backup for a virtual file system. Configuration backup artifact is created with a standard name in the admin file store under the virtual file server in a directory configbackup. There will be one backup file created per virtual file server. The subsequent backup will overwrite the backup artifact.

Options

`-fpg <fpgname>`

Specifies the file provisioning group (FPG) for the virtual file server.

Specifiers

`<vfs>`

Specifies the name of the virtual file server.

Authority

- Super, Service.
- Any role granted `fsconf_backup` right

Usage

- Access to all domains is required to run this command.
- The Config backup artifact is created in the ".admin/configbackup" folder under the virtual file server (VFS). If there is not adequate space available to create the artifact the backup creation will fail.
- HPE recommends that you create the configuration backup before creating a data backup. The same must be applied when creating a restore.
- The following actions take place when creating a configuration backup or restore:
  - Backups of all the file stores, VFS IP, Share information of NFS, CIFS, and Object/Certificates, AC configuration, and Quota information.
  - Restores each of the backup points maintaining a restoration order.
- The following manual steps must be taken when creating a configuration backup:
  - When the configuration backup is finished, use NDMP or other means to back up the config backup folder located in the ".admin" folder (.admin/configbackup).
  - This folder contains the backup artifact.
  - Manually set up the FPG and the VFS, and then restore the configbackup folder under the VFS in the ".admin" file store (the folder containing the backup artifact).
  - Folder creation for each of the CIFS & NFS shares is handled by the NDMP data restore.
  - Manually configure the clusterwide AV VSE data before doing a configuration restore.

See **restorefsconf** for additional information on configuration restore.

Examples

The following example creates a configuration backup on FPG named "samplefpg" under virtual server named "samplevfs":

```
cli% backupfsconf -fpg samplefpg samplevfs
```

Cancel Command

canceltask

Syntax

```
canceltask [option] {all|<task_ID>...}
```

Description

The `canceltask` command cancels one or more tasks.

Options

`-f`

Forces the command. The command completes the process without prompting for confirmation.

Specifiers

`all`

Cancels all active tasks. If not specified, a task ID(s) must be specified.

`<task_ID>`

Cancels only tasks identified by their task IDs. Task_ID must be an unsigned integer within 1-29999 range.

Authority

- Super, Service, Edit
- Any role granted the `task_cancel` right

Usage

- The `canceltask` command can return before a cancellation is completed. Thus, resources reserved for a task might not be immediately available. This can prevent actions like restarting the canceled task. Use the `waittask` command to ensure orderly completion of the cancellation before taking other actions. See **`waittask`** for more details.
- The `-restart` option associated with the `tunevv` command enables a canceled `tunevv` task to be resumed. See **`tunevv`** for more information.

Examples

The following example shows how to cancel a task using the task ID:

```
cli% canceltask 1
Are you sure you want to cancel task 1?
select q=quit y=yes n=no: y
```

Change Command

changedomain

Syntax

```
changedomain [<domain>]
```

Description

The `changedomain` command changes the currentdomain CLI environment parameter.

Options

None.

Specifiers

<domain>

Name of the domain to be set as the working domain for the current CLI session. If the <domain> parameter is not present or is equal to -unset, then the working domain is set to no current domain.

Authority

Any role in the system

Usage

This command is equivalent to `setclienv currentdomain <domain>`.

Examples

The following example shows how to change the currentdomain Win:

```
cli% changedomain Win
```

The following example shows two methods to unset the currentdomain:

```
cli% changedomain -unset
cli% changedomain
```

Check Commands

checkhealth

Syntax

```
checkhealth [<options> | <Component>...]
```

Description

The `checkhealth` command checks the status of system hardware and software components, and reports any issues.

Options

`-lite`

Performs a minimal health check.

`-svc`

Performs a thorough health check. This is the default option.

`-full`

Performs the maximum health check. This option cannot be used with the -lite option.

`-list`

List all components that will be checked.

`-quiet`

Do not display which component is being checked. Do not display the footnote with the -list option.

`-d`

Displays detailed information regarding the status of the system.

`-pelcheck`

Used for the cabling component to show a subset of errors that may have been caused by SAS PEL errors. You cannot use this option with the `-recommended` option.

`-recommended`

Perform additional checks against the recommended configuration. This option is applied only to the cabling component. A limit of eight additional errors can be reported by the cabling component.

Specifiers

`<Component>`

Indicates the component to check. Use `--list` option to get the list of components.

Authority

- Any role in the system
- Users with Edit and Browse roles will not evaluate the system as fully as those with Super and Service roles, and will only evaluate items available to the domains in which the user has privileges.

Usage

None.

Examples

In the following example, the overall health status of all the components is checked:

```
cli% checkhealth
Checking alert
Checking ao
Checking cabling
Checking cage
Checking cert
Checking dar
Checking date
Checking file
Checking fs
Checking host
Checking ld
Checking license
Checking network
Checking node
Checking pd
Checking pdch
Checking port
Checking qos
Checking rc
Checking security
Checking snmp
Checking task
Checking vlun
Checking vv
Checking sp
Component -----Summary Description----- Qty
Alert     New alerts                     8
Cabling   Bad SAS connection              1
Cage      Cages not on current firmware   8
LD        LDs with reduced availability   3
-------------------------------------------
       4 total                          20
```

In the following example, the file, ld, and task components are checked and the detailed information is displayed:

```
cli% checkhealth -d file ld task
Checking file
Checking ld
Checking task
Component --------Summary Description-------- Qty
File      Nodes with Dump or HBA core files     1
Task      Failed Tasks                          2
-------------------------------------------------
       2 total 3

Component -Identifier- ----Detailed Description----
File      node:1       Dump or HBA core files found
Task      Task:529     Failed Task
Task      Task:530     Failed Task
-------------------------------------------------
       3 total
```

checkld

Syntax

```
checkld [options] <LD_name>...
```

Description

Executes consistency checks of data on LDs in the event of an uncontrolled system shutdown, and repairs inconsistent LDs (optional).

Options

-y|-n

Specifies modifying any errors found (-y) or leaves the error unmodified (-n). Defaults to unmodified (-n).

-progress

Specifies polling the sysmgr to get an ldck report.

-recover <pdid:pdch>

Attempts to recover the chunklet, which is specified by physical disk (<pdid>) and position on the disk (<pdch>). Requires use of -y.

-rs <raidset_number>

Checks the specified RAID set only.

Specifiers

<LD_name>...

Checks the integrity of the specified LD. Repeat this specifier to execute validity checks on multiple LDs.

Authority

- Super, Service
- Any role granted the ld_check right

Usage

- Requires access to all domains.
- Repairing LDs refers to making LDs consistent.
- Defines consistency for RAID-1 as: all mirrors in the set have the same data.
- Consistency for RAID-5 or RAID-6 means that parity is consistent with the data in the set.
- Using the -recover option allows one LD only and requires use of the-y option.
- Enter the checkld command on any LD, whether started or not.

Examples

Displays a validity check of LD tp-6-sa-0.0, with the -progress option to be informed of the check's progress:

```
cli% checkld -n -progress tp-6-sa-0.0
Performing a kernel level consistency check, will not try to fix inconsistencies
  ld tp-6-sa-0.0 has 5 rsets, checked 0, inconsistent 0
  ld tp-6-sa-0.0 has 5 rsets, checked 0, inconsistent 0
  ld tp-6-sa-0.0 has been checked
```

Attempts to recover chunklet 1032:10 of LD vv.adm.0, and displays progress output:

```
cli% checkld -y -progress -recover 1032:10 vv.adm.0
```

checkpassword

Syntax

```
checkpassword [<user>]
```

Description

The `checkpassword` command prompts for the user's password and then displays the steps the system uses to authenticate the user and determine the user's role. The information includes whether the user is local to the system or authenticated and authorized based on the configuration parameters set with the **`setauthparam`** command for use with LDAP.

Options

None.

Specifiers

`<user>`

If the `<user>` parameter is not specified, then the current user is used.

Only users with Super role with access to all domains can specify `<user>` names other than their own.

Authority

- Super, Edit, Browse, Service
- Any role granted the `password_checkany` right can check the password of any user.
- Any role granted the `password_checkown` right can only check their own password.

Usage

- The output of the `checkpassword` command is based on current authentication and authorization parameters and might differ from the user's actual authorization level if the **`setauthparam`** command has been used to change parameters or data in the LDAP server has changed since the user last logged in.
- The **`showuserconn`** command can be used to verify the authorization levels assigned at login.

Examples

The following example displays the authentication and authorization of user User1:

```
cli% checkpassword User1
password:
+ attempting authentication and authorization using system-local data
user User1 is authenticated and authorized
```

The output of the `checkpassword` command is a series of information statements, each starting with a plus sign ("+") that indicates the steps the system is using to authenticate the user and determine the role.

The last line of output shows the summary of the steps with the user either being unauthenticated or authenticated and authorized.

checkpd

Syntax

```
checkpd scrub|diag [options ] <pd_ID>...
```

Description

The `checkpd` command executes surface scans or diagnostics on physical disks (PD).

Options

The following options can only be used with the scrub specifier:

-ch <number>

To scan a specific chunklet rather than the entire disk.

-count <number>

To scan a number of chunklets starting from -ch.

The following options can only be used with the diag specifier:

-path <p>

Specifies a PD path as a, b, both, or system.

-test <type>

Specifies read, write, or verify test diagnostics. If no type is specified, the default is read. The "verify" diag test option fills the test range with a pattern then reads back the data and verifies that the data is correct.

-iosize <size>

Specifies I/O size, valid ranges are from 1s to 1m. If no size is specified, the default is 128k.

-threads <num>

Specifies number of I/O threads, valid ranges are from 1 to 4. If the number of threads is not specified, the default is 1.

-time <secs>

Indicates the number of seconds to run, from 1 to 36000. The -time option takes precedence if used with other options such as -total or -range.

-total <size>

Indicates total bytes to transfer per disk. If a size is not specified, the default size is 1g.

-retry <number>

Specifies the total number of retries on an I/O error. If the number of retries is not specified, the default is 4.

-range <size>

Limits diagnostic regions to a specified size, from 2m to 2g.

Specifiers

scrub

Scans one or more chunklets for media defects.

diag

Performs read, write, or verify test diagnostics.

<pd_ID>...

The ID of the PD to be checked. Only one `pd_ID` can be specified for the "scrub" test.

Authority

- Super, Service
- Any role granted the `pd_check` right

Usage

- Access to all domains is required to run this command.
- The <size> specifier can include a letter to indicate units:
  - g = gigabytes (2^30)
  - t = terabytes (2^40)

- p = petabytes (2^50)
- m = 1048576 bytes
- k = 1024 bytes
- s = 512 bytes

- I/O errors will be reported even if the eventual I/O succeeds due to retries.
- Up to 40 PD IDs can be specified for the diag test type.

Examples

In the following example, chunklet 400 on PD 1 is scanned for media defects:

```
cli% checkpd scrub -ch 400 1
2015-06-09 04:41:37 MDT {22401} PD Scrubber Start: max chunk id 400 pd id 1 starting chunk id 400 system cycle
168hr chunk cycle 60s log level 5
2015-06-09 04:41:37 MDT {22401} Open system device...
2015-06-09 04:41:37 MDT {22401} Attach to system manager...
2015-06-09 04:41:37 MDT {22401} Waiting for system manager ready...
2015-06-09 04:41:37 MDT {22401} PD Scrubber 2.0 started
2015-06-09 04:41:37 MDT {22401} Thread 1863816960 started for node 0
2015-06-09 04:41:47 MDT {22401} Finished scanning pd 1 ch 400 status 1
2015-06-09 04:41:47 MDT {22401} No media errors detected
```

checkport

Syntax

```
checkport [options <arg>] <node:slot:port>
```

Description

The `checkport` command performs loopback tests on Fibre Channel ports.

Options

`-time <seconds_to_run>`

Specifies the number of seconds for the test to run using an integer from 0 to 300.

`-iter <iterations_to_run>`

Specifies the number of times for the test to run using an integer from 1 to 1000000.

Specifiers

`<node:slot:port>`

Specifies the port to be tested.

`node`

Specifies the node using a number from 0 through 7.

`slot`

Specifies the PCI slot in the specified node. Valid range is 0 - 9.

`port`

Specifies the port using a number from 1 through 4.

Authority

- Super, Service
- Any role granted the `port_check` right

Usage

- Access to all domains is required to run this command.
- When both the `-time` and `-iter` options are specified, the first limit reached terminates the program. If neither are specified, the default is 1,000 iterations. The total run time is always limited to 300 seconds even when not specified.
- The default loopback is an ELS-ECHO sent to the HBA itself.

Examples

In the following example, the loopback test is performed on port 0:0:1 a total of five times:

```
cli% checkport -iter 5 0:0:1
Starting loopback test on port 0:0:1
Port 0:0:1 completed 5 loopback frames in 0 seconds
Passed
```

checkrclink

Syntax

```
checkrclink subcommand [options] <N:S:P> [<dest_addr>] [<port>]
```

Description

The `checkrclink` command performs a connectivity, latency, and throughput test between two connected HPE 3PAR storage systems.

Options

`-time <secs>`

Specifies the number of seconds for the test to run using an integer from 300 to 172800. If not specified this defaults to 172800 seconds (48 hours).

The -time <secs> value for the server must be longer than the <time> for the client, otherwise the test starts to produce invalid results when the server exits early.

`-fcip`

Specifies if the link is running over fcip. Should only be supplied for FC interfaces.

Specifiers

`<N:S:P>`

Specifies the interface from which to check the link, expressed as node:slot:port.

`<dest_addr>`

Specifies the address of the target system (for example, the IP address).

`<time>`

Specifies the test duration in seconds.

Specifies the number of seconds for the test to run using an integer from 300 to 172800.

`<port>`

Specifies the port on which to run the test. If this specifier is not used, the test automatically runs on port 3492.

Subcommands

```
checkrclink startclient [options] <N:S:P> <dest_addr> <time> [<port>]
```

```
checkrclink stopclient <N:S:P>
```

```
checkrclink startserver [options] <N:S:P> [<dest_addr>] [<port>]
checkrclink stopserver <N:S:P>
checkrclink portconn <N:S:P>
```

startclient

Starts the link test.

stopclient

Stops the link test.

startserver

Starts the server.

stopserver

Stops the server.

portconn

Uses the Cisco Discovery Protocol Reporter to show display information about devices that are connected to network ports. Requires CDP to be enabled on the router.

Authority

- Super, Service
- Any role granted the `rclink_check` right

Usage

- Access to all domains is required to run this command.
- Do not check the links if CPU usage is already close to 100%. The link check temporarily increases CPU usage.
- The remote copy ports must be configured through the **controlport** command before this test can be carried out; however, this test can be performed before links are admitted into the remote copy target system. The test will run on IP or FC links. This test must be performed in two stages.
- To run a test between System A and System B:

  1. Start the server on System B, specifying the `<N:S:P>` for System B and the <dest_addr> for System A.

     ```
     checkrclink startserver <N:S:P> <dest_addr>
     ```

     If you are running the `checkrclink` test on RCFC links, note the IP address displayed in the output.

  2. Start the client on the target system (System A), specifying the `<N:S:P>` for System A; the <dest_addr> for System B (as an IP address for both RCIP and RCFC links); and the duration of the test in seconds.

     ```
     checkrclink startclient <N:S:P> <dest_addr> <time>
     ```

  3. (Optional) To end the test before it completes, on System A, issue the following command:

     ```
     checkrclink stopclient
     ```

  4. After the test is complete, on System B, issue the following command:

```
checkrclink stopserver
```

- To run a test between storage System A and B:
  1. Start the server on storage System B, specifying the relevant <N:S:P>.

     For IP ports only the IP address of storage System A should be specified.

     If you are running the checkrclink test on RCFC links, note the IP address displayed in the output.

  2. Start the client on storage System A, specifying the relevant <N:S:P>, the IP Address of storage System B, and the duration of the test in seconds. For FC Links, the <dest_addr> is also specified as an IP address, the required IP address is displayed in the output of the startserver subcommand.
- By default this test will use port number 5001, however there is an optional parameter to allow the use to specify their own port number.

Examples

The following example starts a client test on port 0:0:1 to IP address 10.100.10.10 with a duration of 60 seconds:

```
cli% checkrclink startclient 0:0:1 10.100.10.10 60
```

checksnmp

Syntax

checksnmp

Description

The checksnmp command allows a user to send an SNMPv2 test trap to the list of managers as specified in showsnmpmgr.

Options

None.

Specifiers

None.

Authority

- Super, Service
- Any role granted the snmp_check right

Usage

- Managers must be registered first using the **addsnmpmgr** command.
- The test trap is in the form of an alert. Notify trap with the following varbinds:
  - component = "test_trap"
  - details = "This is a test trap sent from Storage System <system name>, Serial Number <s/n>"
  - nodeID = ID of the node from which the trap was issued
  - severity = debug (6)

- timeOccurred = time when the command was issued
- id = 4294967295
- messageCode= 4294967295
- state = autoFixed (5)

Examples

```
cli% checksnmp
Trap sent to the following managers:
  192.168.17.10:162
  192.168.17.111:1000
```

The following is an example for the case when there are no SNMP managers configured.

```
cli% checksnmp
Cannot send test trap as no manager exists
```

checkvv

Syntax

```
checkvv [options] <VV_name|pattern|VV_Set>...
```

Description

The `checkvv` command executes validity checks of VV administration information in the event of an uncontrolled system shutdown and optionally repairs corrupted virtual volumes. It can also be used to check one or more virtual volumes to preview the potential space savings through deduplication technology.

Options

`-y|-n`

Specifies that if errors are found they are either modified so they are valid (`-y`) or left unmodified (`-n`). If not specified, errors are left unmodified (`-n`).

`-offline`

Specifies that VVs specified by `<VV_name>` be offline before validating the VV administration information. The entire VV tree will be offline if this option is specified.

`-fixsd`

Specifies that VVs specified by `<VV_name>` be checked for compressed data consistency. The entire tree will not be checked; only those VVs specified in the list will be checked.

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

`-dedup_dryrun`

Launches a dedup ratio calculation task in the background that analyzes the potential space savings with HPE 3PAR Deduplication technology if the VVs specified were in a same deduplication group. The VVs specified must be TPVVs only.

-compr_dryrun

Launches a compression ratio calculation task in the background that analyzes the potential space savings with HPE 3PAR Compression technology of specified VVs. Specified volumes can be TPVVs, TDVVs, fully provisioned volumes, and snapshots.

-dedup_compr_dryrun

Launches background space estimation task that analyzes the overall savings of converting the specified VVs into a compressed TDVVs. Specified volumes can be TPVVs, TDVVs, compressed TPVVs, fully provisioned volumes, and snapshots.

Specifiers

<VV_name|pattern|VV_Set>

Requests that the integrity of the specified VV is checked. This specifier can be repeated to execute validity checks on multiple VVs. Only base VVs are allowed.

The patterns are glob-style (shell-style) patterns. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**).

Authority

- Super, Service, Edit
- Any role granted the `vv_check` right

Usage

- Access to all domains is required to run this command.
- If `-n` initiated without `-offline` for the VVs which are started, there will likely be spurious inconsistencies due to I/O activity.

Examples

The following example offlines and fixes virtual volume test1:

```
cli% checkvv -y -offline test1
```

The following example launches a deduplication estimation task on a VV set set1:

```
cli% checkvv -dedup_dryrun set:set1
```

CLI Command

cli

Syntax

```
cli [options] [commands]
```

Description

The `cli` command provides a means to set up your CLI session or to enter directly into a CLI shell.

Options

`-v`

Displays the CLI client version.

`-b`

Displays the CLI client build level.

`-h`

Displays help for the `cli` command.

`-tpdportnum <portnum>`

Uses `<portnum>` for socket instead of default. For SSL sockets default port is 5783, otherwise default port is 5782.

`-sockssl`

Uses SSL for socket connection. SSL will be used if either this option is specified or the TPDSOCKSSL environment variable is set. This option cannot be used in conjunction with the `-nosockssl` option or the TPDNOSOCKSSL environment variable set.

`-nosockssl`

Do not use SSL for socket connection. SSL will not be used if either this option is specified or the TPDNOSOCKSSL environment variable is set. This option cannot be used in conjunction with the `-sockssl` option or the TPDSOCKSSL environment variable set.

`-certfile <certfile>`

Specifies the certificate file which is used to validate authenticity of CLI server. Certificate in this file must be in PEM (Privacy Enhanced Mail) format. The default file is `<certdir>/cert`. See `-certdir` option.

NOTE: The Windows client requires that paths containing spaces must be enclosed in quotation marks.

`-certdir <certdir>`

Specifies the directory in which to save the certificate exception file ("excp"). When CLI server certificate is not verified by the certificate file ("cert"), the cli looks for its information in the certificate exception file. If not found, the cli prompts user to accept and save it in this file. Next time the cli connects to the same CLI server without prompt. The default is `$HOME/.hpe3par` on UNIX based systems and `%USERPROFILE%\.hpe3par` on Windows.

NOTE: The Windows client requires that paths containing spaces must be enclosed in quotation marks.

`-nocertprompt`

Do not prompt user for CLI server certificate which is not verified. The cli exits with an error message. The default is to prompt user.

-sys <sysname>

Connects to system named <sysname>. If this option is not used cli uses the value of the TPDSYSNAME environment variable. If the TPDSYSNAME environment variable is not set, the cli prompts the user for the system name.

-cmdloop

Entesr interactive command loop after executing commands specified on command line.

-hafter <nlines>

Prints a header after <nlines> of data. (Not available before version 2.2)

-nohdtot

Do not print header and totals. (Not available before version 2.2)

-csvtable

Prints table data as Comma Separated Values (CSV). (Not available before version 2.2)

-listdom

Includes domain column where relevant. (Not available before version 2.2.3)

-clientname

Name of the client application. (Not available before version 3.1.1)

-clientver

Versions of the client application. (Not available before version 3.1.1)

-clienthost

Name of the host the client application runs on. (Not available before version 3.1.1)

-conntimeout

The maximum time in seconds to establish the connection to the CLI server. It defaults to 36 seconds if not set. (Not available before version 3.1.2)

-enherror

If set, perform enhanced error handling. (Not available before version 3.1.3)

-noenherror

If set, do not perform enhanced error handling. (Not available before version 3.1.3)

-matchbulkobjs

If set, perform operations on bulk VV patterns.

-nobanner

If set, suppresses CLI login banner.

Specifiers

[commands]

Any CLI command. This specifier is not required when issuing the cli command. If commands are specified, the CLI exits after executing the commands unless the -cmdloop option is specified. If no commands are specified in the command line, CLI enters the command loop.

Environment Variables

TPDSOCKSSL

If set, use an SSL socket.

TPDNOSOCKSSL

If set, do not use an SSL socket.

TPDCERTFILE

Certificate file. Overridden by the `-certfile` option.

NOTE: The Windows client requires that paths containing spaces must be enclosed in quotation marks.

TPDCERTDIR

Directory in which to save certificate exception file. Overridden by the `-certdir` option.

NOTE: The Windows client requires that paths containing spaces must be enclosed in quotation marks.

TPDNOCERTPROMPT

If set, do not prompt user for CLI server certificate which is not verified. The cli exits with an error message.

TPDSYSNAME

System to connect to. Overridden by the `-sys` option.

TPDCACHEDIR

Directory in which to cache the client bytecode.

TPDSTARTFILE

File containing code that is executed before any command line commands are executed or the interactive command loop is entered. This file normally contains customization scripts.

TPDHAFTER

Number of lines after which the header is printed again. Overridden by the `-hafter` option. (Not available before version 2.2)

TPDNOHDTOT

If set, do not print headers and totals. (Not available before version 2.2)

TPDCSVTABLE

If set, print table data as comma separated values (CSV). (Not available before version 2.2)

TPDLISTDOM

If set, include domain column where relevant. (Not available before version 2.2.3)

TPDCLIENTNAME

Client application name. Overridden by the `-clientname` option. (Not available before version 3.1.1)

TPDCLIENTVER

Client application version. Overridden by the `-clientver` option. (Not available before version 3.1.1)

TPDCLIENTHOST

The host the client application runs on. Overridden by the `-clienthost` option. (Not available before version 3.1.1)

TPDCONNTIMEOUT

The maximum time in seconds to establish the connection to the CLI server. Overridden by the `-conntimeout` option. This timer defaults to 36 seconds if not set. (Not available before version 3.1.2)

TPDFORCE

If set, do not ask for confirmation.

TPDCLIENHERROR

If set to 1, perform enhanced error handling. If set to 0, do not perform enhanced error handling. (Not available before version 3.1.3)

TPDMATCHBULKOBJS

If set, perform operations on bulk VV patterns.

Authority

Super, Edit, Browse, Service

Usage

None.

Examples

None.

Compact Commands

compactcpg

Syntax

```
compactcpg [options] <CPG_name>...
compactcpg -pat [options] <pattern>...
```

Description

The `compactcpg` command consolidates logical disk space in Common Provisioning Groups (CPGs) into as few logical disks as possible, allowing unused logical disks to be removed and their space reclaimed.

Options

`-pat`

Compacts CPGs that match any of the specified patterns. This option must be used if the pattern specifier is used.

`-waittask`

Waits for any created tasks to complete.

`-trimonly`

Removes unused logical disks after consolidating the space. This option will not perform any region moves.

`-nomatch`

Removes only unused logical disks whose characteristics do not match the growth characteristics of the CPG. Must be used with the `-trimonly` option. If all logical disks match the CPG growth characteristics, this option has no effect.

`-dr`

Specifies that the operation is a dry run, and the tasks are not actually performed.

`-f`

Does not ask for confirmation before compacting the logical disks.

Unless the `-f` option is specified, the command asks for confirmation before compacting each CPG.

Specifiers

`<pattern>...`

Specifies a glob-style pattern. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.) This specifier can be repeated to compact multiple CPGs. If this specifier is not used, the <CPG_name> specifier must be used.

`<CPG_name>...`

Specifies the name of the CPG. Multiple CPGs can be specified.

Authority

- Super, Edit
- Any role granted the `cpg_compact` right

Usage

- Logical disks that are consolidated must have the same owner and backup nodes.
- If one logical disk exists with a different owner and backup node, that logical disk will be compacted individually.
- Use -dr (dry run option) to determine whether or not to use the -nomatch option. The -nomatch option should be used to preserve space in matching logical disks so that a future `compactcpg` can move data from non-matching logical disks into those matching logical disks.

Examples

The following example shows how to compact a single CPG named testcpg and remove any unused logical disks after consolidating the space:

```
cli% compactcpg -trimonly testcpg
Are you sure you want to compact CPG 'testcpg' ?
select q=quit y=yes n=no: y
Task 612 started
```

compactld

Syntax

```
compactld [options <arg>] <LD_name>...|<pattern>...
```

Description

The `compactld` command consolidates space on the LDs.

Options

-pat

Compacts the LDs that match any of the specified patterns.

-cons

This option consolidates regions into the fewest possible LDs.

When this option is not specified, the regions of each LD will be compacted within the same LD.

-waittask

Waits for any created tasks to complete.

-taskname <taskname>

Specifies a name for the task. When not specified, a default name is chosen.

-dr

Specifies that the operation is a dry run, and the tasks will not actually be performed.

-f

Does not ask for confirmation before compacting the LDs. Unless the `-f` option is specified, the command asks for confirmation before compacting each LD.

-trimonly

Only unused LD space is removed. Regions are not moved.

Specifiers

<LD_name>...

Specifies the name of the LD to be compacted. Multiple LDs can be specified.

<pattern>...

Specifies a glob-style pattern. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.) This specifier can be repeated to compact multiple LDs. If this specifier is not used, the <LD_name> specifier must be used.

Authority

- Super, Edit
- Any role granted the `ld_compact` right

Usage

None.

Examples

In the following example, LD pdsld0.0 is compacted:

```
cli% compactld -f pdsld0.0
Compacting LDs:
     pdsld0.0
```

controlencryption

Syntax

```
controlencryption backup <filename>

controlencryption checkekm [-all]

controlencryption enable [-ekm] <filename>

controlencryption rekey <filename>

controlencryption restore <filename>

controlencryption setekm -setserver <hostname|IP-port> [-port <port>] -
ekmuser <username>

controlencryption setekm -addserver <hostname|IP> <filename>

controlencryption setekm -ekmpass <filename>

controlencryption setekm -ekmuser <username> <filename>

controlencryption setekm -kmipprotocols
<kmip_protocol_version>[,<kmip_protocol_version>] <filename>

controlencryption setekm -port <port> <filename>

controlencryption setekm -removeserver <hostname|IP> <filename>

controlencryption setekm -setserver <hostname|IP> <filename>

controlencryption status [options]
```

Description

The `controlencryption` command provides various commands to control encryption.

Options

`-d`

Provides details on the encryption status.

Specifiers

`<filename>`

Specifies the name of the file in which to place the backup or the source of the file to restore.

`<hostname|IP>`

Specifies a server fully qualified domain name or IP address. Multiple servers can be specified by using a comma separated list.

`<port>`

Specifies port to be used on External Key Management server.

Subcommands

`backup`

Specifies to backup the keystore.

checkekm

Verifies communication with at least one External Key Management server. When `-all` is specified, verifies all External Key Management servers.

enable

Specifies to enable encryption. When `-ekm` is specified, the External Key Management server is used instead of Local Key Management.

rekey

Specifies to rekey the Self Encrypting Drives (SEDs).

restore

Specifies to restore the keystore.

setekm (creation)

For creation, the following options must be specified:

```
-setserver -ekmuser
```

Where `-port` is optional.

setekm (modification)

setekm -addserver

Specifies hostname(s) or IP(s) of External Key Management server(s). These servers are added to the list of already configured servers. Multiple EKM servers can be specified by using a comma separated list. Up to four EKM servers can be defined.

setekm -ekmpass

Prompts for the password for the External Key Management server user.

setekm -ekmuser

Specifies the External Key Management server username.

setekm -port

Specifies port used to communicate with External Key Management server.

setekm -removeserver

Specifies hostname(s) or IP(s) of External Key Management server(s) These servers are removed from the list of already configured servers. Multiple EKM servers can be specified using a comma separated list.

setekm -setserver

Specifies hostname(s) or IP(s) of External Key Management server(s). Multiple EKM servers can be specified by using a comma separated list.

status

Specifies to list the current encryption status.

Authority

- Super
- Any role granted the `control_encryption` right
- Super, Service (for `status` subcommand only)
- Any role granted the `control_encryption_status` right

Usage

- Access to all domains is required to run this command.
- The keystore MUST be backed up to prevent total loss of data. A password must be supplied (twice) on backup. The same password must be supplied on restore.
- Because the backup must be stored off the nodes, all subcommands other than `status` must be executed using a version of the CLI running on a remote system.

  IMPORTANT: Backups when running with an External Key manager are for configuration information ONLY! While this file is still important, as it is necessary to recover from a disaster, the keys are stored ONLY on the EKM, and must be backed up independently. Also, when configuring the EKM and before encryption is enabled, a backup filename is not required.

- The Data Encryption license is required to use the `enable` subcommand. Also all disks must be SED.
- For `enable` and `rekey` subcommands, there must be no failed disks.
- If the state shows recovery_needed, the `enable` or latest `rekey` operation failed. Re-issue the operation to fix this state.
- Required parameters for initial `setekm` operation are shown in the syntax section.

Examples

The following example backs up the keystore to file backup1:

```
cli% controlencryption backup backup1
```

The following example restores the keystore from file backup1:

```
cli% controlencryption restore backup1
```

The following example configures an External Key Management server:

```
cli% controlencryption setekm -setserver ekm.example.com -port 9001
-ekmuser username -kmipprotocols 1.2,1.3
```

Use the following to get a description of each column in the above examples:

```
clihelp -col controlencryption
```

controliscsiport

Syntax

The syntax for the `controliscsiport` command can be one of the following:

```
controliscsiport addr <IP_address> <netmask|prefix_len> [-vlan <vlan_tag>] [-
f] <node:slot:port>

controliscsiport gw <gw_address> [-vlan <vlan_tag>] [-f] <node:slot:port>

controliscsiport mtu <mtusz_bytes> [-vlan <vlan_tag>] [-f] <node:slot:port>

controliscsiport isns <isns_address> [-vlan <vlan_tag>] [-f] <node:slot:port>

controliscsiport isnsport <isns_port> [-vlan <vlan_tag>] [-f]
<node:slot:port>

controliscsiport delete [-vlan <vlan_tag>] [-f] <node:slot:port>
```

```
controliscsiport ping [<count>] <ipaddr> [-vlan <vlan_tag>] <node:slot:port>
controliscsiport stgt <sendtgt_grp> [-vlan <vlan_tag>] [-f] <node:slot:port>
```

Description

The `controliscsiport` command is used to set properties of an iSCSI port.

Options

`-f`

Do not ask for confirmation. The default is to ask for confirmation.

Specifiers

`<IP_address>`

Indicates the IP address of the iSCSI target.

`<netmask>`

Indicates the IP netmask of the IPv4 iSCSI target.

`<prefix_len>`

Indicates the IP prefix length of the IPv6 iSCSI target.

`<gw_address>`

Indicates the IP address of the gateway.

`<mtusz_bytes>`

Indicates the MTU size in bytes.

`<isns_address>`

Indicates the IP address of the iSNS server.

`<isns_port>`

Indicates the TCP port number of the iSNS server. The default port number is 3205.

`<count>`

Indicates the number of ping packets to send. If a value is specified, it should be an integer from 1 through 64. If a value is not specified, the default is one packet.

`<sendtgt_grp>`

Indicates the STGT of the iSCSI target. <sendtgt_grp> is an integer in the range [0-65535].

`<vlan_tag>`

Indicates the VLAN tag number (1 to 4094).

`<node:slot:port>`

The physical location of the iSCSI target port.

Subcommands

`addr`

Sets the IPv4 address and netmask, or IPv6 address and prefix length of the iSCSI port.

`gw`

Sets the gateway address of the iSCSI port.

`mtu`

Sets the maximum transmission unit (MTU) size for the iSCSI port.

isns

Sets the iSNS server IP address.

isnsport

Sets the TCP port number for the iSNS server. By default, the default iSNS port number is used.

delete

Deletes the iSCSI port configuration.

ping

Pings the specified IP address a specified number of times from the iSCSI port.

stgt

Sets the SendTargets Group Tag (STGT) for the iSCSI port.

Authority

- Super, Service
- Any role granted the `port_control` right

Usage

- Access to all domains is required to run this command.
- Issuing `controliscsiport` will return an error on a dual-port QLogic CNA card if either port is configured in FCoE protocol.

Examples

In the following example, iSCSI port configuration is deleted:

```
cli% controliscsiport delete -f 1:1:1
```

controlmag

Syntax

```
controlmag offloop|onloop [options] <cage_name> <magazine>
```

Description

The `controlmag` command takes drive magazines, or disk drives within a magazine, either on-loop or off-loop. Use this command when replacing a drive magazine or disk drive within a drive magazine.

Options

-disk <disk_number>

Specifies that the operation is performed on the disk as determined by its position within the drive magazine. If not specified, the operation is performed on the entire drive magazine.

-port a|b|both

Specifies that the operation is performed on port A, port B, or both A and B. If not specified, the operation is performed on both ports A and B.

-f

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

Specifiers

<cage_name>

Specifies the name of the drive cage. Drive cage information can be viewed by issuing the **showcage** command.

<magazine>

Specifies the drive magazine number within the drive cage. Valid formats are <drive_cage_number>.<drive_magazine> or <drive_magazine> (for example 1.3 or 3, respectively).

Subcommands

offloop|onloop

Specifies that the specified drive magazine or disk drive is either taken off-loop or brought back on-loop.

Authority

- Super, Service
- Any role granted the `mag_control` right

Usage

- Access to all domains is required to run this command.
- Taking a drive magazine off-loop has the following consequences:
  - Relocation of chunklets.
  - Affected logical disks are put into write-through mode.
  - Momentary dip in throughput, but no loss of connectivity.

Examples

The following example requests that drive magazine 1 in drive cage cage0 be put on loop:

```
cli% controlmag onloop cage0 1
```

controlpd

Syntax

```
controlpd spinup|spindown [options] <WWN>...
```

Description

The `controlpd` command spins a PD up or down. This command is used when replacing a PD in a drive magazine.

Options

-ovrd

Specifies that the operation is forced, even if the PD is in use.

Specifiers

<WWN>...

Specifies the World Wide Name of the PD. This specifier can be repeated to identify multiple PDs.

Subcommands

spinup

Specifies that the PD is to spin up. If this subcommand is not used, then the spindown subcommand must be used.

spindown

Specifies that the PD is to spin down. If this subcommand is not used, then the spinup subcommand must be used.

Authority

- Super, Service
- Any role granted the `pd_control` right

Usage

- Access to all domains is required to run this command.
- The spin down operation cannot be performed on a PD that is in use unless the -ovrd option is used.
- Issuing the `controlpd` command puts the specified disk drive in a not ready state. Further, if this command is issued with the spindown subcommand, data on the specified drive becomes inaccessible.

Examples

The following example instigates the spin up of a PD identified by its WWN of 2000000087002078:

```
cli% controlpd spinup 2000000087002078
```

controlport

Syntax

```
controlport cl2 <cl2val> [-f] <node:slot:port>... (deprecated usage)

controlport config <connmode> [-ct <ctval>] [-unique_nwwn enable|disable] [-
virt_ports <num_virt_ports>] [-f] <node:slot:port>...

controlport ct <ctval> [-f] <node:slot:port>... (deprecated usage)

controlport failback <node:slot:port>...

controlport failover <node:slot:port>...

controlport fs add [-f] <node:slot:port>...

controlport fs delete [-f] <node:slot:port>...

controlport fs ping [-c <count>] [-node <sourcenode>] [-source <sourceaddr>]
<target>

controlport fs speed auto <node:slot:port>...

controlport fs speed <value> {half|full} <node:slot:port>...

controlport fs state {up|down} [-f] <node:slot:port>...

controlport fs traceroute [-node <sourcenode>] [-source <sourceaddr>]
<target>

controlport intcoal <intcoalval> [-f] <node:slot:port>...

controlport label <port_label> [-f] <node:slot:port>...

controlport lip [-c <cagename>] [-portwwn <wwn>] [-f] <node:slot:port>...

controlport login <portwwn> [-f] <node:slot:port>...

controlport logout <portwwn> [-f] <node:slot:port>...

controlport nssync [-f] <node:slot:port>...
```

```
controlport offline <node:slot:port>...
controlport rate <rateval> [-f] <node:slot:port>...
controlport rcfc delete [-f] <node:slot:port>
controlport rcfc init [-f] <node:slot:port>
controlport rcip add [-f] <IP_address> <netmask> <node:slot:port>...
controlport rcip addr [-f] <IP_address> <netmask> <node:slot:port>...
controlport rcip delete [-f] <node:slot:port>...
controlport rcip gw [-f] <gateway_address> <node:slot:port>...
controlport rcip mtu <MTU_size> <node:slot:port>...
controlport rcip ping [-c <count>|-w <wait>|-s <size>|-pf] <IP_address>
<node:slot:port>...
controlport rcip speed auto <node:slot:port>...
controlport rcip speed <value> {half|full} <node:slot:port>...
controlport rcip state {up|down} [-f] <node:slot:port>...
controlport rst [-m <mode>|-l|-i] [-f] <node:slot:port>...
controlport tmwo (enable|disable) [-f] <node:slot:port>...
controlport vcn <vcnval> [-f] <node:slot:port>...
```

Description

Controls all aspects of a Fibre Channel, SAS, Remote Copy, or File Persona port, including the port connection type and data rate. To set up Remote Copy interfaces, use the `controlport rcip` and `controlport rcfc` commands.

Options

`-c`

- When used with the LIP subcommand:

  `-c <cage_name>`

  If using a private loop topology, a LIP command is issued from the port. If a cage is specified using the <*cage_name*> argument, the LIP is issued through the cage controller. If using a point-to-point topology, the link is reset. If using a public loop or fabric topology, a RSCN message is issued to the fabric controller.

- When used with the `rcip ping` subcommand:

  `-c <count>`

  Specifies the number of replies accepted by the system before terminating the command. The default is 1; the maximum value is 25.

- When used with the `fs ping` subcommand:

  -c <count>

  Specifies the number of replies accepted by the system before terminating the command. <count> must be in the range 1-25. Default value is 4.

`-ct`

Sets the connection type. The specified port is reset.

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

`-i`

Forces a mode change for a port even if the port's mode change value is prohibited. The partner port's mode is changed with this option as well. Use with the `controlport rst` command. This option cannot be used with the -l option.

`-l`

Forces the port to reload firmware. This option can only be used with the rst subcommand and cannot be used with the -i option.

`-m <mode>`

This option can only be used with the rst subcommand. Resets the mode of the port. The port can be reset to a target or initiator mode. If not specified, the port is reset to its current mode. If the port's mode change value is prohibited, this command fails when attempting to reset to a different mode. Use `showport -c` to see whether mode change is allowed or prohibited for a particular port. See NOTES for additional information regarding port pair protection. This option is deprecated and will be removed in a later release. Use `controlport config` to set target (host) or initiator (disk) modes.

`-node <sourcenode>`

FS node from which ping/traceroute packets are to be sent. `<sourcenode>` must be in the range 0-7. If not provided, the active node is used.

`-pf`

Prevents packet fragmentation. This option can only be used with the rcip ping subcommand.

`-portwwn <wwn>`

Specifies the Fibre Channel worldwide port name of an attached device for which selective LIP reset will be performed.

`-s <size>`

Specifies the packet size. If no size is specified, the option defaults to 64. This option can only be used with the rcip ping subcommand.

`-source <sourceaddr>`

Sources IP address for identifying the interface on the source node from which ping/traceroute packets are to be sent. If not provided, the default outgoing interface is used.

`-unique_nwwn enable|disable`

Enables or disable the use of a unique node wwn on the specified port.

`-virt_ports <num_virt_ports>`

Specifies the creation/removal of virtual peer ports while configuring a port using the `controlport config peer` command.

If the number of virtual peer ports specified is greater than the number of existing virtual peer ports, then additional virtual peer ports will be created to match the total num_virt_ports. The maximum number of virtual peer ports that can be created is 8. The `showport -peer` command displays the virtual peer ports.

If the number of virtual peer ports specified is zero, then the virtual peer ports are removed.

`-w <wait>`

Specifies the maximum amount of time to wait for replies. The default is the number of requested replies plus 5. The maximum value is 30. This option can only be used with the rcip ping subcommand.

Specifiers

`<cl2val>`

Specifies the Fibre Channel Class-2 parameter of the port. Parameters can be one of ack0, ack1, or disable. This specifier must be provided when issuing the cl2 subcommand.

Note that this specifier is deprecated and will be removed in a subsequent release.

`<connmode>`

Specifies whether the port is used for a disk, host, peer, Remote Copy over Fibre Channel (rcfc), or iSCSI (iscsi) connection. This specifier must be used when issuing the `config` command.

`<ctval>`

Specifies the connection parameter of the port. Parameters can be loop, or point. The loop parameter sets the port to arbitrated loop mode, the point parameter sets the port to point-to-point mode. This specifier must be provided when issuing the ct subcommand.

`<gateway_address>`

Specifies the gateway address for a Remote Copy interface.

`<impval>`

Specifies the IMP port attribute. The IMP value can be set as enable or disable. This specifier must be used when issuing the imp subcommand.

`<intcoalval>`

Specifies if interrupt coalescing is enabled or disabled. The value can be set as enable or disable. This specifier must be used with the intcoal subcommand.

`<IP_address>`

Specifies the IP address for a Remote Copy interface.

`<MTU_size>`

Specifies the MTU size for a Remote Copy interface using an integer from 1000 through 9000.

`<netmask>`

Specifies the netmask for a Remote Copy interface.

`<node:slot:port>`

Specifies the port to be controlled.

`node`

Specifies the node using a number between 0 and 7.

`slot`

Specifies the PCI slot in the specified node. Valid range is 0 - 9.

`port`

Specifies the port using a number from 1 through 4.

`<port_label>`

Specifies the value of the port label. The port label value is a human-readable string used for easy identification of the port.

`<portwwn>`

Specifies the Fibre Channel worldwide port name of an attached port. This specifier is used with the login and logout subcommands.

`<rateval>`

Specifies the data rate of the Fibre Channel port. Rates can be 4, 8, 16, 32, or auto. 4 sets the data rate to 4 Gbps. And similarly for the other rates. The auto parameter sets the port to auto detect the data rate. This specifier must be used when issuing the rate subcommand.

<target>

IP address or name of the target host.

<value> {half|full}

Uses only with the rcip speed subcommand. Specifies the speed setting (1000) and duplex setting (half or full) for a Remote Copy interface. In addition to this specifier, you must also specify an interface using <node:slot:port>. If no speed or duplex settings are specified, or if you specify auto with the rcip speed subcommand, the system auto-negotiates the speed and duplex.

<vcnval>

Specifies the value of the VCN. The VCN value can be set to enable or disable. This specifier must be used when issuing the vcn subcommand.

Subcommands

cl2

Specifies the Fibre Channel Class-2 parameter of the port. The specified port will be reset. Note that this subcommand is deprecated and will be removed in a subsequent release.

config

Sets the specified connection mode and type on FC ports. Also allows setting the unique node WWN option for the port. When unique_nwwn is enabled, the port presents a unique node name on the connection, which is needed by certain initiators such as ONTAP. This command is also used to configure the mode of a CNA port. When configuring a peer port, the command can be used to create or remove virtual peer ports.

ct

Sets the connection type. The specified port will be reset.

Note that this subcommand is deprecated and will be removed in a subsequent release.

failback

Reverts a port failover using the persistent ports mechanism, if port failover is currently active. The failed-over port should be specified with this command.

failover

Fails the specified port over to a failover partner port using the virtual ports mechanism, if a valid partner port is present.

fs add

Adds one or more specified ports as File Persona interfaces.

fs delete

Deletes one or more specified File Persona interfaces.

fs ping

Runs the Linux ping command for the target host specified in the command.

fs speed

Instructs the specified File Persona interface(s) to use the specified speed and duplex, or to auto negotiate speed and duplex. The default is auto. Only for File Persona interfaces.

fs state {up|down}

Sets the specified File Persona interface(s) as either up or down.

fs traceroute

Runs the Linux traceroute command for the target host specified in the command.

intcoal

Enables or disables interrupt coalescing. The specified port will be reset.

label

Associates a human-readable label with the port. This subcommand is not supported with RCIP ports.

lip

Specifies that a Loop Initialization Primitive (LIP) command is issued from the port if there is a private loop topology. If the -c option is specified, then the LIP command is issued through the specified drive cage. If there is a point-to-point topology, then the link is reset. If there is a public loop or fabric topology, then a Registered State Change Notification (RSCN) is issued to the fabric controller. If the -portwwn option is specified, then a selective LIP or RSCN is issued to the specified attached port.

login

Performs Fibre Channel login to an attached port, identified by port WWN.

logout

Performs Fibre Channel logout from an attached port, identified by port WWN.

nssync

Verifies current port database against the Name Server when a fabric is attached. Entries present in the database but missing from the Name Server are removed. Using this command is not required under normal circumstances.

offline

Holds the specified port offline indefinitely. Issue `controlport rst` to bring the port back online.

rate

Specifies the data rate of the Fibre Channel or iSCSI port. The specified port will be reset.

rcfc delete

Deletes the configuration for the Remote Copy interface on the local port. This will bring the interfaces down.

rcfc init

Sets the specified Remote Copy interface on the local port. Only for RCFC ports.

rcip add

Configures the given network port for Remote Copy with the specified IP address and netmask.

rcip addr

Sets the given Remote Copy interface to use the specified IP address and netmask.

rcip delete

Deletes the configuration for one or more specified Remote Copy interfaces.

rcip gw

Sets the gateway for one or more Remote Copy interfaces. Only for RCIP ports.

rcip mtu

Sets the Maximum Transfer Unit (MTU) size for the specified Remote Copy interface(s), overriding the default of 1500. The largest supported value is 9000 and the smallest is 1000. Only for RCIP ports.

rcip ping

Performs a ping from the specified interface to the specified address.

Only for RCIP ports. Permitted for Super, Service, Edit, and Browse users.

rcip speed

Instructs the specified Remote Copy interface(s) to use the specified speed and duplex, or to auto negotiate speed and duplex. The default is auto. Only for RCIP ports.

rcip state {up|down}

Sets the specified Remote Copy interface(s) as either up or down. Only for RCIP ports.

rst

Resets a port. This subcommand is not applicable to RCIP ports.

tmwo

Enables or disables target mode write optimization. The specified port will be reset.

vcn

Sets the VLUN Change Notification (VCN) generation support (enable or disable). When VCN generation support is enabled with a public loop or fabric topology, a RSCN message is issued to the fabric controller whenever a VLUN is created or removed. In addition, if enabled with a public loop topology, a LIP is issued from the port whenever a VLUN is created or removed.

Authority

- Super, Service
- Any role granted the `port_control` right

Usage

- Access to all domains is required to run this command.
- The `-m <mode>` option for the `rst` subcommand cannot be specified if there are active connections already using the port (port online), except as noted in the following section under port pair protection.
- Port pair protection:
  - For dual-port LSI Fibre Channel adapters, both ports in the pair must use the same mode (initiator or target).
  - For quad-port LSI Fibre Channel adapters, each port pair (ports 1 and 2, ports 3 and 4) must use the same mode.
  - Changing the mode of one port in a pair (for example, from initiator to target) causes the other port in the pair to undergo a mode as well.
- If there are active hosts or physical disks when issuing the `controlport rst` or `offline` commands, a warning is returned and you are prompted for confirmation to complete the execution of the commands.
- When issuing the `controlport ct`, `cl2`, `rate`, or `vcn` commands, if there are active disks on the port, an error is returned. If there are active hosts on the port, the -f option is overridden (if specified), a warning is returned, and you are prompted for confirmation to complete the execution of the commands.
- When issuing the `controlport cl2` or `vcn` commands, if the port is configured as peer connection mode, an error is returned.
- The `controlport rcip add` command is only allowed for node/slot/port combinations where there is an interface installed.
- The `controlport rcip addr` command is only allowed for node/slot/port combinations where there is an interface installed.

- Controlport config mode changes are not allowed when ports are online.
- Removing virtual peer ports through `controlport config peer -virt_ports 0` is not allowed when the port is online.
- Controlport config mode change to FCoE on dual-port QLogic CNA card will only successfully complete if both iSCSI ports are unconfigured using `controliscsiport addr` (setting IP address and netmask to zero values) or delete subcommands.
- The ct, cl2, or rate subcommands automatically resets the port for the changes to take effect.
- Issue the **`showport`** command with either the -i or -par options to verify the success of the `controlport` command.
- Resetting a port causes a momentary dip in throughput, but no loss of connectivity.
- The -f flag forces the operation. Without the flag, the command prompts for confirmation. In some cases, unless the TPDFORCE_OVERRIDE environment variable is set, the command will ask for confirmation even if the -f flag is specified because the operation may disrupt the system operation.
- Use caution when changing modes for ports in LSI Fibre Channel adapters.
- Changing the mode of one port in a pair (for example, from initiator to target) causes the other port in the pair to also undergo a mode change. In the case where one port in the pair is offline (and therefore a mode change can be allowed), but the partner port is online, changing the mode of the offline port causes the online partner port to undergo mode change as well.
- This results in loss of use of the partner port because it is no longer online.
- Issuing the `controlport rst -l` command affects both ports of a port pair.
- Only use this command when irreversible damage has been done to a port or port pair.
- If specified, <ctval> will depend on <connmode> value:
  - If <connmode> is "disk" then <ctval> can only be "loop".
  - If <connmode> is "rcfc", then <ctval> can only be "point".
  - If <connmode> is "host", then <ctval> can be either "loop" or "point".
  - If <connmode> is "peer", then <ctval> can only be "point".
  - If <connmode> is "iscsi", then <ctval> can only be "point".

Examples

The following example shows how to increase MTU to 9000 on Ethernet port in node 0, slot 1:

```
cli% controlport rcip mtu 9000 0:1:1
Remote Copy change successful.
```

The following example shows how to set Remote Copy interface on Ethernet port in node 0, slot 1, configure IP address 172.16.1.11 using a netmask of 255.255.255.0:

```
cli% controlport rcip addr 172.16.1.11 255.255.255.0 0:1:1
Are you sure you want to change the address for 0:1:1?
select q=quit y=yes n=no: y
Remote Copy interface change successful.
```

The following example shows how to set the gateway for Ethernet port in node 0, slot 1 using a gateway address of 172.16.1.1:

```
cli% controlport rcip gw 172.16.1.1 0:1:1
Are you sure you want to change the gateway for 0:1:1?
select q=quit y=yes n=no: y
Remote Copy interface change successful.
```

The following examples show usage of `controlport config`:

```
controlport config host -ct point 1:3:1
controlport config host -unique_nwwn enable 1:2:3
controlport config rcfc 0:0:1
controlport config disk 2:2:2
controlport config iscsi 1:8:1
controlport config peer -ct point 1:5:1
controlport config peer -virt_ports 8 3:2:1
```

The following example shows the usage of `controlport fs ping`:

```
cli% controlport fs ping -source 10.35.4.120 -node 0 -c 2 10.135.4.227
PING 10.135.4.227 (10.135.4.227) from 10.35.4.120 : 56(84) bytes of data.
64 bytes from 10.135.4.227: icmp_seq=1 ttl=62 time=0.286 ms
64 bytes from 10.135.4.227: icmp_seq=2 ttl=62 time=0.275 ms

--- 10.135.4.227 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 999ms
rtt min/avg/max/mdev = 0.275/0.280/0.286/0.017 ms
```

The following example shows the usage of `controlport fs traceroute`:

```
cli% controlport fs traceroute -source 10.35.4.121 -node 1 10.35.16.222
traceroute to 10.35.16.222 (10.35.16.222), 30 hops max, 60 byte packets
1  10.35.4.254 (10.35.4.254)  18.264 ms  18.239 ms  18.229 ms
2  10.35.4.254 (10.35.4.254)  2992.447 ms !H  3033.789 ms !H *
```

controlrecoveryauth

Syntax

```
controlrecoveryauth ciphertext <user>[,<user>]

controlrecoveryauth rollcred <user>[,<user>]

controlrecoveryauth setmethod ciphertext|totp

controlrecoveryauth status
```

Description

The `controlrecoveryauth` command controls the method used to authenticate recovery accounts.

Options

None.

Specifiers

None.

Subcommands

`ciphertext <user>[,<user>]`

Exports the user's Ciphertext for decryption by HPE Support.

`rollcred <user>[,<user>]`

Generates a new random password for the user if the 3PAR is in ciphertext mode.

setmethod ciphertext|totp

Changes the method of authentication. The two methods are described in the notes.

status

Returns the current authentication method.

Authority

- Super
- Any role granted the `recoveryauth_control` right
- Super, Service (for status subcommand only)
- Any role granted the `recoveryauth_status` right

Usage

- Time-based One-Time Passwords (TOTP)—Random passwords are generated by HPE Support, they are only valid for a single HPE 3PAR storage system and for a short period of time. TOTP is the default method.
- Ciphertext—Random passwords are generated and encrypted by the 3PAR storage system. The Ciphertext must be decrypted by HPE Support.
- The `controlrecoveryauth` command only manages accounts used by HPE support for console port access.

Examples

None.

controlsecurity

Syntax

```
controlsecurity fips enable [-f]
controlsecurity fips disable [-f]
controlsecurity fips restart [-f]
controlsecurity fips status
controlsecurity ssh-keys generate [-f]
controlsecurity ssh-keys sync [-f]
```

Description

Controls security parameters of an HPE 3PAR storage system. **Options**

The following options are valid for `fips` and `ssh-keys` subcommands:

-f

Suspends user interaction and forces the operation to proceed. If not used, the command requires confirmation before proceeding with the operation.

Subcommands

fips enable

Enables the use of FIPS 140-2 validated cryptographic modules on system management interfaces.

fips disable

Disables the use of FIPS 140-2 validated cryptographic modules on system management interfaces.

fips restart

Restarts all services that are in `Enable failed` status.

fips status

Shows the status of security parameters of system management interfaces.

ssh-keys generate

Regenerates the SSH host keys and distributes them to all nodes.

ssh-keys sync

Copies the SSH host keys from the current node to all other nodes.

Authority

- Super
- Any role granted the `security_control` right.
- Super, Service (for status option only)
- Any role granted the `security_status_control` right.

Usage

The Management Interfaces are CIM, CLI, EKM used for Data at Rest Encryption, LDAP Authentication, QW, RDA, SNMP, Syslog, SSH, WSAPI, and VASA.

EKM and Syslog interfaces always have FIPS mode enabled.

WARNING: Enabling FIPS mode will terminate ALL existing management interfaces/connections/services.

WARNING: Regenerating or syncing the SSH host keys will terminate ALL existing SSH connections.

Examples

Showing the current mode of FIPS and status of services:

```
cli% controlsecurity fips status
FIPS mode: Enabled

Service Status
CIM     Disabled
CLI     Enabled
EKM     Enabled
LDAP    Enabled
QW        Enabled
SSH     Enabled
SYSLOG  Enabled
VASA    Disabled
WSAPI   Disabled
----------------
9       6 Enabled
```

Enabling FIPS mode:

```
cli% controlsecurity fips enable
    Warning: Enabling FIPS mode requires restarting all system management interfaces,
```

```
    which will terminate ALL existing connections including this one.
    When that happens, you must reconnect to continue.
    Continue enabling FIPS mode (yes/no)?
```

Disabling FIPS mode:

```
cli% controlsecurity fips disable
    Warning: Disabling FIPS mode requires restarting all system management interfaces,
    which will terminate ALL existing connections including this one.
    When that happens, you must reconnect to continue.
    Continue disabling FIPS mode (yes/no)?
```

Restarting services that are not enabled:

```
cli% controlsecurity fips restart
    Warning: Will restart all services that are not enabled,
    which may terminate ALL existing connections including this one.
    When that happens, you must reconnect to continue.
    Continue restarting (yes/no)?
```

Regenerating the SSH host keys and distributing them to the other nodes:

```
cli% controlsecurity ssh-keys generate
    Warning: This action will restart the ssh service,
    which may terminate ALL existing connections including this one.
    When that happens, you must reconnect to continue.
    Continue restarting (yes/no)?
```

Synchronizing the SSH host keys from the current node to all other nodes:

```
cli% controlsecurity ssh-keys sync
    Warning: This action will restart the ssh service,
    which may terminate ALL existing connections including this one.
    When that happens, you must reconnect to continue.
    Continue restarting (yes/no)?
```

controlsr

Syntax

```
controlsr export [options]

controlsr grow <size>

controlsr grow -pct <percentage>

controlsr restore

controlsr setperiod [options] {hires|ldrg} <target_period>

controlsr setretention [options] <category> <target>[-<max>]

controlsr enable <data_type>[,<data_type>...]

controlsr disable <data_type>[,<data_type>...]
```

Description

The `controlsr` command manipulates the System Reporter utility.

Options

The `controlsr` export subcommand requires one of the following options:

-ldrg

Exports the LD region access data which Adaptive Optimization uses. This output can be required for HPE 3PAR to analyze unexpected AO behavior. The output file can be large, so it is important to carefully select the required time range.

-hires

Exports all the System Reporter space and performance data at the highest available resolution, 5 minute samples.

-hourly

Exports all of the System Reporter space and performance data at hourly resolution.

-daily

Exports all of the System Reporter space and performance data at daily resolution.

-all

Exports all of the System Reporter space and performance data for the hires, hourly, and daily categories together.

-cleanup

Uses this option to remove an export file from the node if there was a problem with the automatic deletion. The -cleanup option will only remove a previous export archive file if it is not being accessed. If the file is being accessed, it may still be in the process of copying off of the node. Use the -f option to force the file deletion if it is not being copied, but-cleanup will not proceed.

These options are also available to the `controlsr` export subcommand:

-btsecs <secs>

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

If it is not specified then the time at which the report begins depends on the sample category (-hires, -hourly, -daily):

- For hires, the default begin time is 12 hours ago (-btsecs -12h).
- For hourly, the default begin time is 7 days ago (-btsecs -7d).
- For daily, the default begin time is 90 days ago (-btsecs -90d).

If begin time and sample category are not specified then the time the report begins is 12 hours ago and the default sample category is hires. If -btsecs 0 is specified then the report begins at the earliest sample.

-etsecs <secs>

Selects the end time in seconds for the report. If -attime is specified, select the time for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

If it is not specified then the report ends with the most recent sample.

-estimate

Performs calculations to estimate the size of the output file created by the export operation. Use in conjunction with one of: -ldrg, -ldrgsum, -hires, -hourly, -daily, or -all. To create an accurate estimate, also include -btsecs and -etsecs.

-save

Rather than sending the export file to HPE 3PAR Central remote infrastructure, the -save option exports the file to a location local to the machine that has initiated the remote CLI connection. This option must also be used with -file, which specifies the export file location and name. The -save option is only supported when the session is initiated via the Remote CLI Client.

-file <file>

Specifies the path and name to be used for the export file copied to the client system. If the value is a directory, the default file name of exportsr_<category>_<serialnum>_<date>_<time>.tbz will be used. This option must also be used with the -save option.

-f

When used in conjunction with the -cleanup option, the `-f` option forces cleanup of previous export files that may still exist. The -f option will have no affect on any other operation of `controlsr` export.

-waittask

Specifies that the command will wait for any created tasks to complete.

-v

Displays the detailed status of the task as it executes. This option is only valid when combined with `-waittask`.

This option is available to the `controlsr` grow subcommand:

-pct <percentage>

Allows the SR volume to be grown by a percentage of its current size.

<percentage> must be between 5 and 100.

The following options are available to the `setperiod` and `setretention` subcommands:

-estimate

Estimates the amount of history retention which will be possible if the given settings are changed. Changing the settings of one category will affect the retention of all categories. An estimate will be displayed for each category of data.

-f

Forces the setting change without requiring confirmation when the new period or retention setting is not estimated to be achievable.

Specifiers

<size>[g|G]

Only applicable to the grow subcommand.

Specifies the size in MB to be added to the SR volume. The size must be an integer in the range from 8g to 1T. The volume size is rounded up to the next multiple of 256M. If the value is followed by a g or G (no whitespace before the suffix), the size is in GB.

<period>

Only applicable to the `setperiod` subcommand.

The `hires` data can have a period between 1 minute and 5 minutes in full minute increments, specified with an m suffix: `1m`, `2m`, `3m`, `4m`, `5m`. Depending on the size, configuration, and load on the system, short `hires` periods might not be possible, and a warning is presented.

<category>

In the `setretention` subcommand, the available categories are `hires`, `hourly`, and `daily`.

In the `setperiod` subcommand, the available category is `hires`.

<target>[-<max>]

Only applicable to the setretention subcommand.

The algorithm will attempt to meet all category target retentions. The retention max specifies a maximum retention duration for that category's data, beyond which historical data will be removed.

The target and max retention durations are specified in either days (suffix d) or years (suffix y), e.g. "10d", "90d", or "2.5y" To remove the max setting from a category use a value of 0.

Subcommands

export

Exports a portion of the System Reporter data for external analysis.

grow

Increases the SR data storage space, allowing for longer data retention.

restore

Promotes the .srdata snapshot to the .srdata volume.

setperiod [-estimate] hires <target_period>

Changes the target period between samples of the hires data. The hires data may have a period between 1 minute and 5 minutes in full minute increments, specified as "1m" or "2m", etc.

setretention [-estimate] <category> <target>[-<max>]

Changes the historical retention characteristics of different categories of System Reporter data. The available categories are `hires`, `hourly`, and `daily`. The target retention specifies a duration of history for the data age-out algorithm to balance storage space between categories. The retention max specifies a maximum retention duration for that category's data, beyond which historical data will be removed. The data age-out algorithm balances the requests for multiple category retention targets, and may not be able to achieve one or more of the retention targets. Retention target and max durations are specified in either days (suffix d) or years (suffix y), e.g. "10d", "90d", or "2.5y" By default there is not a max limit placed on any category. To remove the max setting from a category use a duration of 0, e.g. "setretention hourly 90d-0"

enable <data_type>[,<data_type>...]

disable <data_type>[,<data_type>...]

To save storage space and processing time, you can disable collection and reporting by System Reporter for certain types of data. The data types that you can disable and enable are `fs` and `bulkobjs`, and `ldrg`.

Authority

Super, Service.

Usage

- There is a size limit of 10GB of SR data to be exported per request. If the command is estimated to output more than 10GB, the command will be rejected. Use the -btsecs and -etsecs options to define a small enough request window such that less than 10GB of data is exported. The -estimate option can be used to see the estimated size of output for a combination of options.
- By default, the export file is copied to HPE 3PAR Central by a linked Service Processor. Use the -save option to instead copy the file to the CLI client machine.

Examples

The following example exports the hourly System Reporter data for the past 24 hours:

```
cli% controlsr export -hourly -btsecs -24h
```

The following example grows the SR volume with an increase of 40GB:

```
cli% controlsr grow 40G
```

The following example restores the latest .srdata snapshot to the .srdata volume:

```
cli% controlsr restore
```

To set the high resolution data retention target to 14 days:

```
cli% controlsr setretention hires 14d
```

To set a target and max retention of ldrg data:

```
cli% controlsr setretention ldrg 14d-31d
```

To remove a retention max use 0. The target must also be specified:

```
cli% controlsr setretention hourly 90d-0
```

To change the high resolution data sampling from the default of 5 minutes to 2 minutes:

```
cli% controlsr setperiod hires 2m
```

To change the period between ldrg samples at the target retention time from the default 30 minutes to 3 hours (leaving space for longer overall ldrg data retention):

```
cli% controlsr setperiod ldrg 3h
```

To disable the collection of any File Persona data in the System Reporter history:

```
cli% controlsr disable fs
```

createaocfg

Syntax

```
createaocfg [options <arg>] <aocfg_name>
```

Description

Creates an Adaptive Optimization configuration.

Options

-t0cpg <cpgname>

Specifies the Tier 0 CPG for this AO config.

-t1cpg <cpgname>

Specifies the Tier 1 CPG for this AO config.

-t2cpg <cpgname>

Specifies the Tier 2 CPG for this AO config.

-mode <mode>

Specifies the optimization bias for the AO config and can be one of the following:

- Performance: Move more regions towards higher performance tier.
- Balanced: Balanced between higher performance and lower cost.
- Cost: Move more regions towards lower cost tier.

The default is Balanced.

-t0min <size>[g|G|t|T]

-t1min <size>[g|G|t|T]

-t2min <size>[g|G|t|T]

Specifies the minimum space utilization of the tier CPG for AO to maintain when optimizing regions between tiers. The size can be specified in MB (default) or GB (using g or G) or TB (using t or T).

Setting a minimum to 0 (default) indicates that no minimum space utilization will be enforced.

-t0max <size>[g|G|t|T]

-t1max <size>[g|G|t|T]

-t2max <size>[g|G|t|T]

Specifies the maximum space utilization of the tier CPG. AO will move regions into and out of the CPG based on their relative access rate history, but will not exceed this maximum size in the CPG.

The size can be specified in MB (default) or GB (using g or G) or TB (using t or T). Setting a max to 0 (default) indicates that AO will use other indicators to decide the maximum CPG space utilization:

Either the CPG sdgl, sdgw, or maximum possible growth size.

Specifiers

<aocfg_name>

Specifies an AO configuration name up to 31 characters in length.

Authority

Super, Edit

Usage

- Two or more Tier CPGs must be defined.
- If domains are used, all CPGs must be in the same domain or not in any.
- A CPG can only belong to one AO configuration.
- Active use of Adaptive Optimization requires an Adaptive Optimization license. Contact your local 3PAR representative for information.

Examples

The following example creates a new AO configuration called aocfg1.

The optimization bias (mode) is set to Cost. Tier 0 will use CPG R1SSDcpg, Tier 1 will use CPG R5FCcpg, and Tier 2 will use CPG R6NLcpg.

```
cli% createaocfg -t0cpg R1SSDcpg -t1cpg R5FCcpg -t2cpg R6NLcpg -mode Cost aocfg1
```

The following example creates a new AO configuration called aocfg2.

The optimization bias is defaulted to Balanced, and no Tier 1 CPG is specified.

```
cli% createaocfg -t0cpg R1SSDcpg2 -t2cpg R6NLcpg2 aocfg2
```

createcert

Syntax

```
createcert <SSL_service> {-csr|-selfsigned} [-f] [-keysize <keysize>] [-days
<days>] [<subject_attribute_options>] [<CSR_filename>]
```

Description

Creates a self-signed certificate or a certificate signing request for a specified service.

Options

-f

Used with `-selfsigned` to create a self-signed certificate, and then restart the services that use the new certificate without prompting the user. If not specified, user interaction is required to restart the corresponding services.

-csr

Creates a certificate signing request for the service. Does not modify certificates or restart services.

-selfsigned

Creates a self-signed certificate for the service. Removes the previous certificate and restarts the service. Does not remove the intermediate and/or root certificate authorities for a service.

-keysize <keysize>

Specifies the encryption key size (in bits) of the self-signed certificate. Valid values are 1024 and 2048. Defaults to 2048.

-days <days>

Specifies the valid days of the self-signed certificate. Valid values are between 1 and 3650 days (10 years). Defaults to 1095 days (3 years).

-C <country>

Specifies the value of country/region (C) attribute of the subject of the certificate.

-ST <state>

Specifies the value of state (ST) attribute of the subject of the certificate.

-L <locality>

Specifies the value of locality (L) attribute of the subject of the certificate.

-O <organization>

Specifies the value of organization (O) attribute of the subject of the certificate.

-OU <organizational_unit>

Specifies the value of organizational unit (OU) attribute of the subject of the certificate.

-CN <common_name>

Specifies the value of common name (CN) attribute of the subject of the certificate. Over ssh, `-CN` must be specified.

-SAN <subject_alt_name[,subject_alt_name]...>

Specifies the subject alternative name (SAN). The SAN is an X509 extension that allows other pieces of information to be associated with the certificate. Specify multiple SANs by separating with a comma.

Specifiers

<SSL_service>

Specifies the SSL service name. Valid service names are `cim`, `cli`, `ekm-client`, `ekm-server`, `ldap`, `qw-client`, `qw-server`, `syslog-gen-client`, `syslog-gen-server`, `syslog-sec-client`, `syslog-sec-server`, `wsapi`, `vasa`, and `unified-server`.

Authority

Super

Usage

- The unified-server establishes a common certificate among `cim`, `cli`, and `wsapi`. Restarts the CIM and wsapi services when generating a self-signed certificate. Also, the CIM and wsapi services are restarted when those certificates are changed.
- Data at rest encryption using an external key manager employs the `ekm-client` and `ekm-server` services. The external key manager services are not covered by the `unified-server` service.
- Specify `-CN` when using ssh.
- Use the **importcert** command to import a signed CSR. Use the **showcert** command to display the certificates.

Examples

Creating a self-signed certificate for encryption key size of 2,048 bits and 365-day validity:

```
cli% createcert unified-server -selfsigned -keysize 2048 -days 365
```

Creating a self-signed certificate with multiple subject alternative names (SANs):

```
cli% createcert unified-server -selfsigned \
    -SAN DNS:theresnoplacelikehome.com,IP:127.0.0.1
```

createcpg

Syntax

```
createcpg [options <arg>] CPG_name
```

Description

Creates a Common Provisioning Group (CPG).

Options

-templ <template_name>

Uses the options defined in template <template_name>. Create the template using the createtemplate command. Options specified in the template are read-only or read-write. You can overwrite read-write options with new options at the time of creation. You cannot overwrite read-only options.

Options not explicitly specified in the template use default values. All of these options are either read-only or read-write (using the -nro or -nrw options of the createtemplate command).

-f

Forces the command. The command completes the process without prompting for confirmation.

-aw <percent>

Specifies the percentage of used snapshot administration or snapshot data space that results in a warning alert. A percent value of 0 disables the warning alert generation. Defaults to 0.

This option is deprecated and will be removed in a subsequent release.

-sdgs <size>[g|G|t|T]

Specifies the growth increment, the amount of logical disk storage created on each auto-grow operation. The default growth increment can vary according to the number of controller nodes in the system. If <size> is non-zero it must be 8G or bigger. You can specify the size in MB (default) or GB (using g or G) or TB (using t or T). A size of 0 disables the auto-grow feature. The following table displays the default and minimum growth increments per number of nodes:

Number of Nodes	Default	Minimum
1-2	32 GB	8 GB
3-4	64 GB	16 GB
5-6	96 GB	24 GB
7-8	128 GB	32 GB

-sdgl <size>[g|G|t|T]

Specifies that the auto-grow operation is limited to the specified storage amount. Specify the storage amount in MB (default) or GB (using g or G) or TB (using t or T). A size of 0 (default) means no limit is enforced. To disable auto-grow, set the limit to 1.

-sdgw <size>[g|G|t|T]

Specifies that the threshold of used logical disk space, when exceeded, results in a warning alert. Specify the size in MB (default) or GB (using g or G) or TB (using t or T). A size of 0 (default) means no warning limit is enforced. To set the warning for any used space, set the limit to 1.

-sa <LD_name>...

Specifies that existing logical disks are added to the CPG and used for snapshot admin (SA) space allocation. Repeat the <LD_name> argument to specify multiple logical disks.

This option is deprecated and will be removed in a subsequent release.

-sd <LD_name>..

Specifies that existing logical disks are added to the CPG and are used for snapshot data (SD) space allocation. The <LD_name> argument can be repeated to specify multiple logical disks.

This option is deprecated and will be removed in a subsequent release.

-domain <domain>

Specifies the name of the domain in which the object will reside. A domain member with Edit or Super role must create the object. Defaults to creation in the current domain, or no domain if the current domain is not set.

-t <RAID_type>

Specifies the RAID type of the logical disk: r0 for RAID-0, r1 for RAID-1, r5 for RAID-5, or r6 for RAID-6. If no RAID type is specified, then the default is r6.

-ssz <size_number_chunklet>

Specifies the set size in terms of chunklets. The default depends on the RAID type specified: 2 for RAID-1, 4 for RAID-5, and 8 for RAID-6.

-rs <size>

Specifies the number of sets in a row. The <size> is a positive integer.

If not specified, the option imposes no row limit.

-ss <size_KB>

Specifies the step size from 32 KB to 512 KB. The step size should be a power of 2 and a multiple of 32. The default value depends on raid type and device type used. If no value is entered and FC or NL drives are used, the step size defaults to 256 KB for RAID-0 and RAID-1, and 128 KB for RAID-5. If SSD drives are used, the step size defaults to 32 KB for RAID-0 and RAID-1, and 64 KB for RAID-5. For RAID-6, the default is a function of the set size.

-ha port|cage|mag

Specifies that the layout must support the failure of one port pair, one cage, or one drive magazine (mag). The default is based on the use of the set size parameter. If the -ssz parameter is not specified, then the availability is automatically calculated based on the current system configuration. If the -ssz parameter is specified, then the default value is cage availability.

-ch first|last

Specifies the chunklet location characteristics: either first (attempt to use the lowest numbered available chunklets) or last (attempt to use the highest numbered available chunklets). With no argument specified, defaults to first.

-p <pattern>

Specifies a pattern for candidate disks. Uses patterns to select disks for creating logical disks. With no pattern specified, defaults to Fast Class (FC) disks. If specified multiple times, each instance of the specified pattern adds additional candidate disks that match the pattern.

Cannot use the -devtype pattern to mix Nearline (NL), FC, and Solid State Drive (SSD). Specify an item specified as an integer, a comma-separated list of integers, or a range of integers specified from low to high.

You can use the following arguments as patterns for this option. Specify an item as an integer, a comma-separated list of integers, or a range of integers from low to high.

-nd <item>

Specifies one or more nodes. Nodes are identified by one or more integers (item). Multiple nodes are separated with a single comma (e.g. 1,2,3). A range of nodes is separated with a hyphen (e.g. 0- 7). The primary path of the disks must be on the specified node(s).

-st <item>

Specifies one or more PCI slots. Slots are identified by one or more integers (item). Multiple slots are separated with a single comma (e.g. 1,2,3). A range of slots is separated with a hyphen (e.g. 0- 7). The primary path of the disks must be on the specified PCI slot(s).

-pt <item>

Specifies one or more ports. Ports are identified by one or more integers (item). Multiple ports are separated with a single comma (e.g. 1,2,3). A range of ports is separated with a hyphen (e.g. 0- 4). The primary path of the disks must be on the specified port(s).

-cg <item>

Specifies one or more drive cages. Drive cages are identified by one or more integers (item). Multiple drive cages are separated with a single comma (e.g. 1,2,3). A range of drive cages is separated with a hyphen (e.g. 0-3). The specified drive cage(s) must contain disks.

-mg <item>

Specifies one or more drive magazines. The "1." or "0." displayed in the CagePos column of `showpd` output indicating the side of the cage is omitted when using the `-mg` option. Drive magazines are identified by one or more integers (item). Multiple drive magazines are separated with a single comma (e.g. 1,2,3). A range of drive magazines is separated with a hyphen (e.g. 0-7). The specified drive magazine(s) must contain disks.

-pn <item>

Specifies one or more disk positions within a drive magazine. Disk positions are identified by one or more integers (item). Multiple disk positions are separated with a single comma (e.g. 1,2,3). A range of disk positions is separated with a hyphen(e.g. 0-3). The specified position(s) must contain disks.

-dk <item>

Specifies one or more physical disks. Disks are identified by one or more integers(item). Multiple disks are separated with a single comma (e.g. 1,2,3). A range of disks is separated with a hyphen (e.g. 0-3). Disks must match the specified ID(s).

-tc_gt <number>

Specifies that physical disks with total chunklets greater than the number specified be selected.

-tc_lt <number>

Specifies that physical disks with total chunklets less than the number specified be selected.

-fc_gt <number>

Specifies that physical disks with free chunklets greater than the number specified be selected.

-fc_lt <number>

Specifies that physical disks with free chunklets less than the number specified be selected.

-devid <model>

Specifies that physical disks identified by their models be selected. Models can be specified in a comma-separated list. Models can be displayed by issuing the "`showpd -i`" command.

-devtype <type>

Specifies that physical disks must have the specified device type (FC for Fast Class, NL for Nearline, SSD for Solid State Drive, SCM for Storage Class Memory) to be used. Device types can be displayed by issuing the `showpd` command. If it is not specified, the default device type is FC.

-rpm <number>

Disks must be of the specified speed. Device speeds are shown in the RPM column of the showpd command. For FC and NL drives, the number corresponds to both a performance measure and actual rotational speed. For SSD & SCM drives, this field is "N/A".

Uses disks that satisfy all of the specified characteristics.

For example, `-p -fc_gt 60 -fc_lt 230 -nd 2` specifies all the disks that have greater than 60 and less than 230 free chunklets and ar connected to node 2 through their primary path.

Specifiers

<CPG_name>

Specifies the name of the common provisioning group being created.

Authority

- Super, Edit
- Any role granted the `cpg_create` right

Usage

- Requires access to all domains.
- For this command:
  - KB = 1024 bytes
  - MB = 1048576 bytes
  - GB = 1073741824 bytes
- When using the `createcpg` command, the size of the logical disk space created is the first integer multiple of the RAID set size that is large enough to accommodate the requested virtual volume size.

  For example, with the default RAID-5 layout with a set size of 768 MB, a requested virtual volume size of 8192 MB causes the creation of logical disks with a total size rounded up to an integer multiple of 768 that is 8448 MB. The growth increment of CPGs is similarly rounded up because the growth is done by creating logical disks that must be created in units of the logical disk RAID set size.
- By default, logical disks are created using only physical disks with the same device type. (By default, the Fast Class device type is used). Use the "`-p -devtype NL`", "`-p -devtype SSD`" to override this default. Use **`showpd`** to see the device types of physical disks in the system.
- The options `-sdgs`, `-sdgl` and `-sdgw` control the auto logical disk creation for the Common Provisioning Group's snapshot data regions. Auto logical disk creation occurs when the amount of free logical disk space falls below the specified grow(enlarge) size setting options (`-sdgs`, `-sdgl`).
- If auto-grow is enabled, new logical disks will be created automatically in addition to any logical disks specified with `-sa` or `-sd` options.
- The following options, `-t`, `-ssz`, `-rs`, `-ss`, `-ha`, `-ch`, and `-p`, are used to control auto logical disk creation (if auto-grow is enabled).

- If no device type is specified using the `-p devtype` option, Fast Class is assumed.
- The following arguments, `-tc_gt`, `-tc_lt`, `-fc_gt`, `-fc_lt`, `-devid`, and `-devtype` are used to select the disks that are used to create CPG based on the characteristics of the disk.

Examples

Creating a CPG (`cpg1`) and setting logical disk storage to auto-grow at 16 GB, setting a growth limit of 32 GB, and providing a growth warning at 24 GB:

```
cli% createcpg -sdgs 16g -sdgl 32g -sdgw 24g cpg1
```

Creating a CPG (`cpg2`) for all disks with greater than 60 and fewer than 1230 free chunklets, and that connect to node 2 through the primary path:

```
cli% createcpg -p -fc_gt 60 -fc_lt 1230 -nd 2 cpg2
```

Creating the CPG (`cpg3`) for all disks in cages 1 and 2 where magazine 4 of each cage is used in the CPG:

```
cli% createcpg -p -cg 1-2 -p -mg 4 cpg3
```

createdomain

Syntax

`createdomain [options] <domain_name>`

Description

The `createdomain` command creates system domains.

Options

`-comment <comment>`

Specifies any comments or additional information for the domain. The comment can be up to 511 characters long. Unprintable characters are not allowed. The comment must be placed inside quotation marks if it contains spaces.

`-vvretentiontimemax <time>{h|H|d|D}`

Specifies the maximum value that can be set for the retention time of a volume in this domain. <time> is a positive integer value and in the range of 0 - 43,800 hours (1825 days). Time can be specified in days or hours providing either the 'd' or 'D' for day and 'h' or 'H' for hours following the entered time value.

To disable setting the volume retention time in the domain, enter 0 for <time>.

Specifiers

`<domain_name>`

Specifies the name of the domain you are creating. The domain name can be no more than 31 characters. The name "all" is reserved.

Authority

- Super
- Any role granted the `domain_create` right

Usage

- Access to all domains is required to run this command.
- Use of this command requires a Domains license. Contact your local HPE representative for information.

Examples

The following example displays the creation of domain Engineering with an informational comment:

```
cli% createdomain -comment "This is a test domain." Engineering
```

createdomainset

Syntax

```
createdomainset [options] <setname> [<domain | pattern>...]
```

Description

The `createdomainset` command defines a new set of domains and provides the option of assigning one or more existing domains to that set. The command also allows the addition of domains to an existing set by use of the -add option.

Options

`-add`

Specifies that the domains listed should be added to an existing set. At least one domain must be specified.

`-comment <comment>`

Specifies any comment or additional information for the set. The comment can be up to 255 characters long. Unprintable characters are not allowed.

Specifiers

`<setname>`

Specifies the name of the domain set to create or add to, using up to 27 characters in length.

`<domain | pattern>...`

Optional list of domains or glob-style patterns matching domains that are to be included in the domain set. If no <domain> is specified a domain set with no domains is created.

Authority

- Super, Edit
- Any role granted the `domainset_create` right can create a domain set
- Any role granted the `domainset_set` right can add domains to a domain set

Usage

Access to all domains is required to run this command.

Examples

To create a domain set with 2 domains in it:

```
cli% createdomainset domset domain1 domain2
```

An alternative method to create a domain set:

```
cli% createdomainset set:domset domain1 domain2
```

To create an empty domain set with a comment:

```
cli% createdomainset -comment "A placeholder set" emptyset
```

createflashcache

Syntax

```
createflashcache [options] <size>
```

Description

Creates flash cache LDs for each node pair.

Options

-nocheck_scm_size

Overrides the size comparison check to allow the creation of Adaptive Flash Cache when SCM devices size are mismatched.

-sim

Specifies running the Adaptive Flash Cache in simulator mode. The simulator mode does not require the use of SSD drives.

-t *<RAID_type>*

Specifies the RAID type of the logical disks for Flash Cache (r0 for RAID-0 or r1 for RAID-1). With no RAID type specified, the storage system chooses the default.

Specifiers

***<size>*[g|G|t|T]**

Specifies the size for the flash cache in MiB for each node pair. Flash cache size is an integer that is a multiple of 16384 (16GiB). The minimum size of the flash cache is 64GiB. The maximum usable size of the flash cache is based on the node types, ranging from 384GiB up to 12288GiB (12TiB). An optional suffix (with no whitespace before the suffix) modifies the units to GiB (g or G suffix) or TiB (t or T suffix).

Authority

- Super, Edit
- Any role granted the `flashcache_create` right

Usage

This command creates flash cache on either SSD or SCM, depending on the array setup. The command uses SSD when FC or NL drives are present, and requires the size option. Support for SCM is available on all flash arrays. You cannot use the size option when SCM is present. Instead, the system selects size automatically to maximize usage.

Examples

The following examples show how to create Flash Cache for SSD and SCM.

For SSD:

```
createflashcache 128g
```

For SCM:

```
createflashcache
```

createfpg

Syntax

```
createfpg -recover [-wait] [-altname <altname>] {<vv> ... | set:<vvset>}
createfpg [options] <cpgname> <fpgname> <size>{t|T|g|G}
```

Description

The `createfpg` command creates an FPG of the given name and size within the specified CPG.

Options

`-comment <comment>`

Provides the textual description of the FPG.

`-full`

Creates the FPG using fully provisioned volumes.

`-tdvv`

Creates the FPG using thinly deduplicated volumes.

`-node <nodeid>`

Binds the created FPG to the specified node.

`-recover`

Recovers the FPG which is involved in Remote DR and that was removed using the `-forget` option.

If VVs are specified after the `-recover` option, they will be attached and any FPGs on them will be discovered. VV sets may also be specified using the `set:<setname>` modifier.

`-altname`

This option can only be used in combination with the `-recover` option. This option renames the FPG as part of the recovery process, to allow the newly recovered FPG to coexist with another FPG sharing the same name. This is generally used to enable recovery of individual files from a snapshot copy of an FPG, while the original FPG remains activated. When this option is used, only a single FPG can be recovered at once, so either a single VV or a VV set containing multiple VVs for a single FPG can be specified.

`-wait`

Waits until the associated task is completed before proceeding. This option produces verbose task information.

`-rcopygroup <rcopy_group>`

Creates the FPG and automatically add it to the specified Remote Copy group.

Specifiers

<cpgname>

Specifies the CPG where the VVs associated with the FPG are created.

<fpgname>

Specifies the name of the FPG to create.

<size>

Specifies the size of the FPG to create, between 1T and 128T.

A suffix (with no whitespace before the suffix) modifies the units to TiB or GiB (t, T, g, or G suffix).

Authority

- Super, Edit
- Any role granted the `fpg_create` right.

Usage

- Use of this command requires access to all domains.
- Each `createfpg` command creates the following elements. The command fails if any of these elements are not available.

  Number of VVs (size determined by the FPG size) *<fpgname>*`.n`.
- If no node is specified, `createfpg` will automatically determine the node with least amount of storage attached and bind the new FPG to this node.
- The maximum length of a FPG name is 21 characters.
- The maximum length of a FPG name in a Remote Copy group is 18 characters.
- FPG names must not end in a `.n` format.
- FPG names must be unique across systems using Remote Copy for replication of those FPGs. Using duplicate names across systems will result in NFS and Object file shares being unavailable upon recovery on the target system.
- Thin Persistence is disabled as it is not applicable to fully provisioned volumes. If the FPG is later converted to thinly provisioned volumes, Thin Persistence can be enabled using `setfpg -unmap enable`. If an FPG is later converted to fully provisioned volumes, Thin Persistence should be disabled first using `setfpg -unmap disable` to avoid unnecessary overhead.
- The minimum size for an FPG is 1TiB.

Examples

Creating an 18T FPG:

```
cli% createfpg examplecpg examplefs 18T
```

Attaching `fpg.1` to File Person, and recovering the FPG:

```
cli% createfpg -recover fpg.1
```

Attaching the VVs in set `fpg1` to File Persona, and recovering any discovered FPGs:

```
cli% createfpg -recover set:fpg1
```

createfsgroup

Syntax

`createfsgroup [options] <groupname>`

Description

The `createfsgroup` command creates a new, local group associated with File Persona.

Options

`-gid <number>`

Specifies the group ID to be used. GIDs are numbers starting from zero and grouped as follows. Groups can be referred to by either groupname or GID.

- **0-99** – Globally allocated by the Debian project and used for system accounts. The `root` account, for instance, is GID 0.
- **100-4294967294** – Normal groups.
- **65534** – The **nobody** group, with no rights or permissions.

`-memberlist <string>`

Specifies members of the group.

Specifiers

`<groupname>`

Specifies the local group name using up to 256 characters. Valid characters are alphanumeric characters, periods, dashes (except first character), and underscore.

Authority

- Super
- Any role granted the right to `fsgroup_create`

Usage

- Use of this command requires access to all domains.
- The `-gid` option can have any value between 100 and 4294967294.
- To access an SMB share, specify the group as `LOCAL_CLUSTER\<groupname>`.

Examples

Creating a local group with the name `accounting`, and specifying a GID:

```
cli% createfsgroup -gid 3000 accounting
```

createfshare

Syntax

`createfshare {smb|nfs|obj|ftp} [options <arg>] <vfs> <sharename>`

Description

The `createfshare` command creates file shares for supported protocols.

Options

`-fpg <fpgname>`

Specifies the File Provisioning Group (FPG) to which the `<vfs>` belongs.

With no FPG specified, the command discovers the FPG based on the specified `<vfs>`. If the `<vfs>` exists under multiple FPGs, you must specify the `-fpg <fpgname>`.

`-fstore <fstore>`

Specifies the file store under which to create the share. With no `fstore` specified, the command uses the `<sharename>` as the file store name. The command creates the file store if it does not exist. If you specify this option to create a file share, you must specify it when you set or remove the share using **`setfshare`**/**`removefshare`**.

-sharedir *<sharedir>*

Specifies the directory path to share. It can be a full path starting from "/", or subdirectory of the file store that the share will expose network clients. Do not specify "." or ".." for the subdirectory. If not specified, the share created will be rooted at the file store. If this option is specified, option `-fstore` must be specified.

-comment *<comment>*

Specifies any comments or additional information for the share. The comment can be up to 255 characters long. Unprintable characters are not allowed.

-f

Specifies that the command is forced. When creating a share of a second protocol type for a given file store, if this option is not used, the command requires confirmation before proceeding with its operation.

-ftp

Creates an FTP file share.

The following options are specific to **smb** subcommand:

-abe {true|false}

Access-Based Enumeration. Specifies whether users can see only the files and directories to which they have been allowed access on the shares. Defaults to `false`.

-allowip *<iplist>*

Specifies client IP addresses that are allowed access to the share. Use commas to separate the IP addresses. Defaults to " ", which allows all IP addresses (empty means all are allowed).

-denyip *<iplist>*

Specifies the client IP addresses that are denied access to the share. Use commas to separate the IP addresses. Defaults to " ", which denies access to no IP addresses (empty means none are denied).

-allowperm *<permlist>*

Specifies the type of permission that a user or group is allowed on the share. *<permlist>* must use the format: *<user1>*:*<perm>*,*<user2>*:*<perm>*,*...*. The *<user>* variable can be a user name or a group name. The *<perm>* variable must use **fullcontrol**, **read**, or **change**.

`Everyone` is a special user for all users and groups.

If the user is configured locally using `createfsuser`, use `<user>` to specify the user (for example, `-allowperm user1:fullcontrol`).

If the user is configured on Active Directory, use `setfs ad` to join Active Directory domain with *<domain>* if it has not been done, and use *<domain>**<user>* or *<ad-netbios>**<user>* to specify the user name (for example, `-allowperm example.com\\aduser:fullcontrol`). The `<ad-netbios>` can be found by running "`showfs -ad`".

If the user is configured on the LDAP server, use "`setfs ldap`" to create LDAP configuration with `<ldap-netbios>` if it has not been done, and use "`<ldap-netbios>\\<user>`" to specify the user (for example, `-allowperm ldaphost\\ldapuser:read`).

If not specified, no default permissions are allowed for the new shares, which set the same default as that of a Windows Server 2012 R2 server. This prevents inadvertent access to the SMB share by a user who is not explicitly specified.

-denyperm *<permlist>*

Specifies a denied permission for a user or group with access to the share. Specify the *<permlist>* variable in the format *<user1>*:*<perm>*,*<user2>*:*<perm>*,*...*". The `<user>` variable can be a user name or a group name. The *<perm>* variable must use **fullcontrol**, **read**, or **change**.

`Everyone` is a special user for all users and groups.

If the user is configured locally using "`createfsuser`", use `<user>` to specify the user (for example, `–denyperm user1:fullcontrol`).

If the user is configured on Active Directory, use "`setfs ad`" to join Active Directory domain with `<domain>` if it has not been done, and use "`<domain>\\<user>`" or "`<ad-netbios>\\<user>`" to specify the user (for example, `–denyperm example.com\\aduser:fullcontrol`). The `<ad-netbios>` can be found running "`showfs -ad`".

If the user is configured on the LDAP server, use "`setfs ldap`" to create LDAP configuration with `<ldap-netbios>` if it has not been done, and use "`<ldap-netbios>\\<user>`" to specify the user (for example, `–denyperm ldaphost\\ldapuser:read`).

`-audit {operation1:value1[,operation2:value2]...}`

Specifies the operations to be audited at share level for SMB protocol on specified VFS. The event can be audited based on values: success, failure, none or all. The operations and values are case insensitive, and generate events under the following circumstances:

- TreeConnect: Connected to a share. To create an audit record for a user trying to access a share that is not configured on the VFS, use the command: "`setfsaudit smb -audit Global:failure`".
- TreeDisconnect: Disconnected from a share. Note: Certain clients that use SMBv1 may not trigger a TreeDisconnect when the client logs off.
- Open: Open/Create a file or directory.
- Close: Close a file or directory.
- Rename: Rename a file or directory.
- Delete: Delete a file or directory.
- Read: Read from a file. Only the first read after the file was opened is reported.
- Write: Write to a file. Only the first write after the file was opened is reported.
- ChangeSecurity: Change security attributes.

`-cache {off|manual|optimized|auto}`

Specifies client-side caching for offline files. Valid values are:

- "off": The client must not cache any files from this share. The share is configured to disallow caching.
- "manual": The client must allow only manual caching for the files open from this share.
- "optimized": The client may cache every file that it opens from this share. Also, the client may satisfy the file requests from its local cache. The share is configured to allow automatic caching of programs and documents.
- "auto": The client may cache every file that it opens from this share. The share is configured to allow automatic caching of documents. If this is not specified, the default is "manual".

`-ca {true|false}`

Specifies if SMB3 continuous availability features must be enabled for this share. If not specified, the default is "true".

`-options <options>`

Specifies options to use for the share to be created. Standard NFS export options except "no_subtree_check" are supported. Do not enter option "fsid", which is provided. If not specified, the following options will be automatically set: sync, auth_nlm, wdelay, sec=sys, no_all_squash, crossmnt, secure, subtree_check, hide, root_squash, ro.

The following options are specific to **nfs** subcommand:

- See the `linux exports(5)` man page for detailed information.
- `-audit {operation1:value1[,operation2:value2]...}`

  The following operations can be audited on a share level for NFS protocol on specified VFS. Acceptable value is either on or off. Defaults to off. Operations are case insensitive.

  These operations generate an event under the following circumstances:

  - audit_open: Open a file. This is NFSv4 only.
  - audit_close: Close a file. This is NFSv4 only.
  - audit_meta: Rename a file or directory, creating a hardlink for a file, creating a file or directory and removing a file or a directory.
  - audit_attr: Changing file/directory attributes.
  - audit_read: Read of a file/directory. Subsequent file reads are suppressed for a specific time (currently 10m). Note that reads from a directory are never suppressed.
  - audit_write: Write to a file. Subsequent writes are suppressed for a specific time (currently 10m).
  - audit_acl: Allow audit for ACL (AccessControlList) changes on a file or directory.

- `-clientip <clientlist>`

  Specifies the clients that can access the share. The NFS client can be specified by the name (for example, sys1.example.com), the name with a wildcard (for example, *.example.com), or by its IP address. Use comma to separate the IP addresses. If this is not specified, the default is "*".

The following options are specific to **obj** subcommand:

- `-ssl {true|false}`

  Specifies if SSL is enabled. The default is "false".

- `-urlpath <urlpath>`

  Specifies the URL that clients will use to access the share. If this is not specified, the command uses `<sharename>` as `<urlpath>`.

The following options are specific to **ftp** subcommand:

- `-shareip <iplist>`

  Specifies the IP addresses to be assigned to the FTP share. The IP addresses must be assigned to the specified VFS before creating the share. Use commas to separate multiple IP addresses. If this option is not specified all the available IPs from the VFS which are not assigned to any other FTP shares will be used to create the share.

- `-ssl {true|false}`

  Specifies if SSL is enabled. The default is "false".

- `-options <option1=value1>[,<option2=value2>]...`

  Specifies the configuration options to be used for the FTP share. If not specified, all options will be set to default values as shown below.

  - accept_timeout=60
  - allow_anon_ssl=YES
  - anon_max_rate=0

- anon_mkdir_write_enable=YES
- anon_other_write_enable=NO
- anon_umask=0077
- anon_upload_enable=YES
- anon_world_readable_only=NO
- ascii_mode=both
- chown_upload_mode=0600
- connect_timeout=60
- connection_mode=passive
- data_connection_timeout=300
- dirlist_enable=YES
- file_access_mode=both
- file_open_mode=0777
- force_anon_logins_ssl=NO
- force_anon_data_ssl=NO
- force_local_data_ssl=NO
- force_local_logins_ssl=YES
- hide_ids=NO
- home_dir_prefix=/
- home_dir_support=NO
- idle_session_timeout=300
- implicit_ssl=YES
- local_max_rate=0
- local_umask=0022
- lock_upload_files=YES
- login_access_mode=local
- max_clients=250
- mdtm_write=YES
- no_anon_password=YES
- pasv_max_port=49500
- pasv_min_port=49251
- require_ssl_reuse=NO
- ssl_version=tlsv1
- text_userdb_names=NO

All the above values can be modified after the share is created except "home_dir_prefix=/" and "home_dir_support=NO".

Specifiers

The following specifiers are for **all** subcommands:

<vfs>

The virtual file server (VFS) under which the file store, if it does not exist, and the share will be created.

<sharename>

The share name to be created.

Subcommands

smb

Creates an SMB file share.

nfs

Creates an NFS file share.

obj

Creates an Object file share.

NOTE: IMPORTANT: The Object Access API has been deprecated and will be removed in a future release.

---

ftp

Creates an FTP file share.

Authority

- Super, Edit
- Any role granted the right to `fshare_create`

Usage

- Access to all domains is required to run this command.
- The FPG and it's underneath VFS must be created before creating file shares.
- For SMB share permissions, the same user cannot be specified with the same permission in both "`-allowperm`" and "`-denyperm`". Share permissions are different from file system permissions, that is, ACLs on files/folders. Option "`-allowperm`" grants users to access the share, which does not set the folder permissions of the folder exported by the share. If a user wants to read a file, he must have read share permissions and an ACE that grants him read permissions. Even if "Everyone:full control" is allowed, that does not mean that anyone can do anything, due to the additional restrictions of the folder ACLs. To modify the folder ACL or permissions, connect to the share through the protocol and use the appropriate mechanism for that protocol to modify the ACL or permissions to allow different access than what is provided by default.
- To access an SMB share: for users configured locally, specify "LOCAL_CLUSTER\<user>", for users configured on Active Directory, specify "<domain>\<user>" or "<ad-netbios>\<user>", for users configured on the LDAP server, specify "<ldap-netbios>\<user>".
- For NFS shares, it is not allowed to create two shares which have identical clients (that is, specified by -clientip) and share directory (that is, specified by `-sharedir`). If you create NFS shares without specifying different `-clientip` and `-sharedir` options, the second "`createfshare`" will fail.
- To create FTP/Object share, the VFS specified by `<vfs>` must have an associated IP address.
- This command displays a warning message:

- If user tries to create NFS/OBJ/FTP shares on a file store with "ntfs" security mode.
- If user tries to create shares of different protocols on a file store with "legacy" security mode.

- For the commands which involve implicit creation of file stores:
  - If the share is SMB, a warning message is displayed and the file store will be created in "ntfs" security mode after the user confirmation.
  - If the share is NFS/OBJ/FTP, no warning message is displayed and the file store will be created in "legacy" security mode.
- For obj/ftp shares, the VFS must have a certificate associated with it to enable SSL.

Examples

The following example creates SMB file share "myshare":

```
cli% createfshare smb -fstore myfstore myvfs myshare
```

The following example creates SMB share share1 with auditing enabled for successful TreeConnect operation through SMB protocol:

```
cli% createfshare smb -audit TreeConnect:success vfs1 share1
```

The following example creates NFS file share:

```
cli% createfshare nfs -options ro -fstore myfstore myvfs myshare
```

The following example creates NFS share "myshare" with auditing options for open, write operations through NFS protocol:

```
cli% createfshare nfs -audit audit_open:on,audit_write:on -fstore \
myfstore myvfs myshare
```

The following example creates Object file share:

```
cli% createfshare obj -fstore myfstore myvfs myshare
```

The following example creates FTP file share with implicit file store creation:

```
cli% createfshare ftp -ssl true myvfs myshare
```

The following example creates FTP file share with share IP as 1.1.1.1:

```
cli% createfshare ftp -shareip 1.1.1.1 -fstore myfstore myvfs myshare
```

createfsip

Syntax

```
createfsip [options] <ipaddr> <subnet> <vfs>
```

Description

The `createfsip` command assigns an IP address to a virtual file system (VFS).

Options

`-vlantag <tag>`

Specifies the VLAN Tag to be used.

-fpg <fpg>

Specifies the file provisioning group (FPG) in which the VFS was created.

-subnetoverlap

Allow overlapping of subnets.

NOTE:

This is not a typical configuration. If not properly configured, the configuration can cause data unavailability. If option is not specified, overlapping subnets are not allowed.

-f

Specifies that the command is forced. If this option is not specified, the command requires confirmation before proceeding with its operation.

Specifiers

<vfs>

Specifies the VFS to which the IP address will be assigned.

<ipaddr>

Specifies the IP address to be assigned to the VFS.

<subnet>

Specifies the subnet mask to be used.

Authority

- Super, Edit
- Any role granted the `fsip_create` right.

Usage

Access to all domains is required to run this command.

Examples

Assigning the IP address 10.10.10.1 to VFS vfs_1:

```
cli% createfsip 10.10.10.1 255.255.255.0 vfs_1
```

Assigning specified overlapping subnets and an IP address to a specified VFS:

```
cli% createfsip -fpg examplefpg -subnetoverlap 10.10.10.10 255.255.0.0
examplevfs
```

createfsnap

Syntax

```
createfsnap [options <arg>] <vfs> <fstore> <tag>
```

Description

The `createfsnap` command creates a file store snapshot for File Persona.

Options

-retain <rcnt>

Number of snapshots to retain with the specified tag. Snapshots exceeding the count will be deleted, oldest first. The valid range of <rcnt> is from 1 to 1024.

-f

Do not ask for confirmation before creating a snapshot with retention count (-retain). This option is ignored if option -retain is not specified.

-fpg <fpgname>

Specifies the file provisioning group (FPG) that <vfs> belongs.

If this is not specified, the command will find out the FPG based on the specified <vfs>. However, if <vfs> exists under multiple FPGs, -fpg must be specified.

Specifiers

<vfs>

Specifies the name of the virtual file server (VFS).

<fstore>

Specifies the name of the file store that the snapshot will be taken.

This is the path relative to <vfs>.

<tag>

Specifies the suffix to be appended to the timestamp of snapshot creation time in ISO 8601 date and time format, which will become the name of the created file store snapshot (for example: if "snapshot1" is being used as <tag>, the snapshot name will be 2013-12-17T215020_snapshot1).

The name can be used as the value of option -snapname to display or remove a snapshot.

Authority

- Super, Edit
- Any role granted the `fsnap_create` right

Usage

- Access to all domains is required to run this command.
- If option -retain is specified and the file store already has the maximum number of snapshots taken, the oldest snapshot will be deleted first before the new snapshot is created. If the command fails to create the new snapshot, the deleted snapshot will not be restored.

Examples

The following example creates a file store snapshot:

```
cli% createfsnap -fpg myfpg myvfs myfstore snapshot1
```

createfspn

Syntax

```
createfspn [options] <vfs> <user> <spnname>
```

Description

The `createfspn` command creates SPN for a given VFS to enable Kerberos authentication for File Persona.

Options

-fpg <fpgname>

Specifies the file provisioning group (FPG) that <vfs> belongs to. If not specified, the command will find out the FPG based on the specified <vfs>. However, if <vfs> exists under multiple FPGs, -fpg must be specified.

-passwd <password>

Specifies the password of the authorized user. If not specified, the user is prompted for a cleartext password.

Specifiers

<vfs>

Specifies the VFS name.

<user>

Specifies the authorized user to access Active Directory domain controller.

<spnname>

Specifies the name of SPN to be created.

Authority

- Super, Edit
- Any role granted the fspn_create right.

Usage

- Access to all domains is required to run this command.
- The Kerberos realm must be a Microsoft AD.
- DNS must hold an entry for the hostname that corresponds to SPN name.

Examples

Creating an SPN for a specified VFS:

```
cli% createfspn -fpg fpg1 -passwd password vfs1 aduser spn1
```

createfsroute

Syntax

```
createfsroute [-vlan <vlantag>] <targetaddr> <subnet>|<prefixlen> <gateway>
```

Description

The createfsroute command creates a route for a target with non-default gateway.

Options

-vlan <vlantag>

A VLAN tag with values ranging from 0 to 4094. Defaults to 0.

Specifiers

<targetaddr>

The target IPv4/IPv6 address for which the gateway is to be assigned.

<subnet>|<prefixlen>

The subnet or prefix length for the target IP Address.

<gateway>

Gateway to be assigned to the target IP address.

Authority

- Edit
- Any role granted `fsroute_set` right

Usage

Access to all domains is required to run this command.

Examples

The following example creates a route with vlan tag 10:

```
cli% createfsroute -vlan 10 10.16.101.11 255.255.0.0 10.16.101.1
```

The following example creates a route at vlan 0:

```
cli% createfsroute 10.16.101.11 255.255.0.0 10.16.101.1
```

createfstore

Syntax

```
createfstore -secmode {ntfs|legacy} [-f] [-comment <comment>] [-fpg
<fpgname>] [-secop_errsuppress {true|false}] <vfs> <fstore>
```

Description

The `createfstore` command creates a new file store with the specified name for the specified file provisioning group (FPG) and the virtual file server (VFS).

Options

-fpg <fpgname>

Specifies the name of the FPG.

-comment <comment>

Specifies the textual description of the file store.

-secmode {ntfs|legacy}

Specifies the security mode of the file store. Valid values are:

- `legacy`: File stores upgraded from 1.0 or 1.1 will be in this mode. New file stores can also be created in this mode by specifying "legacy".
- `ntfs`: Windows clients will get native NTFS ACL enforcement. Non-windows clients can create/read/write/delete/rename files. Non-windows clients can also read the Permissions. But they will not be able to set Permissions, change owner or change group except when the user trying to perform this operation is "root".

-secop_errsuppress {true|false}

In "ntfs" security mode, permission changing operations will report error messages for non-root users when this option is "false". This may not be the ideal behavior for some applications which are trying to perform permission

changing operations such as "chmod" implicitly over an NFSv3 client. This option can be set to "true" to allow those applications to work with NFSv3 clients. If not specified, the option will be set to "false". This option can be used only while creating file stores in "ntfs" security mode.

Specifiers

<vfs>

Specifies the name of the virtual file server.

<fstore>

Specifies the name of the file store to be created.

Authority

- Super, Edit
- Any role granted `fstore_create` right

Usage

- Access to all domains is required to run this command.
- The file store name cannot exceed 255 characters.
- If clients from multiple protocols access NTFS file store simultaneously, it is recommended to set "enableoplocks" and "enablesmbleases" to "false" (use "`setfs smb`" subcommand). Otherwise SMB clients can cache file data locally. This could cause a client from a non-SMB protocol to see stale data (If an SMB client updated a file but is keeping the update in its cache) or it could cause an SMB client to see stale data (If a non-SMB protocol client updated a file after a SMB client had already read its contents and cached the data). This behavior can be prevented by setting "enableoplocks" and "enablesmbleases" to "false". Note that changing these settings may have a performance impact for SMB, depending on the workload, but it is recommended if there will be simultaneous access by clients of multiple protocols.

Examples

The following example creates a file store named "samplefstore" in "ntfs" security mode under virtual server named "samplevfs":

```
cli% createfstore -secmode ntfs samplevfs samplefstore.
```

The following example creates a file store named "samplefstore" in "legacy" security mode on FPG named "samplefpg" under virtual server named "samplevfs" and sets a comment "Sample fstore" with the created file store:

```
cli% createfstore -secmode legacy -comment "Sample fstore" -fpg samplefpg \samplevfs samplefstore
```

createfsuser

Syntax

```
createfsuser [options] <username>
```

Description

The `createfsuser` command creates a new local user.

Options

-passwd <password>

Specifies the user's password.

-primarygroup <groupname>

Specifies the user's primary group.

-enable {true | false}

Specifies the user is enabled or disabled on creation.

-uid <userid>

Specifies the user ID to be used. UIDs are numbers from 0 to 4294967294. The numbers are grouped as follows:

- UIDs 0-99 are globally allocated by the Debian project and are used for system accounts. The "root" account, for instance, is UID 0.
- UIDs 100- 4294967294 are normal user accounts.
- UID 65534 is user "nobody", an account with no rights or permissions.

Users may usually be referred to by either username or UID.

-grplist <grouplist>

Specifies a list of additional groups the user is to be a member.

Specifiers

<username>

Specifies the local user name using up to 20 characters. Valid characters are alphanumeric characters, periods, dashes (except as the first character), and underscores.

Authority

- Super
- Any role granted the `fsuser_create` right.

Usage

- Access to all domains is required to run this command.
- If not specified, `-uid` will be given a default value.
- The `-uid` option can have any value between 1000 and 4294967294.
- User's primary group cannot be Built-in groups: "Administrators", "Users", "Guests" and "Backup Operators".
- The user and group names are case insensitive.
- If the `-enabled` option is not supplied the user will be enabled by default.
- If a password is not supplied the user will be prompted to enter one.
- To access an SMB share, specify the user as "**LOCAL_CLUSTER\<username>**".

Examples

Creating a local user named jdoe:

```
cli% createfsuser -grplist finance,marketing
    -primarygroup "local users" jdoe
```

createfsuser

Syntax

```
createsv [option] <copy_of_VV>[:[<snapshot_VV>][:[<ID>][:[<WWN>][:{RO|
RW}]]]]...
```

Description

The `createsv` command creates consistent group snapshots of a list of virtual volumes. Consistent group snapshots are all created at the same point in time so that related structures on different volumes in the group remain consistent.

Options

`-ro -match`

By default, all snapshots are created read-write. The `-ro` option instead specifies that all snapshots created will be read-only. The `-match` option specifies that snapshots are created matching each parent's read-only or read-write setting. The `-ro` and `-match` options cannot be combined. Either of these options can be overridden for an individual snapshot VV in the colon-separated specifiers.

`-exp <time>{d|D|h|H|m|M}`

Specifies the relative time from the current time that volume will expire. `<time>` is a positive integer value and in the range of 1 minute - 1825 days. Time can be specified in days, hours, or minutes. Use "d" or "D" for days, "h" or "H" for hours, and "m" or "M" for minutes following the entered time value.

`-comment <comment>`

Specifies any additional information up to 511 characters for the volume.

`-f`

Do not ask for confirmation before creating a volume with retention time (`-retain`).

`-retain <time>{d|D|h|H|m|M}`

Specifies the amount of time, relative to the current time, that the volume will be retained. `<time>` is a positive integer value and in the range of 1 minute - 1825 days. Time can be specified in days, hours, or minutes. Use "d" or "D" for days, "h" or "H" for hours, and "m" or "M" for minutes following the entered time value.

NOTE: If the volume is not in any domain, then its retention time cannot exceed the value of the system's VVRetentionTimeMax. The default value for the system's VVRetentionTimeMax is 14 days. If the volume belongs to a domain, then its retention time cannot exceed the value of the domain's VVRetentionTimeMax, if set. The retention time cannot be removed or reduced once it is set. If the volume has its retention time set, it cannot be removed within its retention time. If both expiration time and retention time are specified, then the retention time cannot be longer than the expiration time.

This option requires the HPE 3PAR Virtual Lock license. Contact your local service provider for more information.

`-addtoset <VV_set>`

Specifies that the snapshots will be added to the specified set. The set will be created if it does not exist.

`-noblock`

When a read-only snapshot is taken, the storage system blocks host i/o to the parent virtual volume while the snapshot is being created. If the `-noblock` option is specified then the storage system will not block host i/o to the parent virtual volume during snapshot creation. This option is only applicable for read-only snapshots.

Specifiers

`<copy_of_VV>`

Specifies the name of the Virtual Volume being copied (the parent volume).

<snapshot_VV>

Optional name of the snapshot Virtual Volumes. If <snapshot_VV> is not specified, the system tries to automatically generate the snapshot name of the form <copy_of_VV>.<type><number> where <type> is either ro or rw and <number> is either empty or the lowest number starting from 0 that does not conflict with an existing volume name. If the generated name is too long (because the <copy_of_VV> name is too long) the command will fail, and the user is required to specify <snapshot_VV> explicitly.

<ID>

Optional ID of the <snapshot_VV>. If the <ID> is not specified, an ID is chosen automatically.

<WWN>

Optional WWN of the snapshot VV. If <WWN> is not specified, a WWN will be chosen automatically.

{RO|RW}

Optional read-only or read-write setting to be applied only to the individual snapshot VV. Specifying this setting overrides the –ro and –match options for this snapshot.

Authority

- Super, Edit
- Any role granted the groupsv_create right

Usage

- If the volume is not in any domain, then its retention time cannot exceed the value of the system's VVRetentionTimeMax. The default value for the system's VVRetentionTimeMax is 14 days. If the volume belongs to a domain, then its retention time cannot exceed the value of the domain's VVRetentionTimeMax, if set. The retention time cannot be removed or reduced once it is set. If the volume has its retention time set, it cannot be removed within its retention time. If both expiration time and retention time are specified, then the retention time cannot be longer than the expiration time. This option requires the HPE 3PAR Virtual Lock License. Contact your local HPE 3PAR Authorized Service Provider for more information.
- Consistent group snapshots are all created at the same point in time so that related structures on different volumes in the group remain consistent.
- A maximum of 2048 virtual copies per volume are allowed. Creating a read-only snapshot from an existing read-only snapshot, or creating a read-write snapshot from a read-write snapshot or the base will create a hidden snapshot, therefore further reducing the available number of virtual copies for that base volume.

Examples

In the following example, virtual volumes VV1 and VV2 are created using the **createvv** command. Then the createsnapgroup command is used to create a read-only snapshot of each of the volumes and display a list of the snapshots:

```
cli% createvv -snp_cpg cpg1 cpg1 VV1 4g
cli% createvv -snp_cpg cpg1 cpg1 VV2 4g
cli% creategroupsv -ro VV1 VV2
```

In the following example, snapshots of VV1 and VV2 are created again, one read-only and one read-write. The first snapshot has an automatic ID and WWN, while the second example snapshot specifies an ID and WWN.

```
cli% creategroupsv VV1:vv1.ro:::RO VV2:vv2.rw:321:50002AC0001A0024:RW
```

creategroupvvcopy

Syntax

The syntax for the createvvcopy command can be one of the following:

```
createvvcopy -p [options] <parent_VV>:<destination_VV>...

createvvcopy -p -online [options]
<parent_VV>:<destination_cpg>:<VV_name>[:<wwn>]...

createvvcopy -r [options] <destination_VV>...

createvvcopy -halt <destination_VV>...
```

Description

The createvvcopy command creates consistent group physical copies of a list of virtual volumes.

Options

-p

Starts a copy operation from the specified parent volume (as indicated using the <parent_VV> specifier) to its destination volume (as indicated using the <destination_VV> specifier). First a set of consistent group snapshots of all the <parent_VV> specifiers are taken and then each snapshot is copied to the corresponding <destination_VV> specifier.

After the copies are complete, the snapshots are deleted unless the -s option is specified. Each <parent_VV> specifier must be a base virtual volume or a read/write snapshot.

-r

Resynchronizes the set of destination volumes (as indicated using the <destination_VV> specifier) with their respective parents using saved snapshots so that only the changes made since the last copy or resynchronization are copied. The resynchronization operation can only be performed if snapshots of the parents saved using the -s option in an earlier instance of the createvvcopy command are present. These old snapshots are replaced by a new set of consistent group snapshots for the next resynchronization operation (such as the -s option need not be specified with the -r option).

-halt

Cancels an ongoing physical copy. This causes the destination volume (as indicated using the <destination_VV> specifier) to be marked with the "cpf" status, which can be cleared up when a new copy is started.

-s

Saves snapshots of the parent volume (as indicated with the <parent_VV> specifier) for quick resynchronization and to retain the parent-copy relationships between each parent and destination volume. The -s option is implied and need not be specified when the -r option is used. Each <destination_VV> specifier remains marked as a physical copy of its <parent_VV> specifier until it is promoted to a base virtual volume using the promotevvcopy command, which also removes the saved snapshot of the <parent_VV> specifier. The saved snapshots should not be removed manually. If the same <parent_VV> specifier is copied to different <destination_VV> specifiers with the -s option, a different snapshot of the <parent_VV> specifier is saved for each <destination_VV> specifier.

-b

Uses this specifier to block until all the copies are complete. Without this option, the command completes before the copy operations are completed (use the showvv command to check the status of the copy operations).

-pri <high|med|low|auto>

Specifies the priority of the copy operation when it is started. This option allows the user to control the overall speed of a particular task. If this option is not specified, the createvvcopy operation is started with default priority of medium. Auto priority can be used only if -online option is specified and this operation is started with auto as the default priority. High priority indicates that the operation will complete faster. Low priority indicates that

the operation will run slower than the default priority task. If the priority is set to auto, then the array controls the speed of the operation based on the load in the array for all flash models. Auto priority behaves like medium priority for the converged flash models of the arrays containing Nearline (NL) and Fast Class (FC) drives. This option cannot be used with -halt option.

-online

Specifies that the copy is to be performed online. This means that the destination can be immediately exported and is automatically created by this command, meaning that it cannot exist prior to executing this command. This option requires the specification of a CPG for the destination volume and allows for the options listed below. Finally, `-online` requires the `-p` option and is incompatible with the `-r`, `-halt`, `-s`, and `-b` options.

-skip_zero

When copying from a thin provisioned source, only copy allocated portions of the source VV. Using this option can save time, but only use it if the destination is newly created or has been re-initialized to zero. If there is pre-existing data in the destination VV it will not be overwritten to match the source VV if the same offset is not allocated in the source. This option cannot be combined with `-r` or `-halt`.

The following options can only be used when the `-online` option is specified:

-tpvv

Indicates that the VV the online copy creates should be a thinly provisioned volume. Cannot be used with the `-dedup` option.

-tdvv

This option is deprecated, see -dedup.

-dedup

Indicates that the VV the online copy creates should be a thinly deduplicated volume, which is a thinly provisioned volume with inline data deduplication. This option can only be used with a CPG that has SSD (Solid State Drive) device type. Cannot be used with the `-tpvv` option.

-compr

Indicates that the VV the online copy creates should be a compressed virtual volume.

-compr

Indicates that the VV the online copy creates should be a compressed virtual volume.

-snp_cpg <snp_cpg>

Specifies the name of the CPG from which the snapshot space will be allocated.

Specifiers

<destination_VV>

Indicates the destination virtual volume.

<parent_VV>

Indicates the parent virtual volume.

<destination_cpg>

Specifies the destination CPG to use for the destination volume if the `-online` option is specified.

<VV_name>

Specifies the virtual volume name to use for the destination volume if the `-online` option is specified.

<wwn>

Specifies the WWN to use for the destination volume if the `-online` option is specified.

Authority

- Super, Edit
- Any role granted the `groupvvcopy_create` right

Usage

- The source and destination volumes must be writable.
- The destination volume cannot be a system volume or a File Persona volume.
- The destination volume cannot be exported as a VLUN before or during the **`createvvcopy`** command process.
- The destination volume must be greater than or equal in size to the source volume.
- If the `-s` option is specified to save a snapshot for fast resynchronization and the snapshot goes stale, the copy fails.
- A physical copy of a virtual volume fails in any situation that a snapshot copy fails or when there is insufficient snapshot space or I/O errors.
- The `creategroupvvcopy` command can be issued multiple times. However, the HPE 3PAR storage system allows only two active physical copy tasks to run concurrently. Any additional physical copy tasks are queued, pending the completion of the active physical copy tasks. Consistent group snapshots are all created at the same point in time so that related structures on different volumes in the group remain consistent.
- The `-online` option for an online copy automatically creates a destination volume, so unlike other physical copies, the destination volume should not exist. The destination volume requires a CPG to be specified, and optionally it can be either a TPVV, TDVV, or have a snapshot CPG associated with it.
  Note that the `-tpvv`, `-dedup`, `-compr`, and `-snp_cpg` options apply to all copies. The destination volume can be immediately exported (hence the reason for the online option). The online option is incompatible with the `-r`, `-halt`, `-s`, and `-b` options.

Examples

None.

createhost

Syntax

```
createhost [options] <hostname> [<WWN>...]
createhost -iscsi [options] <hostname> [<iscsi_name>...]
```

Description

The `createhost` command defines a new system host and provides the option of assigning one or more paths to that host. Paths can be either Fibre Channel WWNs or iSCSI names. The command also provides options to annotate the host with descriptor information such as physical location, IP Address, Operating System, Model, etc. In addition, if the Smart SAN license is installed, the desired relationship between the host and array port for target-driven zoning can be specified.

Options

`-add`

Adds the specified WWN(s) or iscsi_name(s) to an existing host (at least one WWN or iscsi_name must be specified). Do not specify host persona.

-domain <domain | domain_set>

Creates the host in the specified domain or domain set. The default is to create it in the current domain, or no domain if the current domain is not set. The domain set name must start with "set:".

-f

Forces the tear down of lower priority VLUN exports if necessary.

-persona <hostpersonaval>

Sets the host persona that specifies the personality for all ports which are part of the host set. This selects certain variations in scsi command behavior which certain operating systems expect.

<hostpersonaval> is the host persona id number with the desired capabilities. These can be seen with `showhost -listpersona`.

-loc <location>

Specifies the host's location.

-ip <IP address>

Specifies the host's IP address.

-os <OS>

Specifies the operating system running on the host.

-model <model>

Specifies the host's model.

-contact <contact>

Specifies the host's owner and contact information.

-comment <comment>

Specifies any additional information for the host.

-port <node:slot:port>...|<pattern>...

Specifies the desired relationship between the array port(s) and host for target-driven zoning. Multiple array ports can be specified by either using a pattern or a comma-separated list. This option is used only when the Smart SAN license is installed. At least one WWN needs to be specified with this option.

Specifiers

<hostname>

Specifies the name of the host using up to 31 characters.

<WWN>

Specifies the World Wide Name (WWN) to be assigned or added to an existing host. This specifier can be repeated to specify multiple WWNs.

This specifier is optional.

<iscsi_name>

Host iSCSI name to be assigned or added to a host. This specifier is optional.

Authority

- Super, Edit
- Any role granted the `host_create` right

Usage

- Access to all domains is required to run this command.
- If the host is running the HPE Host Explorer then `createhost` may be called with no paths and the hostname reported by the explorer to create the host with all reported paths. If no agent is running then `createhost` with no paths creates a host without assigning a host path. The -add option modifies an existing host by associating a host path with that host. If the -add option is not used, the `createhost` command defines a new host.
- If assigning paths to a host being created, specify the -f option to remove any existing VLUNs associated with those paths to avoid any conflicting VLUNs.
- Verify the creation of a host by issuing the **showhost** command.
- The options that allow for adding descriptive information are for annotation purposes only; the information provided here is not actively used by the storage server.
- Although it is optional to specify a host persona, a host must have one.
- If one is not specified, it will default to 1 (Generic). This will not operate correctly with all types of hosts. The host persona can be changed with `sethost -persona` (see **sethost**).
- Host personas should not be specified with the -add argument.
- A port's FC parameters should be set with `controlport config` (see **controlport**).

Examples

The following example creates system host test01:

```
cli% createhost test01 2000000087041F72
```

The following example creates an iSCSI host test02:

```
cli% createhost -iscsi test02 iqn.1991-06.com.microsoft:dt-391-xp.hq.3par.com
```

The following example creates host ibm520-13 with an AIX host persona:

```
cli% createhost -persona 3 ibm520-13 10000000C97B142E
```

The following example requests that host lynx be zoned with port 1 in node 0, slot 8:

```
cli% createhost -port 0:8:1 lynx 1000A0B3CC1C68BE
```

createhostset

Syntax

```
createhostset [options] <setname> [<host | pattern>...]
```

Description

The `createhostset` command defines a new set of hosts and provides the option of assigning one or more existing hosts to that set. The command also allows the addition of hosts to an existing set by use of the -add option.

Options

-add

Specifies that the hosts listed should be added to an existing set. At least one host must be specified.

-comment <comment>

Specifies any comment or additional information for the set. The comment can be up to 255 characters long. Unprintable characters are not allowed.

-domain <domain>

Creates the host set in the specified domain. For an empty set the default is to create it in the current domain, or no domain if the current domain is not set. A host set must be in the same domain as its members; if hosts are specified as part of the creation then the set will be created in their domain. The -domain option should still be used to specify which domain to use for the set when the hosts are members of domain sets. A domain cannot be specified when adding a host to an existing set with the -add option.

Specifiers

<setname>

Specifies the name of the host set to create/add to, using up to 27 characters in length.

<host | pattern>...

Optional list of hosts or glob-style patterns matching hosts that are to be included in the host set. If no <host> is specified a host set with no hosts is created.

Authority

- Super, Edit
- Any role granted the `hostset_create` right can create a host set
- Any role granted the `hostset_set` right can add hosts to a host set

Usage

None.

Examples

To create an empty hostset:

```
cli% createhostset hostset
```

To add a host to the set:

```
cli% createhostset -add hostset hosta
```

To create a host set with hosts in it:

```
cli% createhostset hostset host1 host2
```

or

```
cli% createhostset set:hostset host1 host2
```

To create a host set with a comment and a host in it:

```
cli% createhostset -comment "A host set" oraset ora1
```

createmaint

Syntax

```
createmaint [options] <mainttype>
```

Description

The createmaint command creates a maintenance window record with the specified options and maintenance type.

Options

-f

Force the creation of the maintenance window record even if an active maintenance window record already exists with the same maintenance type. If this option is not specified, the command requires confirmation before proceeding with its operation to create the second maintenance window record with the same type.

-comment <comment>

Specifies any comment or additional information for the maintenance window record. The comment can be up to 255 characters long. Unprintable characters are not allowed.

-duration <time> {m|h}

Sets the duration of the maintenance window record. May be specified in minutes (e.g. 20m) or hours (e.g. 6h). Value is not to exceed 24 hours. The default is 4 hours.

Specifiers

<mainttype>

Specify the maintenance type.

Maintenance type can be `Other`, `Node`, `Restart`, `Disk`, `Cage`, `Cabline`, `Upgrade`, `DiskFirmware`, or `CageFirmware`.

Authority

- Super, Service
- Any role granted the `maint_set` rights

Usage

If an active maintenance window record exists with the same maintenance type, its reference count will be increased and the window record will be updated. Specifically, if the new end time is later than the existing one, the duration will be extended as appropriate and the comment will be updated. If the new end time is not later than the existing one, only the reference count will be increased.

Examples

To create a maintenance window record with 240 minutes as default duration:

```
cli% createmaint -comment "Other Maint" Other
```

To create a maintenance window record of 21 hours:

```
cli% createmaint -comment "Other Maint" -duration 21h Other
```

To create maintenance window record forcefully:

```
cli% createmaint -f Other
```

creatercopygroup

Syntax

```
creatercopygroup [options] <group_name> <target_name>:<mode>
[<target_name>:<mode>...]
```

Description

The `creatercopygroup` command creates a Remote Copy volume group.

Options

-domain <domain>

Creates the Remote Copy group in the specified domain. The volume group must be created by a member of a particular domain with Super or Edit privileges.

-usr_cpg <cpg name> <target_name>:<cpg_name>...

Specifies the local user CPG and target user CPG that will be used for volumes that are auto-created. The local CPG will only be used after failover and recovery.

-snp_cpg <cpg name> <target_name>:<cpg_name>...

Specifies the local snap CPG and target snap CPG that will be used for volumes that are auto-created. The local CPG will only be used after failover and recovery.

Specifiers

<group_name>

Specifies the name of the volume group, using up to 22 characters if the mirror_config policy is set, or up to 31 characters otherwise. This name is assigned with this command.

<target_name>

Specifies the target name associated with this group. This name should already have been assigned using the **creatercopytarget** command. The <target_name>:<mode> pair can be repeated to specify multiple targets.

<mode>

Specifies that the mode of the created group, the available modes are:

- sync—synchronous replication
- async—asynchronous streaming replication
- periodic—periodic asynchronous replication

The <target_name>:<mode> pair can be repeated to specify multiple targets.

Authority

- Super, Edit
- Any role granted the `rcopygroup_create` right

Usage

- This command requires the HPE 3PAR Remote Copy Software license. Contact your local service provider for more information.
- When a Remote Copy group is created, a VV set will also be created. Its name will be the Remote Copy group's name prepended with "RCP_", with numbers appended if necessary to make the set name unique. Volumes will automatically be added to and removed from this set when volumes are admitted to and dismissed from the group.

Examples

The following example creates an asynchronous periodic mode volume group named Group1 whose target system is System2 (target name System2_out):

```
creatercopygroup Group1 System2_out:periodic
```

If you are using domains, the syntax is as follows:

```
creatercopygroup -domain domain2 Group1 System2_out:periodic
```

If you are using volume auto-create:

```
createcopygroup  -usr_cpg Localcpg target:Targetcpg -snp_cpg Localsnp
  target:Targetsnp groupname target:periodic
```

createcopytarget

Syntax

Syntax for remote copy over IP (RCIP) is as follows:

```
createcopytarget [options] <target_name> IP [node:slot:port:IP_address...]
```

Syntax for remote copy over fibre channel (RCFC) is as follows:

```
createcopytarget [options] <target_name> FC <node_WWN>
[node:slot:port:WWN...]
```

Description

The `createcopytarget` command creates a remote copy target definition.

Options

`-disabled`

Creates the target in disabled mode

Specifiers

`<target_name>`

The name of the target definition to be created, specified by using up to 23 characters.

`<node_WWN>`

The node's World Wide Name (WWN) on the target system (Fibre channel target only).

`<node:slot:port:IP_address>...`

Specifies the node, slot, and port of the Ethernet port on the local system and an IP address of the peer port on the target system.

`<node:slot:port:WWN>...`

Specifies the node, slot, and port of the Fibre Channel port on the local system and World Wide Name (WWN) of the peer port on the target system.

Authority

- Super, Edit
- Any role granted the `rcopytarget_create` right

Usage

- Functionality of this command requires the HPE 3PAR Remote Copy Software license. Contact your local service provider for more information.
- Access to all domains is required to run this command.
- IP targets are made up of pairs composed of the node, slot and port of the ethernet port on the local system and an IP address of the peer port on the target system.
- FC targets are made up of sets with the node, slot, and port of the FC port on the local system and WWN of the peer port on the target system.

Examples

The following example creates a target definition named System2 on the primary system, where the Ethernet ports on the primary system are located at 0:3:1 and 1:3:1, and the corresponding IP addresses on the backup system are 10.1.1.11 and 10.1.2.11:

```
createcopytarget System2 IP 0:3:1:10.1.1.11 1:3:1:10.1.2.11
```

The following example creates a primary target definition named System1 on the backup system, where the Ethernet ports on the backup system are located at 2:3:1 and 3:3:1, and the corresponding IP addresses on the primary system are 10.1.1.96 and 10.1.2.96:

```
createcopytarget System1 IP 2:3:1:10.1.1.96 3:3:1:10.1.2.96
```

The following example creates a target definition named System2 on the local system where the node WWN of the target system is 20010002AC000060:

```
createcopytarget System2 FC 20010002AC000060 0:4:1:22410002AC000060
1:4:1:23410002AC000060
```

createsched

Syntax

```
createsched [options] <cmd> <taskschedule> <schedname>
```

Description

Specify details to run created tasks using the scheduler.

Options

-run_once

Runs the task one time only, at the specified time.

-no_alert

Dampens alert notifications. Tasks created with `no_alert` that fail do not generate alerts.

-f

Silences user interaction when creating a new scheduled task. Does not detect an event when a new scheduled task exceeds the recommended number of scheduled task starts per minute.

Specifiers

<cmd>

Specifies the CLI command or the HPE-distributed script. Maximum of 1023 characters.

<taskschedule>

Specifies the crontab-style schedule. Maximum of 127 characters.

<schedname>

Specifies the name given to the schedule. Maximum of 31 characters.

Authority

- Super, Service, Edit
- Any role granted the `sched_create` right

Usage

- The names `all` and `upgrade` are reserved.
- Schedule the following commands only:

  ```
  checkhealth compactcpg compactld createfsnap creategroupsv createsv
  creategroupvvcopy createvvcopy importvv moverelocpd removevv setqos
  setsralertcrit setvv startao startfsav startfsnapclean syncrcopy tunepd
  tunesys tunevv updatesnapspace updatevv
  ```

- The *<taskschedule>* string has five fixed fields:
  - `minute`
  - `hour`
  - `day-of-month`
  - `month`
  - `day-of-week`
- To avoid spikes in resource utilization, spread out scheduled task start times as much as possible. `Createsched` predicts out to 30 days when the newly scheduled task exceeds the recommended number of scheduled task starts per minute. When detected, the system requests user confirmation to continue creating the schedule.
- To skip this detection and confirmation, use the `-f` option.
- Separate fields using spaces. The allowed values for the fields are:

Field	Allowed Values
Minute	0-59
Hour	* or 0-23
Day-of-Month	* or 1-31
Month	* or 1-12
Day-of-Week	* or 0-6 (0 is Sunday)

- Scheduling on February 29 is not supported.
- Lists are allowed. A list is a set of numbers (or ranges) separated by commas. Examples: `"1,2,5,9","0-4,8-12"`.
- Ranges of numbers are allowed. Ranges are two numbers separated with a hyphen. The specified range is inclusive. For example, 8-11 for an hour entry specifies execution at hours 8, 9, 10 and 11. Ranges must be specified from low to high.
- One of six special strings may also appear instead of the specification above:

String	Meaning
@yearly	Run every January 1st, "0 0 1 1 *"
@monthly	Run the first day of every month, "0 0 1 * *"
@weekly	Run every Sunday, "0 0 * * 0"
@daily	Run every midnight, "0 0 * * *"
@hourly	Run every hour, on the hour, "0 * * * *"

- If some form of daylight savings or summer/winter time is in effect, then jobs scheduled during the switchover period could be executed once, twice, or not at all.

Examples

Creating a scheduled task that runs every hour on the hour ("0 * * * *") and creates a read only (`-ro`) snapshot of `sample_vv`, that expires in two hours (`-exp 2h`). Name the created snapshot `sample_vv_set.HH.MMM` where HH and MM refer to the hour and minute that the snapshot is created (`@vvname@.@H@.@M@`):

```
createsched "createsv -ro -exp 2h @vvname@.@H@.@M@ vv" "0 * * * *" sv_task
```

Creating a scheduled task that creates a consistent group snapshot.

The read only snapshots are created from the volume set `set:dbora` that expire in two hours (`-exp 2h`), name snapshots using the pattern `@vvname@.@s@`. Includes the volume name, specified in the set, and the epoch in seconds appended. The task name is `snap_odbora_set`, and the task runs every hour on the hour during the working hours from 8am to 5pm (`"0 8-17 * * *"`). For more information, enter the `clihelp -sub objectsets` command, or see, **Object sets and domain sets**.

```
createsched "createsv -ro -exp 2h @vvname@.@s@ set:dbora" "0 8-17 * * *"  sv_dbora_task
```

If just a single vv then:

```
createsched "createsv -ro -exp 2h @vvname@.@s@ vvname" "0 * * * *" snap_large_tpvv_a
```

Adding a comment with spaces to `createsv` by surrounding the comment with escaped double quotes:

```
createsched "createsv -ro -exp 2h -comment \"Comment with spaces\"
   @vvname@.@H@.@M@ vv" "0 * * * *" sv_task
```

createsnmpuser

Syntax

```
createsnmpuser [options <arg>] <username>
```

Description

The `createsnmpuser` command creates the SNMPv3 user. The user will be prompted to provide authentication and privacy passwords, which will be converted to secret keys.

Options

`-p <password>`

Specifies the user password. The length of <password> must be at least 8 characters.

-privPassword <privPassword>

Specifies the user's privacy password. The length of <privPassword> must be at least 8 characters.

Specifiers

<username>

Specifies the SNMPv3 user name.

Authority

- Super
- Any role granted the snmpuser_create right.

Usage

- Access to all domains is required to run this command.
- Both authentication and privacy passwords must be specified when SNMP user is created. If –p or –privPassword is not specified, the user will be prompted to enter those passwords.

Examples

The following example displays the setting of an SNMPv3 user:

```
cli% createsnmpuser -p password1 -privPassword password2 joe1
Please enter the password to generate authentication key for use in SNMPv3:
Password for user authentication key:
Retype password for user authentication key:
Please enter the password to generate privacy key for use in SNMPv3:
Password for user privacy key:
Retype password for user privacy key:
SNMP user successfully set.
```

createspare

Syntax

```
createspare [options <arg>] <chunklet_specifier>...
```

Description

The createspare command allocates chunklet resources as spares. Chunklets marked as spare are not used for logical disk creation and are reserved explicitly for spares, thereby guaranteeing a minimum amount of spare space.

Options

-f

Do not ask for confirmation. The default is to ask for confirmation.

-p

Specifies that partial completion of the command is acceptable.

Additionally, specified chunklets are marked as spare only if they are not spare already.

Specifiers

<chunklet_specifier>...

The chunklet specifier is one of the following arguments:

<PD_ID:chunklet_number>

Specifies the identification of the physical disk and the chunklet number on the disk. This specifier can be repeated.

<PD_ID:a>

Specifies the identification of the physical disk and all chunklets (a) on the disk. This specifier can be repeated.

a:<chunklet_num>

Specifies a chunklet number on all physical disks. This specifier can be repeated.

-pos <cage:mag:disk:chunklet_num>

Specifies the position of a specific chunklet identified by its position in a drive cage, drive magazine, physical disk, and chunklet number. This specifier can be repeated.

-pos <cage:mag:disk:a>

Specifies that all chunklets on a physical disk, identified by drive cage number, drive magazine number, and physical disk number, are marked spare. This specifier can be repeated.

Authority

- Super, Service
- Any role granted the `spare_create` right

Usage

- Access to all domains is required to run this command.
- To verify the creation of a spare chunklet, issue the **showspare** command.

Examples

The following example marks chunklet 1 as spare for physical disk 15:

```
cli% createspare 15:1
```

The following example specifies the position in a drive cage, drive magazine, physical disk, and chunklet number. -pos 1:0.2:3:121, where 1 is the drive cage, 0.2 is the drive magazine, 3 is the physical disk, and 121 is the chunklet numbers.

```
cli% createspare -pos 1:0.2:3:121
```

createsralertcrit

Syntax

`createsralertcrit <type> [options <arg>] <condition>[,<condition>...] <name>`

Description

Creates a criterion that System Reporter evaluates to determine whether to generate a performance alert.

Options

Many options depend upon the type specifier, as shown in the following list.

- Options common to **all** types are:

  -daily

  Evaluates criterion on a daily basis, at midnight.

-hourly

Evaluates criterion on an hourly basis.

-hires

Evaluates criterion on a high resolution (5 minute) basis.

-count *<number>*

The number of matching objects that must meet the criteria in order for the alert to be generated. Note that only one alert is generated in this case and not one alert per affected object.

-recur *<recurrences>/<samples>*

Generates an alert only if the other conditions of the criterion recur repeatedly. The *<recurrences>* value is an integer from 2 to 10. The *<samples>* value is an integer from 2 to 10 that represents the number of previous System Reporter samples in which to examine the recurrences. The *<samples>* value must be at least the requested quantity of recurrences. These samples refer to the selected resolution of the criterion: hires, hourly, or daily.

-btsecs *<relative_time>*

A negative number indicating the number of seconds before the data sample time used to evaluate conditions which compare against an average. Instead of a number representing seconds, you can specify `btsecs` using a suffix of m, h, or d to represent time in minutes (-30m), hours (-1.5h), or days (-7d). The relative time cannot be more than 10 samples ago: 50 minutes for hires, 10 hours for hourly, or 10 days for daily. If you do not specify this option, the average is computed for the most recent data sample only.

Do not combine `-btsecs` with the `-recur` option.

-critical

Highest severity alert.

-major

Requires urgent action.

-minor

No immediate action required.

-info

Default. Informational only.

-comment *<comment>*

Specifies comments or additional information for the criterion, up to 511 characters.

-defer *<duration>*

Specifies the deferral period for repeating and alert. You can either specify the deferral duration in seconds or with a suffix of m, h, or d to represent time in minutes (30m), hours (1.5h), or days (7d). Asingle alert criteria can generate multiple alerts if multiple objects exceed the defined threshold. A deferral period applies to each unique alert. Acknowledging an alert with `setalert ack <id>` ends the deferral period early for that alert.

- Options specific to **PORT** type are:

  -port_type *<type>*[,*<type>*...]

  Limit the ported data to the types specified. Allowed types are:

  - `disk`—Disk port
  - `host`—Host Fibre channel port

- iscsi—Host ISCSI port
- free—Unused port
- fs—File Persona port
- peer—Data Migration FC port
- rcip—Remote copy IP port
- rcfc—Remote copy FC port

-port *<npat>*:*<spat>*:*<ppat>*[,*<npat>*:*<spat>*:*<ppat>*...]

Includes ports with *<port_n>*:*<port_s>*:<port_p> that match any of the specified *<npat>*:*<spat>*:*<ppat>* patterns. Patters are a glob-style. If not specified, includes all ports.

-both | -ctl | -data

Specifies how to display control and data transfers (both, only control, or only data). With no option specified, displays only data transfers.

-groupby *<groupby>*[,<groupby>...]

Consider the summed value of each combination of <groupby> items against the criterion's condition(s). The supported items are: PORT_N PORT_S PORT_P PORT_TYPE GBITPS TRANS_TYPE ALL

If not specified, the default is to group by PORT_N and PORT_S and PORT_P.

Fields applicable to the **PORT** type are:

read_iops, write_iops, total_iops, read_kbps, write_kbps, total_kbps, read_svctms, write_svctms, total_svctms, read_ioszkb, write_ioszkb, total_ioszkb, qlen, sum_read_iops, sum_write_iops sum_total_iops, sum_read_kbps, sum_write_kbps, sum_total_kbps, read_bw_pct, write_bw_pct, total_bw_pct, avg_busy, read_time0..read_time31, write_time0..write_time31, total_time0..total_time31, read_size0..read_size15, write_size0..write_size15, total_size0..total_size15

- Options specific to the **VLUN** type are:

-host *<host_name|host_set|host_wwn|pattern>*[,*<host_name|host_set|host_wwn|pattern>*...]

Limit the data to hosts with names that match one or more of the specified names or glob-style patterns. Host set name must start with "set:" and can also include patterns. To specify the host by WWN, start with "wwn:". A WWN can also include glob-style patterns.

-vv *<VV_name>*|*<VV_set>*|*<pattern>*[,*<VV_name>*|*<VV_set>*|*<pattern>*...]

Limits the data to VVs with names that match one or more of the specified names or glob-style patterns. Prefix VV set name with set: and include patterns, if desired.

-l *<LUN|pattern>*[,*<LUN|pattern>*...]

Limits the data to LUNs that match one or more of the specified LUNs or glob-style patterns.

-port *<npat>*:*<spat>*:*<ppat>*[,*<npat>*:*<spat>*:*<ppat>*...

Ports with <port_n>:<port_s>:<port_p> that match any of the specified <npat>:<spat>:<ppat> patterns are included, where each of the patterns is a glob-style pattern. If not specified, all ports are included.

-vlun <host>:<vv>[:<lun>:<port_n>:<port_s>:<port_p>]
[,<host>:<vv>[:<lun>:<port_n>:<port_s>:<port_p>]]...

Limits the data to VLUNs matching the specified combination of host, VV, lun, and port. Each of these components in this option may be a glob-style pattern. The host and VV components may specify a corresponding object set by prefixing "set:" to the component. The host component may specify a WWN by prefixing the component with "wwn:". The lun and port components are optional, and if not present, data will be filtered to any matching combination of host and VV. This option cannot be combined with `-host`, `-vv`, `-l`, or `-port`.

-groupby <groupby>[,<groupby>...]

Considers the summed value of each combination of <groupby> items against the criterion's condition(s). The supported items are: `DOM_NAME VV_NAME HOST_NAME LUN HOST_WWN PORT_N PORT_S PORT_P VVSET_NAME HOSTSET_NAME VM_NAME VM_ID VM_HOST VVOLSC ALL`

If not specified, the default is to group by `HOST_NAME`, `VV_NAME`, and `LUN`.

Fields applicable to the **VLUN** type are:

```
read_iops, write_iops, total_iops, read_kbps, write_kbps, total_kbps,
read_svctms, write_svctms, total_svctms, read_ioszkb, write_ioszkb,
total_ioszkb, qlen, sum_read_iops, sum_write_iops sum_total_iops,
sum_read_kbps, sum_write_kbps, sum_total_kbps, avg_busy,
read_time0..read_time31, write_time0..write_time31,
total_time0..total_time31, read_size0..read_size15,
write_size0..write_size15, total_size0..total_size15
```

- Options specific to the **LD** type:

-node <node>[,<node>...]

Limits the data to that corresponding to one of the specified nodes.

-cpg <CPG_name|pattern>[,<CPG_name|pattern>...]

Limits the data to LDs in CPGs with names that match one or more of the specified names or glob-style patterns.

{<LD_name>|<pattern>}[,{<LD_name>|<pattern>}]

LDs matching any of the specified LD_name or glob-style pattern are included. If not specified, all LDs are included.

-groupby <groupby>[,<groupby>...]

Considers the summed value of each combination of `<groupby>` items against the criterion's condition(s). The supported items are: `DOM_NAME LDID LD_NAME CPG_NAME NODE ALL`

If not specified, the default is to group by `LDID`.

Fields applicable to the **LD** type are:

```
read_iops, write_iops, total_iops, read_kbps, write_kbps, total_kbps,
read_svctms, write_svctms, total_svctms, read_ioszkb, write_ioszkb,
total_ioszkb, qlen, sum_read_iops, sum_write_iops sum_total_iops,
sum_read_kbps, sum_write_kbps, sum_total_kbps, avg_busy,
read_time0..read_time31, write_time0..write_time31,
total_time0..total_time31, read_size0..read_size15,
write_size0..write_size15, total_size0..total_size15
```

- Options specific to the **PD** type:

-disk_type <type>[,<type>...]

Limits the data to disks of the types specified. Allowed types are:

- FC—Fast Class
- NL—NearLine
- SSD—Solid State Drive
- SCM—Storage Class Memory Drive

-rpm <speed>[,<speed>...]

Limits the data to disks of the specified RPM. Allowed speeds are 7, 10, 15, 100, 150, and 1500

{<PDID>|<pattern>},{<PDID>|<pattern>}

PDs with IDs matching any of the specified PDID or glob-style pattern are included. If not specified, all PDs are included.

-groupby <groupby>[,<groupby>...]

Considers the summed value of each combination of <groupby> items against the criterion's condition(s). The supported items are: PDID PORT_N PORT_S PORT_P DISK_TYPE SPEED ALL

If not specified, the group defaults to PDID.

Fields applicable to the **PD** type:

read_iops, write_iops, total_iops, read_kbps, write_kbps, total_kbps, read_svctms, write_svctms, total_svctms, read_ioszkb, write_ioszkb, total_ioszkb, qlen, sum_read_iops, sum_write_iops sum_total_iops, sum_read_kbps, sum_write_kbps, sum_total_kbps, avg_busy, read_time0..read_time31, write_time0..write_time31, total_time0..total_time31, read_size0..read_size15, write_size0..write_size15, total_size0..total_size15

- Option specific to the **CMP** type is:

  -node <node>[,<node>...]

  Limits the data to that corresponding to one of the specified nodes.

  Fields applicable to the **CMP** type are:

  read_hits, write_hits, reads, writes, totals, lock_blocks, read_hitpct, write_hitpct, delack_nl, delack_fc, delack_ssd, free_page, clean_page, write1_page, writen_page, writesched_page, writing_page, recov_page, dirty_nl, dirty_fc, dirty_ssd, maxdirty_nl, maxdirty_fc, maxdirty_ssd, write_delack_pct

  The following fields are deprecated:

  delack_fc10, delack_fc15, delack_ssd150, delack_ssd100, dirty_fc10, dirty_fc15, dirty_ssd150, dirty_ssd100, maxdirty_fc10, maxdirty_fc15, maxdirty_ssd150, maxdirty_ssd100 Each acts as an alias for the equivalent field without distinction between disk speeds. For example delack_ssd100 or delack_ssd150 report equal to delack_ssd.

- Options specific to the **CPU** type:

  -node <node>[,<node>...]

  Limits the data to that corresponding to one of the specified nodes.

  -groupby <groupby>[,<groupby>...]

  Considers the summed value of each combination of <groupby> items against the criterion's condition(s). The supported items are: NODE CPU ALL

  If not specified, the default is to group by NODE and CPU.

Fields applicable to the **CPU** type:

`user_pct, sys_pct, idle_pct, interrupts, context_switches`

- Options specific to the **LINK** type:

  `-node <node>[,<node>...]`

  Limits the data to that corresponding to one of the specified nodes.

  `-groupby <groupby>[,<groupby>...]`

  Considers the summed value of each combination of `<groupby>` items against the criterion's condition(s). The supported items are: `NODE QUEUE NODE_TO ASIC_FROM ASIC_TO ALL`

  If not specified, the default is to group by `NODE`.

  Fields applicable to the **LINK** type:

  `xcb_sent, kbps, xcb_size`

- Options specific to the **QOS** type:

  `-target {{vvset|domain}:{<name>|<pattern>}|sys:all_others}[,...]`

  Limits the data to the specified QoS target rule(s). Include a target type either `{vvset|domain}`, and a name or glob-style pattern.

  The sys:all_others rule can be selected to display statistics for all other host I/O not regulated by any "on" QoS rule.

  Multiple targets types can be specified as a comma-separated list.

  `-groupby <groupby>[,<groupby>...]`

  Considers the summed value of each combination of `<groupby>` items against the criterion's condition(s). The supported items are: `DOM_NAME QOS_ID TARGET_TYPE TARGET_NAME ALL`

  If not specified, the default is to group by `QOS_ID` and `TARGET_NAME`.

  Fields applicable to the **QOS** type:

  `read_iops, write_iops, total_iops, read_kbps, write_kbps, total_kbps, read_svctms, write_svctms, total_svctms, read_ioszkb, write_ioszkb, total_ioszkb, qlen, sum_read_iops, sum_write_iops sum_total_iops, sum_read_kbps, sum_write_kbps, sum_total_kbps, limit_iops, limit_kbps, read_wait, write_wait, total_wait, rej, wait_qlen`

- Options specific to the **RCOPY** type:

  `-target <TARGET_NAME|pattern>[,<TARGET_NAME|pattern>...]`

  Limits the data to TARGET_NAMEs that match one or more of the specified TARGET_NAMEs or glob-style patterns.

  `-groupby <groupby>[,<groupby>...]`

  Considers the summed value of each combination of `<groupby>` items against the criterion's condition(s). The supported items are: `TARGET_NAME LINK_ID LINK_ADDR PORT_N PORT_S PORT_P ALL`

  Fields applicable to the **RCOPY** type:

  `rc_kbs, rc_kbps, rc_hbrttms`

- Options specific to the **RCVV** type:

-vv <VV_name|VV_set|pattern>[,<VV_name|VV_set|pattern>...]

Limits the data to VVs with names that match one or more of the specified names or glob-style patterns. VV set name must be prefixed by "set:" and can also include patterns.

-groupby <groupby>[,<groupby>...]

Considers the summed value of each combination of `<groupby>` items against the criterion's condition(s). The supported items are: `VV_NAME DOM_NAME TARGET_NAME TARGET_MODE GROUP_NAME GROUP_ROLE PORT_TYPE PORT_N PORT_S PORT_P VVSET_NAME ALL`

If not specified, the default is to group by `VV_NAME`.

Fields applicable to the **RCVV** type:

```
lcl_read_iops, lcl_write_iops, lcl_total_iops, rmt_read_iops,
rmt_write_iops, rmt_total_iops, lcl_read_kbps, lcl_write_kbps,
lcl_total_kbps, rmt_read_kbps, rmt_write_kbps, rmt_total_kbps,
lcl_read_svctms, lcl_write_svctms, lcl_total_svctms, rmt_read_svctms,
rmt_write_svctms, rmt_total_svctms, lcl_read_ioszkb, lcl_write_ioszkb,
lcl_total_ioszkb, rmt_read_ioszkb, rmt_write_ioszkb, rmt_total_ioszkb,
lcl_total_qlen, lcl_busy_pct, rmt_total_qlen, rmt_busy_pct
```

- Options specific to the **CACHE** type:

  -node <node>[,<node>...]

  Limits the data to that corresponding to one of the specified nodes.

  Fields applicable to the **CACHE** type:

  ```
  reads_per_sec, writes_per_sec, cmp_read_hit_pct, cmp_write_hit_pct,
  fmp_read_hit_pct, fmp_write_hit_pct, read_back_iops, read_back_mbps,
  destage_write_iops, destage_write_mbps, fmp_used_pct, fmp_q_dormant,
  fmp_q_cold, fmp_q_norm, fmp_q_warm, fmp_q_hot, fmp_q_destage, fmp_q_read,
  fmp_q_flush, fmp_q_wrtback, cmp_q_free, cmp_q_clean, cmp_q_write1,
  cmp_q_writen, cmp_q_wrtsched, cmp_q_writing, cmp_q_dcowpend, cmp_q_dcowproc
  ```

- Options specific to the **LDSPACE** type:

  -disk_type <type>[,<type>...]

  Limits the data to disks of the types specified. Allowed types are:

  - `FC`—Fast Class
  - `NL`—NearLine
  - `SSD`—Solid State Drive
  - `SCM`—Storage Class Memory Drive

  -owner <node>[,<node>...]

  Limits data to LDs owned by the specified nodes.

  -raid_type <type>[,<type>...]

  Limits the data to RAID of the specified types. Allowed types are 0, 1, 5, and 6.

  -cpg <CPG_name|pattern>[,<CPG_name|pattern>...]

  Limits the data to LDs in CPGs with names that match one or more of the specified names or glob-style patterns.

  -ld <LD_name>|<pattern>[,<LD_name>|<pattern>]

  LDs matching any of the specified LD_name or glob-style pattern are included. If not specified, all LDs are included.

-groupby <groupby>[,<groupby>...]

Considers the summed value of each combination of <groupby> items against the criterion's condition(s). The supported items are: DOM_NAME CPG_NAME LDID LD_NAME DISK_TYPE RAID_TYPE SET_SIZE STEP_SIZE ROW_SIZE OWNER ALL

If not specified, the default is to group by LDID.

Fields applicable to the **LDSPACE** type:

raw_mb, used_mb, free_mb, total_mb

- Options specific to the **PDSPACE** type:

-disk_type <type>[,<type>...]

Limits the data to disks of the types specified. Allowed types are:

- FC—Fast Class
- NL—NearLine
- SSD—Solid State Drive
- SCM—Storage Class Memory Drive

-rpm <speed>[,<speed>...]

Limits the data to disks of the specified RPM. Allowed speeds are 7, 10, 15, 100, 150, and 1500.

-pd <PDID>|<pattern>[,<PDID>|<pattern>]

PDs with IDs matching any of the specified PDID or glob-style pattern are included. If not specified, all PDs are included.

-groupby <groupby>[,<groupby>...]

Considers the summed value of each combination of <groupby> items against the criterion's condition(s). The supported items are: PDID CAGEID CAGESIDE MAG DISK DISK_TYPE SPEED ALL

If not specified, the default is to group by PDID.

Fields applicable to the **PDSPACE** type:

nrm_used_ok, nrm_used_fail, nrm_avail_clean, nrm_avail_dirty, nrm_avail_fail, spr_used_ok, spr_used_fail, spr_avail_clean, spr_avail_dirty, spr_avail_fail, lifeleft_pct, t_degc

- Options specific to the **CPGSPACE** type:

-disk_type <type>[,<type>...]

Limits the data to disks of the types specified. Allowed types are:

- FC—Fast Class
- NL—NearLine
- SSD—Solid State Drive
- SCM—Storage Class Memory Drive

-raid_type <type>[,<type>...]

Limits the data to RAID of the specified types. Allowed types are 0, 1, 5, and 6.

-cpg <CPG_name|pattern>[,<CPG_name|pattern>...]

Limits the data to LDs in CPGs with names that match one or more of the specified names or glob-style patterns.

-groupby <groupby>[,<groupby>...]

Considers the summed value of each combination of <groupby> items against the criterion's condition(s). The supported items are: DOM_NAME CPGID CPG_NAME DISK_TYPE RAID_TYPE ALL

If not specified, the default is to group by CPG_NAME.

Fields applicable to the **CPGSPACE** type:

adm_used_mb, snp_used_mb, usr_used_mb, total_used_mb, adm_free_mb, snp_free_mb, usr_free_mb, total_free_mb, adm_mb, snp_mb, usr_mb, total_mb, growth_avail_mb, private_base_mb, private_snap_mb, free_mb, shared_mb, total_space_mb

- Options specific to the **VVSPACE** type:

-usr_cpg <CPG_name|pattern>[,<CPG_name|pattern>...]

Include VVs with usr space mapped to a CPG with a name matching one of the specified CPG_name or glob-style patterns.

-snp_cpg <CPG_name|pattern>[,<CPG_name|pattern>...]

Include VVs with snp space mapped to a CPG with a name matching one of the specified CPG_name or glob-style patterns.

-prov <prov_type>[,<prov_type>...]

Only include VVs of the specified provisioning type(s).

The possible values are: cpvv dds full peer snp tdvv tpsd tpvv

-vv <VV_name|VV_set|pattern>[,<VV_name|VV_set|pattern>...]

Limits the data to VVs with names that match one or more of the specified names or glob-style patterns. VV set name must be prefixed by "set:" and can also include patterns.

-compr <Yes|No|Off|NA>

Only include VVs of the specified compression setting.

-groupby <groupby>[,<groupby>...]

Considers the summed value of each combination of <groupby> items against the criterion's condition(s). The supported items are: DOM_NAME VVID VV_NAME BSID WWN SNP_CPG_NAME USR_CPG_NAME PROV_TYPE VV_TYPE VVSET_NAME VM_NAME VM_ID VM_HOST VVOLSC VVOL_STATE COMPR ALL

If not specified, the default is to group by VV_NAME.

Fields applicable to the **VVSPACE** type:

userrawrsvd_mb, snaprawrsvd_mb, adminrawrsvd_mb, totalrawrsvd_mb, userused_mb, userfree_mb, userrsvd_mb, snapused_mb, snapfree_mb, snaprsvd_mb, adminused_mb, adminfree_mb, adminrsvd_mb, totalused_mb, totalrsvd_mb, virtualsize_mb

- Options specific to the **SYSSPACE** type:

-disk_type <type>[,<type>...]

Limits the data to disks of the types specified. Allowed types are:

- FC—Fast Class
- NL—NearLine
- SSD—Solid State Drive
- SCM—Storage Class Memory Drive

-groupby <groupby>[,<groupby>...]

Considers the summed value of each combination of <groupby> items against the criterion's condition(s). The supported items are: DOM_NAME DISK_TYPE ALL

If not specified, the default is to group by ALL, considering the system as a whole.

Fields applicable to the **SYSSPACE** type:

total_mb, alloc_mb, used_mb, system_mb, free_mb, failed_mb, compact_ratio, dedup_ratio, overprov_ratio

Specifiers

<type>

Specifies one of the following:

port
vlun
pd
ld
cmp
cpu
link
qos
rcopy
rcvv
ldspace
pdspace
cpgspace
vvspace
sysspace

<condition>

Specifies condition using the format *<field><comparison><value>*.

The *<field>* value corresponds to the type (see above). The *<comparison>* value uses the format <, <=, >, >=, =, !=. The and *<value>* is a numeric value, or a numeric value followed by %_average, to indicate that the field is compared against the average across multiple objects. Specify multiple objects by filtering options or across multiple data sample times as specified by the –btsecs option. See examples.

Some characters, such as and >, are significant in most shells and must be escaped or quoted when running this command from another shell. Separate multiple conditions using a comma (,) to indicate a logical AND requirement (conjunction). Use a tilde (~) to indicate a logical OR requirement (disjunction). AND logic takes precedence over OR logic, and parentheses are not supported to override the natural precedence of the condition terms and logical operators.

<name>

Specifies the name of the SR alert criterion, with a maximum of 31 characters.

Authority

- Super, Service
- Any role granted the `sralertcrit_create` right

Usage

- System Reporter periodically samples the performance of the system. Each time System Reporter takes a sample, it evaluates all of the criteria corresponding to that time interval and generates corresponding alerts if all of the specified conditions are met. If previous samples satisfied the criteria but the current sample does not, the alert is automatically fixed.
- The sum referenced in conditions containing a sum field (sum_read_iops, sum_write_iops, sum_total_iops, sum_read_kbps, sum_write_kbps, sum_total_kbps) is the total for the particular field for everything matching the specified options. In the case of no options, this would be the total for all the given objects.
- Any condition data field, except for those with sum_ prefix, may have the delta_ prefix prepended. The difference in the value of the field will be calculated from the previous sample. The delta can be negative.
- Fields utilizing the sum_ or delta_ prefixes may not be compared to a group average (specified by condition value %_average).
- Conditions that contain histogram fields (read_time0..read_time31, write_time0..write_time31, total_time0..total_time31, read_size0.. read_size15, write_size0..write_size15, total_size0..total_size15) refer to all entries for the given histogram or greater. For example, referring to write_time10 means that the condition is evaluating the sum of the I/Os that have a write_time greater than 8 milliseconds (which is what column 10 refers to).
- Including `-groupby ALL` for a criterion instructs that the condition(s) be compared against a summation of all of the relevant objects.

Examples

Generating an alert for each port with more than 50 write IOPS in a high resolution sample:

```
cli% createsralertcrit port write_iops>50 write_port_check
```

Generating an alert for each port with more than 10000 write IOPS in a high resolution sample if the sum of all write IOPS for all ports exceeds 50000:

```
cli% createsralertcrit port write_iops>10000,sum_write_iops>50000 write_port_check
```

Generating an alert if an hourly sample discovers that the sum of all the write IOPS performed on node 1 exceed 10000:

```
cli% createsralertcrit port -hourly -port 1:*:* sum_write_iops>10000 port_1_writes
```

Generating an alert if a high resolution sample discovers that the average of all host ports service time exceeds 2000 milliseconds:

```
cli% createsralertcrit port -port_type host -groupby ALL total_svctms>2000 host_ports_avg_2s
```

Generating an alert for each LD that has more than 1000 reads with over an 8 ms latency (corresponding to column 10 in the read histogram) discovered during an hourly sample. The 6 hour deferral period will prevent the alert from triggering again until 6 hours has passed:

```
cli% createsralertcrit ld -hourly -defer 6h read_time10>1000 ld_read
```

Generating an alert if node 0 has mode than 100 interrupts and greater than 70% CPU utilization during a high resolution sample:

```
cli% createsralertcrit cpu -node 0 interrupts>100,sys_pct>70 node0busy
```

Generating an alert if a high resolution sample discovers that the sum of all IOPS of all ports in the system exceeds 100000:

```
cli% createsralertcrit port -daily sum_total_iops>100000 sys_busy
```

Generating an alert if a VLUN experiences a sharp increase in total IOPS from the previous hour:

```
cli% createsralertcrit vlun -hourly delta_total_iops>800 vlun_spike
```

Generating an alert if 8 or more PDs experience at least 50ms average service time repeatedly, happening four times in any 6 high resolution periods (5 minutes each):

```
cli% createsralertcrit pd -hires -count 8 -recur 4/6 total_svctms>50 recur8slow_pd
```

Generating an alert for any NL PD where write service time is more than 150% of the average write service times of all the NL PDs:

```
cli% createsralertcrit pd -disk_type NL write_svctms>150%_average slow_nl
```

Generating an alert for any RC VV where remote write service time for the past hour is more than 150% of the average remote write service time of all the RC VVs over the past 5 hours:

```
cli% createsralertcrit rcvv -hourly -btsecs -5h rmt_write_svctms>150%_average slow_rcvv
```

Generating an alert if the average of all CPU cores in a node falls below 5% idle time:

```
cli% createsralertcrit cpu -groupby NODE idle_pct<5 node_CPUs_busy
```

createsv

Syntax

`createsv [options] <SV_name> <copy_of_VV | VV_set | RC_group>`

Description

Creates a point-in-time (snapshot) copy of a VV. Create snapshots to perform tasks such as backing up data on the base volume and allowing users to modify multiple copies of a base volume without affecting the original base volume.

Options

`-ro`

Specifies that the copied volume is read-only. If not specified, the volume is read/write.

`-i <VV_ID>`

Specifies the ID of the copied VV set. This option cannot be used when VV set is specified.

`-rcopy`

Specifies that synchronous snapshots be taken of a volume in a remote copy group. The same snapshot `<SV_name>` will be applied on both the primary and secondary arrays.

`-exp <time>{d|D|h|H|m|M}`

Specifies the relative time from the current time that volume will expire. `<time>` is a positive integer value and in the range of 1 minute - 1825 days. Specify time using `d` or `D` for days, `h` or `H` for hours, and `m` or `M` for minutes following the entered time value.

`-comment <comment>`

Specifies any additional information up to 511 characters for the volume.

When used with the `-rcopy` option for Asynchronous Streaming Remote Copy, the comment is applied only on the primary array.

-f

Silences system prompts to confirm volume creation with a retention time (-retain).

-retain *<time>*{d|D|h|H|m|M}

Specifies the amount of time, relative to the current time, that the system retains the volume. *<time>* is a positive integer value in the range of 1 minute - 1825 days. Use d or D for days, h or H for hours, and m or M for minutes following the entered time value.

If the volume is not in any domain, then its retention time cannot exceed the value of the VVRetentionTimeMax on the system. The default value for the VVRetentionTimeMax is 14 days.

If the volume belongs to a domain, then its retention time cannot exceed the value of the VVRetentionTimeMax of the domain, if set. ou cannot remove or reduce the retention time once it is set. If the volume has its retention time set, you cannot remove the volume within its retention time. If you have specified both expiration time and retention time, the retention time cannot be longer than the expiration time.

This option requires the HPE 3PAR Virtual Lock license. Contact your local service provider for more information.

-addtoset *<VV_set>*

Specifies the snapshots to add to the specified set. The set is created if it does not exist.

Specifiers

<SV_name>

Specifies the snapshot name, using up to 31 characters in length.

<copy_of_VV | VV_set>

Specifies the parent volume name or volume set name, using up to 31 characters. The VV_set name must start with set:. If copy_of VV is a VV set then svname can contain various patterns that are used to form the snapshot volume name. (For more information, issue clihelp sub,vvnamepat.) For remote copy groups, the group name should begin with rcgroup:. You can use @vvname@ replacements for this setting. If you do not provide a pattern, the <SV_name> field is used as a prefix, to which the VV_name is appended. This should be used in conjunction with the -rcopy flag.

Authority

- Super, Edit
- Any role granted the sv_create right

Usage

- If the volume is not in any domain, then its retention time cannot exceed the value of the system's VVRetentionTimeMax. The default value for the system's VVRetentionTimeMax is 14 days. If the volume belongs to a domain, then its retention time cannot exceed the value of the domain's VVRetentionTimeMax, if set. The retention time cannot be removed or reduced once it is set. If the volume has its retention time set, it cannot be removed within its retention time. If both expiration time and retention time are specified, then the retention time cannot be longer than the expiration time. This option requires the HPE 3PAR Virtual Lock License. Contact your local HPE 3PAR Authorized Service Provider for more information.
- A maximum of 2048 virtual copies per volume are allowed. RW snapshots made directly against a base volume or another RW snapshot will reduce the available amount by one for each. RO snapshots made directly against another RO snapshot will also reduce the available amount by one for each.
- For a vv set, the createsv command creates consistent group snapshots of the virtual volumes in the set. Consistent group snapshots are all created at the same point in time so that related structures on different volumes in the set remain consistent.

- For remote copy (`-rcopy`), identical snapshots are created on both the primary and secondary arrays. Only read-only snapshots may be taken with the `-rcopy` flag.
- If the remote copy volume pair has different names on the primary and secondary arrays, any automatic @vvname@ substitution will reflect the name of the primary array volume.

Examples

- Creating a read-only snapshot volume svro_vv0 from volume vv0:

```
cli% createsv -ro svr0_vv0 vv0
```

- Creating snapshot volume svrw_vv0 from the snapshot svro_vv0:

```
cli% createsv svrw_vv0 svro_vv0
```

- Creating a snapshot volume for each member of the VV set vvcopies. Each snapshot will be named svro-<name of parent volume>:

```
cli% createsv -ro svro-@vvname@ set:vvcopies
```

- Creating identical read-only snapshot volumes on both the primary and secondary node:

```
cli% createsv -rcopy -ro sv_vv0 vv0
```

  If vvcopies contained VVs named vv.0, vv.1 and vv.2 this would result in snapshots named svro-vv.0, svro-vv.1 and svro-vv.2.

- Creates the snapshots sv-vv.0, sv-vv.1, and sv-vv.2 and adds them to the VV set vvsnaps:

```
cli% createsv -addtoset vvsnaps sv-@vvname@ set:vvcopies
```

- Creating coordinated snapshots across a full remote copy group:

```
cli% createsv -rcopy -ro sv_@vvname@_1 rcgroup:group1
```

  If there are vv.1, vv.2 ...,vv.n within group1, this results in sv_vv.1_1, sv_vv.2_1, ..., sv_vv.n_1. If however, there is no pattern included, the `<SV_name>` is used as a prefix, appended with a '.':

```
cli% createsv -rcopy -ro sv rcgroup:group1
```

  Assuming the same VVs as before, creates sv.vv.1, sv.vv.2, sv.vv.n.

createtemplate

Syntax

```
createtemplate vv|ld|cpg [options <arg>] <template_name>
```

Description

The `createtemplate` command creates virtual volume (VV), logical disk (LD), and common provisioning group (CPG) templates. Templates are sets of predetermined parameters that are applied to the creation of virtual volumes, logical disks, and CPGs.

Use of this command with the vv specifier to create a VV is deprecated and will be removed in a subsequent release. The **`createvv`**command must be used instead to create a VV.

Options

The following options are valid for VV, LD, and CPG templates:

-nro

Sets unspecified properties in the template to read-only.

-nrw

Sets unspecified properties in the template to read-write.

-desc <description>

Specifies a description for the created template, up to 255 characters.

-rw

Specifies that the list of property values in the remainder of the command line, or until the next `-rw` or `-ro` option, are read-write (values can be modified when the template is used).

-ro

Specifies that the list of property values in the remainder of the command line, or until the next `-rw` or `-ro` option, are read-only (values cannot be modified when the template is used).

With neither `-ro` nor `-rw` is specified, defaults to `-ro`.

-rs <size>

Specifies the number of sets in a row. The <size> is a positive integer.

If not specified, no row limit is imposed.

-ss <size_KB>

Specifies the step size from 32 KB to 512 KB. The step size should be a power of 2 and a multiple of 32. The default value depends on raid type and device type used. If no value is entered and FC or NL drives are used, the step size defaults to 256 KB for RAID-0 and RAID-1, and 128 KB for RAID-5. If SSD drives are used, the step size defaults to 32 KB for RAID-0 and RAID-1, and 64 KB for RAID-5. For RAID-6, the default is a function of the set size.

-ha port|cage|mag

Specifies that the layout must support the failure of one port pair, one cage, or one drive magazine (mag). This option has no meaning for RAID-0. The default is cage availability.

-ch first|last

Specifies the chunklet location characteristics: either first (attempt to use the lowest numbered available chunklets) or last(attempt to use the highest numbered available chunklets). If no argument is specified, the default characteristic is first.

-p <pattern>

Specifies a pattern for candidate disks. Patterns are used to select disks that are used for creating logical disks. If no pattern is specified, the option defaults to Fast Class (FC) disks. If specified multiple times, each instance of the specified pattern adds additional candidate disks that match the pattern. The -devtype pattern cannot be used to mix Nearline (NL), FC, and Solid State Drive (SSD) drives. An item is specified as an integer, a comma-separated list of integers, or a range of integers specified from low to high.

The following arguments can be specified as patterns for this option:

An item is specified as an integer, a comma-separated list of integers, or a range of integers specified from low to high.

-nd <item>

Specifies one or more nodes. Nodes are identified by one or more integers (item). Multiple nodes are separated with a single comma (for example, 1,2,3). A range of nodes is separated with a hyphen (for example, 0- 7). The primary path of the disks must be on the specified node(s).

-st <item>

Specifies one or more PCI slots. Slots are identified by one or more integers (item). Multiple slots are separated with a single comma (for example, 1,2,3). A range of slots is separated with a hyphen (for example, 0- 7). The primary path of the disks must be on the specified PCI slot(s).

-pt <item>

Specifies one or more ports. Ports are identified by one or more integers (item). Multiple ports are separated with a single comma (for example, 1,2,3). A range of ports is separated with a hyphen (for example, 0- 4). The primary path of the disks must be on the specified port(s).

-cg <item>

Specifies one or more drive cages. Drive cages are identified by one or more integers (item). Multiple drive cages are separated with a single comma (for example, 1,2,3). A range of drive cages is separated with a hyphen (for example, 0-3). The specified drive cage(s) must contain disks.

-mg <item>

Specifies one or more drive magazines. The "1." or "0." displayed in the CagePos column of **showpd** output indicating the side of the cage is omitted when using the -mg option. Drive magazines are identified by one or more integers (item). Multiple drive magazines are separated with a single comma (e.g. 1,2,3). A range of drive magazines is separated with a hyphen (e.g. 0-7). The specified drive magazine(s) must contain disks.

-pn <item>

Specifies one or more disk positions within a drive magazine. Disk positions are identified by one or more integers (item). Multiple disk positions are separated with a single comma (e.g. 1,2,3). A range of disk positions is separated with a hyphen(e.g. 0-3). The specified position(s) must contain disks.

-dk <item>

Specifies one or more physical disks. Disks are identified by one or more integers(item). Multiple disks are separated with a single comma (e.g. 1,2,3). A range of disks is separated with a hyphen (e.g. 0-3). Disks must match the specified ID(s).

-tc_gt <number>

Specifies that physical disks with total chunklets greater than the number specified be selected.

-tc_lt <number>

Specifies that physical disks with total chunklets less than the number specified be selected.

-fc_gt <number>

Specifies that physical disks with free chunklets greater than the number specified be selected.

-fc_lt <number>

Specifies that physical disks with free chunklets less than the number specified be selected.

-devid <model>

Specifies that physical disks identified by their models be selected. Models can be specified in a comma-separated list. Models can be displayed by issuing the "`showpd -i`" command.

-devtype <type>

Specifies that physical disks must have the specified device type (FC for Fast Class, NL for Nearline, SSD for Solid State Drive, SCM for Storage Class Memory) to be used. Device types can be displayed by issuing the `showpd` command. If the `devtype` is not specified, the default device type is FC.

-rpm <number>

Disks must be of the specified speed. Device speeds are shown in the RPM column of the `showpd` command. For FC and NL drives, the number corresponds to both a performance measure and actual rotational speed. For SSD & SCM drives, this field is "N/A".

Disks that satisfy all of the specified characteristics are used. For example -p -fc_gt 60 -fc_lt 230 -nd 2 specifies all the disks that have greater than 60 and less than 230 free chunklets and that are connected to node 2 through their primary path.

The following options are valid only for VV templates:

-type {cpvv|tdvv|tpvv|none}

Specifies the type of virtual volume for the template being created. Valid arguments are cpvv, tdvv, tpvv, or none.

Use cpvv if the-cpg <CPG_name> option is used. If creating a template for creating TDVVs, use tdvv. If creating a template for creating TPVVs, use tpvv.

If cpvv, tpvv, or tpvv are not specified, the default value is none meaning the created template is applicable for the creation of any volume type. The "-type cpvv" option is deprecated and will be removed in a subsequent release.

-size <size>[g|G|t|T]

Specifies the size of the virtual volume and logical disk. Size can be specified in megabytes (default), in gigabytes by using the g|G parameter and terabytes by using the t|T parameter.

-spt <sectors_per_track>

Defines the virtual volume geometry sectors per track value that is reported to the hosts through the SCSI mode pages. The valid range is between 4 to 8192 and the default value is 304.

-hpc <heads_per_cylinder>

Allows you to define the virtual volume geometry heads per cylinder value that is reported to the hosts though the SCSI mode pages. The valid range is between 1 to 255 and the default value is 8.

-pol <pol>[,<pol>...]

Specifies the following policies that the created virtual volume follows.

If an argument is not specified, the option defaults to stale_ss.

stale_ss

Specifies that invalid snapshot volumes are permitted. Failure to update snapshot data does not affect the write to the base volume, but the snapshot is considered invalid.

no_stale_ss

Specifies that invalid snapshot volumes are not permitted. Failure to update a snapshot is considered a failure to write to the base volume.

one_host

This constrains the export of a volume to one host or one host cluster (when cluster names may be used as a host name).

no_one_host

This policy should only be used when exporting a virtual volume to multiple hosts for use by a cluster-aware application, or when "port presents" VLUNs are used. This is the default policy setting.

3par_host_dif

This policy enables support for 3PAR implementation of host based DIF.

std_host_dif

This policy enables support for standard SCSI implementation of host based DIF.

no_host_dif

This policy disables support for host DIF.

-usr_cpg <usr_cpg>

Specifies the name of the CPG from which the user space will be allocated.

-usr_aw <percent>

This option enables user space allocation warning. Generates a warning alert when the user data space of the TPVV exceeds the specified percentage of the virtual volume size.

-usr_al <percent>

Indicates the user space allocation limit. The user space of the TPVV is prevented from growing beyond the indicated percentage of the virtual volume size. After this limit is reached, any new writes to the virtual volume will fail.

-snp_cpg <snp_cpg>

The name of the CPG from which the snapshot space is allocated.

-snp_aw <percent>

Enables a snapshot space allocation warning. A warning alert is generated when the reserved snapshot space of the VV exceeds the indicated percentage of the VV size.

-snp_al <percent>

Sets a snapshot space allocation limit. The snapshot space of the VV is prevented from growing beyond the indicated percentage of the virtual volume size.

The following options are only used for CPG templates:

-aw <percent>

Specifies the percentage of used snapshot administration or snapshot data space that results in a warning alert. A percent value of 0 disables the warning alert generation. The default is 0.

This option is deprecated and will be removed in a subsequent release.

-sdgs <size>[g|G|t|T]

Specifies the growth increment, the amount of logical disk storage created on each auto-grow operation. The default growth increment may vary according to the number of controller nodes in the system. If <size> is non-zero, it must be 8G or bigger. The size can be specified in MB (default) or GB (using g or G) or TB (using t or T). A size of 0 disables the auto-grow feature. The following table displays the default and minimum growth increments per number of nodes:

Number of Nodes	Default	Minimum
1–2	32G	8G
3–4	64G	16G
5–6	96G	24G
7–8	128G	32G

-sdgl <size>[g|G|t|T]

Specifies that the auto-grow operation is limited to the specified storage amount. The storage amount can be specified in MB (default) or GB (using g or G) or TB (using t or T). A size of 0 (default) means that no limit is enforced. To disable auto-grow, set the limit to 1.

-sdgw <size>[g|G|t|T]

Specifies that the threshold of used logical disk space, when exceeded, results in a warning alert. The size can be specified in MB (default) or GB (using g or G) or TB (using t or T). A size of 0 (default) means that no warning limit is enforced. To set the warning for any used space, set the limit to 1.

Specifiers

vv|ld|cpg

Specifies that the template is for the creation of a virtual volume (VV), logical disk (LD), or common provisioning group (CPG).

<template_name>

Specifies the name of the template, up to 31 characters in length.

Authority

- Super
- Any role granted the `template_create` right

Usage

- Requires access to all domains.
- Defaults to logical disk creation using only physical disks with the same device type. Device type defaults to Fast Class. Use the `-p -devtype NL`, `-p -devtype SSD` to override the device type default. Use `showpd` to see the device types of physical disks in the system.
- You can set the value for `-ssz` and `-rs` to `-`. Use this in conjunction with the (default) read-only property for options to ensure that the `-ssz` and `-rs` values cannot be changed when a virtual volume, logical disk, or CPG is created and the system is used to set the default values for these options.
- The options `-sdgs`, `-sdgl`, and `-sdgw` control the auto logical disk creation for the Common Provisioning Group snapshot data regions. Auto logical disk creation occurs when the amount of free logical disk space falls below the specified grow (enlarge) size setting options (`-sdgs`, `-sdgl`, and `-sdgw`).
- For this command KB = 1024 bytes, MB = 1024 KB, GB = 1024 MB.

Examples

None.

createuser

Syntax

```
createuser [option] <username> <domainname> <role>
```

Description

The `createuser` command allows a user to create a new user with the specified name and role in the specified domain.

Options

-c <clear-text-password>

Specifies the user's password in clear-text format. The length of <clear-text-password> must be between the admin defined minimum length and 32 characters. If no minimum has been defined, then the default minimum value is 6.

Specifiers

<username>

Name of the user. The username can be up to 31 characters long. Valid characters are alphanumeric (letters and digits), period ("."), dash ("-") and an underscore ("_"). The first character must be alphanumeric or an underscore for non-SSH users. To access the system via SSH, the first character of the username must be alphanumeric.

<domainname>

Name of the domain. The domain name can be up to 31 characters long.

"all" can be used if the user is not to be in a specific domain.

<role>

The user role. Valid values for role are "browse", "edit", "service", "super", "audit" and any extended role.

Super, audit, and service roles cannot be given in individual domains.

Use showrole to list extended roles.

Authority

- Super
- Any role granted the user_create right

Usage

- Access to all domains is required to run this command.
- If no password is specified using -c then the user will be prompted for a clear text password.
- Verify the creation of a new user by issuing the **showuser** command.

Examples

The following example displays the successful creation of a new user user1 with the clear-text password 123456, with access to domain testdomain and with edit role authority:

```
cli% createuser -c 123456 user1 testdomain edit
User created
```

createvfs

Syntax

```
createvfs [options] <ipaddr> <subnet> <vfsname>
```

Description

createvfs creates a virtual file system (VFS). It can optionally create the file provisioning group (FPG) to which the VFS will belong.

Options

Only one of the following certificate options can be specified:

-nocert

Do not create a self-signed certificate associated with the VFS.

-certfile <certfile>

Uses the certificate data contained in this file.

-certdata <certificate string>

Uses the certificate data contained in this string.

-comment

Specifies any additional textual information.

-bgrace <time>

The block grace time in seconds for quotas within the VFS.

-igrace <time>

The inode grace time in seconds for quotas within the VFS.

-snapquota {enable|disable}

Enables or disables the quota accounting flag for Snapshots at VFS level.

-fpg <fpgname>

The name of an existing FPG in which the VFS must be created.

-cpg <cpgname>

The CPG in which the FPG must be created.

-size <size>

The size of the FPG to be created.

-tdvv

Creates the FPG with TDVV volumes.

-full

Creates the FPG with fully provisioned volumes.

-node <nodeid>

The node to which the FPG must be assigned.

Can only be used when creating the FPG with the -cpg option.

-vlan <vlanid>

The VLAN ID associated with the VFSIP.

-wait

Waits until the associated task is completed before proceeding.

This option will produce verbose task information.

-subnetoverlap

Allow overlapping of subnets.

NOTE:

This is not a typical configuration. If used improperly, the configuration can cause data unavailability. If not specified, overlapping subnets will not be allowed.

-f

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

Specifiers

<ipaddr>

The IP address to which the VFS must be assigned.

<subnet>

The subnet for the IP Address.

<vfsname>

The name of the VFS to be created.

Authority

- Super, Edit
- Any role granted `vfs_create` right

Usage

- Access to all domains is required to run this command.
- If an FPG is created, it will be given the same name as the VFS.
- Both names must be available for creation for the command to succeed.
- Either -fpg or the parameters to create an FPG must be specified to create a VFS.
- This command will spawn a task and return the taskid.
- Certificates must be in PEM format, containing both public and private keys.
- When specifying grace times, both block and inode grace times are required.
- Grace times are specified in minutes.
- When specifying -node, the FPG must be created alongside the creation of the VFS.
- The VFS name cannot exceed 212 characters.

Examples

Creating a VFS, examplevfs on examplefpg and generating a self-signed certificate:

```
cli% createvfs -fpg examplefpg 10.10.10.10 255.255.0.0 examplevfs
```

Creating a VFS, Eng, and in doing so creating the associated FPG:

```
cli% createvfs -cpg examplecpg -size 8T 10.10.10.11 255.255.0.0 Eng
```

Creating a VFS on a specific node:

```
cli% createvfs -node 1 -cpg examplecpg -size 1T -snapquota enable 10.10.10.10 255.255.0.0 examplevfs
```

Creating a VFS where overlapping subnets are allowed:

```
cli% createvfs -fpg examplefpg -subnetoverlap 10.10.10.10 255.255.0.0 examplevfs
```

createvlun

Syntax

The syntax for the `createvlun` command can be one of the following:

Port presents:

```
createvlun [options] <VV_name | VV_set> <LUN> <node:slot:port>
```

Host set:

```
createvlun [options] <VV_name | VV_set> <LUN> <host_set>
```

Host sees:

```
createvlun [options] <VV_name | VV_set> <LUN> <host_name>
```

Matched set:

```
createvlun [options] <VV_name | VV_set> <LUN> <node:slot:port> <host_name>
createvlun [options] <VV_name | VV_set> <LUN> <host_name> <node:slot:port>
```

Description

The `createvlun` command creates a VLUN template that enables export of a virtual volume as a SCSI VLUN to a host or hosts. A SCSI VLUN is created when the current system state matches the rule established by the VLUN template.

There are four types of VLUN templates:

- Port presents—created when only the node:slot:port are specified. The VLUN is visible to any initiator on the specified port.
- Host set—created when a host set is specified. The VLUN is visible to the initiators of any host that is a member of the set.
- Host sees—created when the hostname is specified. The VLUN is visible to the initiators with any of the host's WWNs.
- Matched set—created when both hostname and node:slot:port are specified. The VLUN is visible to initiators with the host's WWNs only on the specified port. Conflicts between overlapping VLUN templates are resolved using prioritization, with port presents templates having the lowest priority and matched set templates having the highest.

Options

`-f`

Specifies that the operation is forced and that the VLUN is created even if the specified virtual volume has existing VLUNs. Unless the `-f` option is specified, the command asks for confirmation if a virtual volume is already exported in a VLUN template.

`-cnt <number>`

Specifies that a sequence of VLUNs, as specified by the num argument, are exported to the same system port and host that is created. The num argument can be specified as any integer. For each VLUN created, the .int suffix of the VV_name specifier and LUN are incremented by one.

`-novcn`

Specifies that a VLUN Change Notification (VCN) not be issued after export. For direct connect or loop configurations, a VCN consists of a Fibre Channel Loop Initialization Primitive (LIP). For fabric configurations, a VCN consists of a Registered State Change

Notification (RSCN) that is sent to the fabric controller.

-ovrd

Specifies that existing lower priority VLUNs will be overridden, if necessary. Can only be used when exporting to a specific host.

Specifiers

<VV_name | VV_set>

Specifies the virtual volume or virtual volume set name, using up to 31 characters in length. The volume name is provided in the syntax of basename.int. The VV set name must start with "set:".

<LUN>

Specifies the LUN as an integer from 0 through 16383. Alternatively n+ can be used to indicate a LUN should be auto assigned, but be a minimum of n, or m-n to indicate that a LUN should be chosen in the range m to n. In addition the keyword "auto" may be used and is treated as 0+.

<host_name>

Specifies the host where the LUN is exported, using up to 31 characters.

<host_set>

Specifies the host set where the LUN is exported, using up to 31 characters in length. The set name must start with "set:".

<node:slot:port>

Specifies the system port of the virtual LUN export.

node

Specifies the system node, where the node is a number from 0 through 7.

slot

Specifies the PCI bus slot in the node, where the slot is a number from 0 through 5.

port

Specifies the port number on the FC card, where the port number is 1 through 4.

Authority

- Super, Edit
- Any role granted the `vlun_create` right

Usage

- If a volume is already exported as a VLUN, you will be prompted for confirmation if a new export of the same volume is attempted.
- The host and port can both be supplied when issuing this command (matched set). This is the most restrictive access as both the host name and port must match before access to the VLUN is granted.
- Verify the creation of VLUNs by issuing the **showvlun** command.
- Conflicts between overlapping VLUN templates are resolved by a priority order among templates with matched set being the highest and port presents the lowest.

Examples

The following example exports virtual volume test on LUN 2:

```
cli% createvlun test 2 testhost
```

The following will export a VV test to a host set:

```
cli% createvlun test 1 set:host_set
```

The following will export a VV set test_set to a host:

```
cli% createvlun set:test_set 1 host
```

The following will export a VV set test_set to a host set:

```
cli% createvlun set:test_set 1 set:host_set
```

createvv

Syntax

createvv [options] *<usr_cpg>* *<vvname>*[.*<index>*] *<size>*[g|G|t|T]

Description

The createvv command creates volumes that are provisioned from one or more common provisioning groups. Volumes can be fully provisioned from a CPG or can be thinly provisioned. Users can optionally specify a CPG for snapshot space for fully provisioned volumes and thinly provisioned volumes.

Options

-templ *<tname>*

Uses the options defined in template *<tname>* created using the **createtemplate** command. Options specified in the template are read-only or read-write (using the -nro or -nrw options of the createtemplate command). New options at creation time can overwrite the read-write options. The default values of the read-only options cannot be overwritten. If not included, the -size and -usr_cpg options are automatically treated as read-write even if the other unincluded properties are marked read-only.

-shared

Specifies that the system will try to share the logical disks among the VVs. This option can only be used with the -cnt option. This option cannot be used with the -tdvv or -tpvv option.

-wait *<secs>*

If the command would fail due to the lack of clean space, the -wait option specifies the number of seconds to wait for the system to clean the dirty space before returning. If -wait 0 is issued, the command returns immediately. If this option is not used, the command will keep waiting for dirty chunklets to be cleaned if enough space will be available with the dirty chunklets cleaned. This option cannot be used with the -tdvv or -tpvv option.

-i *<ID>*

Specifies the ID of the volume. By default, the next available ID is chosen.

-cnt *<count>*

Specifies the number of identical VVs to create. This must be an integer from 1 through 999. If <id> is specified with -i option, it is used as the ID of the first VV and incremented by 1 for each subsequent volume.

-exp *<time>*{d|D|h|H|m|M}

Specifies the relative time from the current time that volume expires. *<time>* is a positive integer value and in the range of 1 minute - 1825 days. Specify the time in days, hours, or minutes. Use d or D for days, h or H for hours, and m or M for minutes following the entered time value.

-comment *<comment>*

Specifies any additional information up to 511 characters for the volume.

-f

Silences the system confirmation for creating a volume with retention time (`-retain`) or setting the host DIF policy (`-pol xxx_host_dif`).

-retain <time>{d|D|h|H|m|M}

Specifies the amount of time, relative to the current time, that the volume is retained. `<time>` is a positive integer value and in the range of 1 minute - 1825 days. Use `d` or `D` for days, `h` or `H` for hours, and `m` or `M` for minutes following the entered time value.

If the volume is not in any domain, then its retention time cannot exceed the value of the system `VVRetentionTimeMax`. The default value for the `VVRetentionTimeMax` is 14 days.

If the volume belongs to a domain, then its retention time cannot exceed the value of the domain `VVRetentionTimeMax`, if set. You cannot remove or reduce the retention time after it is set. If the volume has its retention time set, it cannot be removed within its retention time. If both expiration time and retention time are specified, then the retention time cannot be longer than the expiration time.

This option requires the HPE 3PAR Virtual Lock license. Contact your local service provider for more information.

-spt <sectors_per_track>

Defines the virtual volume geometry sectors per track value that is reported to the hosts through the SCSI mode pages. The valid range is between 4 to 8192 and the default value is 304.

-hpc <heads_per_cylinder>

Allows you to define the virtual volume geometry heads per cylinder value that is reported to the hosts though the SCSI mode pages. The valid range is between 1 to 255 and the default value is 8.

-pol <pol>[,<pol>...]

Specifies the following policies that the created virtual volume follows.

With no argument specified, defaults to `stale_ss`.

stale_ss

Specifies that invalid snapshot volumes are permitted. Failure to update snapshot data does not affect the write to the base volume, but the snapshot is considered invalid.

no_stale_ss

Specifies that invalid snapshot volumes are not permitted. Failure to update a snapshot is considered a failure to write to the base volume.

one_host

This constrains the export of a volume to one host or one host cluster (when cluster names may be used as a host name).

no_one_host

This policy should only be used when exporting a virtual volume to multiple hosts for use by a cluster-aware application, or when "port presents" VLUNs are used. This is the default policy setting.

3par_host_dif

This policy enables support for 3PAR implementation of host based DIF.

std_host_dif

This policy enables support for standard SCSI implementation of host based DIF.

no_host_dif

This policy disables support for host DIF.

zero_detect

This policy enables the storage system to scan for zeros in the incoming write data. This feature when used during physical copy to a TPVV will avoid allocating space for blocks containing zero. This feature when used with a Thin Persistence license will reclaim allocated space when zero blocks are written to the TPVV. This policy is only applicable for the base TPVV.

no_zero_detect

This policy disables the storage system to scan for zeros in the incoming write data to reclaim allocated space on the volume. This is the default policy setting.

-snp_cpg <snp_cpg>

Specifies the name of the CPG from which the snapshot space will be allocated.

-snp_aw <percent>

Enables a snapshot space allocation warning. A warning alert is generated when the reserved snapshot space of the VV exceeds the indicated percentage of the VV size.

-snp_al <percent>

Sets a snapshot space allocation limit. The snapshot space of the VV is prevented from growing beyond the indicated percentage of the virtual volume size.

Use the following options when creating thinly provisioned volumes:

-tdvv

Deprecated. Should use `-dedup`.

-compr

Specifies that the volume should be a thinly compressed volume.

-dedup

Specifies that the volume should be a thinly deduplicated volume which is a thinly provisioned volume with inline data deduplication. This option can only be used with CPG that has SSD (Solid State Drive) device type.

-tpvv

Specifies that the volume should be a thinly provisioned volume.

-minalloc <size>

This option specifies the default allocation size (in MB) to be set. Allocation size specified should be at least (number-of-nodes * 256) and less than the CPG grow size.

-usr_aw <percent>

This option enables user space allocation warning. Generates a warning alert when the user data space of the TPVV exceeds the specified percentage of the virtual volume size.

-usr_al <percent>

Indicates the user space allocation limit. The user space of the TPVV is prevented from growing beyond the indicated percentage of the virtual volume size. After this limit is reached, any new writes to the virtual volume will fail.

Specifiers

<usr_cpg>

Specifies the name of the CPG from which the volume user space will be allocated. If the `-tdvv` or `-tpvv` option is specified, the volume is thinly provisioned. Otherwise, the volume is fully provisioned. If the name of the CPG is to be taken from a template, this field should be "-".

<vvname>[.<index>]

Specifies a VV name up to 31 characters in length. If the `-cnt` option is used, the optional decimal number `<index>` specifies the name of the first VV(`<vvname>.<index>`). The `<index>` is incremented by 1 for each subsequent VV. The `<index>` must be an integer from 0 to 999999. All VV names have the same length constraint.

<size>[g|G|t|T]

Specifies the size for the volume in MB. The volume size is rounded up to the next multiple of 256 MB. The size should be an integer. An optional suffix (with no whitespace before the suffix) will modify the units to GB (g or G suffix) or TB (t or T suffix). If the size is to be taken from a template, this field should be "-".

Authority

- Super, Edit
- Any role granted the `vv_create` right

Usage

- Access to all domains is required to run this command with logical disks specified.
- Use of this command with the `-tdvv` or `-tpvv` option requires a Thin Provisioning license. Contact your local HPE representative for information.
- The `-templ` option is not valid for volumes that are fully provisioned.
- Using the `-shared` option can be more efficient when creating a large number of small volumes since the system will attempt to share the underlying logical disks. However, if volumes are subsequently removed, this can lead to unused space in the logical disks that cannot be reclaimed. Using `-cnt` without the `-shared` option avoids this problem, but can be inefficient when creating small volumes due to the overhead of creating a larger number of unshared LDs for the volumes.
- The `-shared` and `-wait` option cannot be used when creating thinly provisioned virtual volumes with the `-tdvv` or `-tpvv` option.
- `-minalloc` can be used only with `-tpvv` or `-tdvv`.
- `-minalloc` can be used only with `-tpvv` or `-tdvv`.
- Verify the creation of virtual volumes by issuing the **showvv** command.
- Maximum size of Dedup and Compressed Volumes is 16 TB. Minimum size of Compressed volumes is 16GB. Maximum VV size of Fully provisioned volumes is 128 TB.

  For this command MB = 1048576 bytes, GB = 1024MB, and TB = 1024GB.

Examples

- Creating a 10G TPVV named tpvv1 whose user space is allocated from the common provisioning group cpg1. The user space allocation warning is set to 50 and the user space allocation limit is set to 75:

```
cli% createvv -tpvv -usr_aw 50 -usr_al 75 cpg1 tpvv1 10G
```

- Creating 3 VVs vv1.2, vv1.3, vv1.4:

```
cli% createvv -cnt 3 cpg1 vv1.2 1g
```

- Creating 2 VVs vva.0 and vva.1 which may share LDs:

```
cli% createvv -cnt 2 -shared cpg1 vva 1g
```

- Creating a TPVV name tpvv1 with the template, temp2:

```
cli% createvv -tpvv -templ temp2 cpg2 tpvv1 2g
```

- Creating a TDVV name tdvv1 with the template, temp2:

```
cli% createvv -tdvv -templ temp2 cpg2 tdvv1 2g
```

- Creating a thinly compressed volume of size 16GB:

```
cli% createvv -compr cpg2 vv1 16g
```

- Creating 3 TPVVs vv1.2, vv1.3, vv1.4:

```
cli% createvv -cnt 3 -tpvv cpg1 vv1.2 1g
```

- Creating a TPVV with default allocation size of 2GB:

```
cli% createvv -tpvv -minalloc 2048 cpg2 tpvv1 1g
```

- The following example creates a thinly compressed volume of size 16GB:

```
cli% createvv -compr cpg2 vv1 16g
```

createvvcopy

Syntax

The syntax for the `createvvcopy` command can be one of the following:

```
createvvcopy [options] -p <parent_volume | parent_volumeset>
<destination_volume | destination_volumeset>
```

```
createvvcopy -p <parent_volume | parent_volumeset> -online [-snp_cpg
<snap_cpg>] [{-tpvv | -dedup | -compr}] [-wwn <wwn>] <destination cpg>
<destination volume>
```

```
createvvcopy -r [options] <destination_volume | destination_volumeset>
```

```
createvvcopy -halt <destination_volume | destination_volumeset>
```

Description

The `createvvcopy` command creates a full physical copy of a virtual volume or a read-write virtual copy on another virtual volume. This enables you to create copies of virtual volumes to perform tasks such as moving data to a larger virtual volume or creating a copy of data for testing.

Options

`-p <parent_volume | parent_volumeset>`

Specifies that a snapshot of the specified parent volume or volume set is created and copied to a specified destination volume or volume set (as indicated with the `<destination_volume | destination_volumeset>` specifier). The set name must start with "set:". The parent volume must be a base volume or read-write snapshot. This option cannot be used with the `-r` option. Also note that a volume set cannot be specified with the `-online` option.

`-r`

Specifies that the destination volume be resynchronized with its parent volume using a saved snapshot so that only the changes since the last copy or resynchronization must be copied. A snapshot of the parent saved using the `-s` option in an earlier `createvvcopy` instance is required for the resynchronization operation. This old snapshot is replaced by a new snapshot for the next resynchronization operation (the `-s` option is implied, and need not be specified). This option cannot be used with the `-p` option.

-halt

Specifies that an ongoing physical copy to be stopped. This will cause the destination volume to be marked with the 'cpf' status, which will be cleared up when a new copy is started.

-s

Saves the snapshot of the source volume after the copy of the volume is completed. This enables a fast copy for the next resynchronization. If not specified, the snapshot is deleted and the association of the destination volume as a copy of the source volume is removed. The `-s` option is implied when the `-r` option is used and need not be explicitly specified.

-b

Specifies that this command blocks until the operation is completed. If not specified, the `createvvcopy` command operation is started as a background task.

-pri <high|med|low|auto>

Specifies the priority of the copy operation when it is started. This option allows the user to control the overall speed of a particular task. If this option is not specified, the `createvvcopy` operation is started with default priority of medium. Auto priority can be used only if `-online` option is specified and this operation is started with auto as the default priority. Auto priority can be used only if `-online` option is specified and this operation is started with auto as the default priority. High priority indicates that the operation will complete faster. Low priority indicates that the operation will run slower than the default priority task. If the priority is set to auto, then the array controls the speed of the operation based on the load in the array for all flash models. Auto priority behaves like medium priority for the converged flash models of the arrays containing Nearline (NL) and Fast Class (FC) drives. This option cannot be used with `-halt` option.

-online

Specifies that the copy is to be performed online. This means that the destination can be immediately exported and is automatically created by this command, meaning that it cannot exist prior to executing this command. This option requires the specification of a CPG for the destination volume and allows for the options listed below. Finally, `-online` requires the `-p` option and is incompatible with the `-r`, `-halt`, `-s`, and `-b` options.

-addtoset <set_name>

Adds the VV copies to the specified VV set. The set will be created if it does not exist. Can only be used with `-online` option.

-skip_zero

When copying from a thin provisioned source, only copy allocated portions of the source VV. Using this option can save time, but only use it if the destination is newly created or has been reinitialized to zero. If there is pre-existing data in the destination VV, it will not be overwritten to match the source VV if the same offset is not allocated in the source. This option cannot be combined with `-r`, `-halt`, or `-online` options.

The following options can only be used when the `-online` option is specified:

-tpvv

Indicates that the VV the online copy creates should be a thinly provisioned volume. Cannot be used with the `-dedup` option.

-tdvv

This option is deprecated, see `-dedup`.

-dedup

Indicates that the VV the online copy creates should be a thinly deduplicated volume, which is a thinly provisioned volume with inline data deduplication. This option can only be used with a CPG that has SSD (Solid State Drive) device type. Cannot be used with the -tpvv option.

-compr

Indicates that the VV the online copy creates should be a compressed virtual volume.

-snp_cpg <snp_cpg>

Specifies the name of the CPG from which the snapshot space will be allocated.

-wwn <wwn>

Specifies that the WWN of the online copy virtual volume use <wwn>. If this option is not used, the system will automatically choose the WWN based on the system serial number, the volume ID, and the wrap counter.

Specifiers

<destination_cpg>

Specifies the destination CPG to use for the destination volume if the -online option is specified.

<destination_volume | destination_volumeset>

Specifies the destination volume name or volume set name for the copy operation using up to 31 characters in length. The volume set name must start with "set:". The destination volume (or each member of the destination volume set) must be a writable base volume (not a snapshot) of equal or greater size than a parent volume (if specified) and it must not be exported as a VLUN.

Authority

- Super, Edit
- Any role granted the vvcopy_create right

Usage

- The source and destination volumes must be writable. If the online option is specified, the destination volume must not already exist.
- The destination volume cannot be a system volume or a File Persona volume.
- The destination volume cannot be exported as a VLUN before or during the createvvcopy command process (except if online, see following).
- The destination volume must be greater than or equal in size to the source volume.
- If both the destination and source are volume sets then they must contain the same number of members.
- If the source is a volume set, then the destination name can contain various patterns that are used to form the copy volume name. (For more information, enter the clihelp sub,vvnamepat command.) The time/date related patterns are not supported in this command.
- If the -s option is specified to save a snapshot for fast resynchronization and the snapshot goes stale, the copy fails.
- A physical copy of a virtual volume fails in any situation that a snapshot copy fails or when there is insufficient snapshot space or I/O errors.
- The createvvcopy command can be issued multiple times. However, only two active physical copy tasks per node in cluster are allowed to run concurrently.
- Any additional physical copy tasks are queued, pending the completion of the active physical copy tasks. Online copy tasks are performed using the region mover, which allows only nine region move tasks to run concurrently.
- Multiple physical copy operations can occur simultaneously. Host-initiated I/O operations and those operations initiated by issuing the createvvcopy command are executed at the same priority level. As a result, noticeable performance degradation from a host perspective can be observed.

- Issuing the `createvvcopy` command results in the creation of a temporary snapshot and, in the case of `-r`, a resynchronization (resync) snapshot.
- The temporary and resynchronization snapshots cannot be deleted while the copy is in progress.
- Upon completion of the copy, the temporary snapshot is automatically deleted if the `-s` option is not specified.
- If the resynchronization snapshot is saved, it can later be manually deleted. If the resynchronization snapshot is deleted, later resynchronization is not possible.
- If the -s option is not specified, the relationship between the destination volume and source volume is not retained.
- Enter the `showvv` command to verify that a virtual volume copy has been made.
- Enter the `showvv -d` command to display the number of blocks remaining to be copied.
- If the source or destination are volume sets, then the `createvvcopy` command creates consistent group snapshots at the same point in time so that related structures on different volumes in the group remain consistent.
- On systems that support zero_detect, if destination volume is a TPVV, zero_detect policy is enabled at start of physical copy. After physical copy completes, the policy is changed to no_zero_detect, even when the destination TPVV had zero_detect policy enabled before the start of physical copy operation.
- When Thin Persistence license is present, `createvvcopy` command will reclaim allocated space when zero blocks are written to the destination TPVV.
- If the `-pri` option is not specified, the physical copy will start at the default priority. If the tunes and host I/O are running then it is recommended to start `vvcopy` with default priority.
- The `-online` option for an online copy automatically creates a destination volume, so unlike other physical copies, the destination volume must not exist. The destination volume requires a CPG to be specified. It will be a TPVV or a dedup, compressed VV. It can optionally have a snapshot CPG associated with it.
- The destination volume can be immediately exported (hence the reason for the online option). The online option is incompatible with the `-r`, `-halt`, `-s`, and `-b` options. The destination volume to be created can be named with pattern substitution based on the source volume name, see `help sub,vvnamepat` for details.

Examples

The following example displays the creation of a copy of the virtual volume vv1:

```
cli% createvvcopy -p vv1 vv2
Started copy. child=vv2 parent=vv1
```

Creating a set of copies for the volumes in set vvcopyset, keeping snapshots around for quick resynchronization:

```
cli% createvvcopy -s -p set:vvcopyset set:copies
Child Parent  Status TaskID
s1    vv1    started     14
s2    vv2    started     15
s3    vv3     queued     16
```

Re-synchronize the volumes in the volume set "copies":

```
cli% createvvcopy -r set:copies
Child Parent  Status TaskID
s1           started     17
s2           started     18
s3            queued     19
```

Creating a set of copies for the volumes in set vvcopyset, mapping the parent volumes to destination volumes with a pattern:

```
cli% createvvcopy -p set:vvcopyset @vvname@-copy
Child     Parent  Status TaskID
vv1-copy  vv1    started     20
vv2-copy  vv2    started     21
vv3-copy  vv3     queued     22
```

Creating an online copy of vv1 that is named vv2 that is a thin-provisioned volume, with cpg1 as the storage space for vv2:

```
cli% createvvcopy -p vv1 -online -tpvv cpg1 vv2
```

Create an online copy of vv1 that is named vv2 which is fully-provisioned, using cpg1 as its user space and cpg2 as its snapshot space:

```
cli% createvvcopy -p vv1 -online -snp_cpg cpg2 cpg1 vv2
```

Creating online copies for each member of vvset, adding each copy to newset:

```
cli% createvvcopy -p set:vvset -online -addtoset newset -tpvv FC_r1 @vvname@-copy
```

createvvset

Syntax

```
createvvset [options] <setname> [<VV | pattern>...]
```

Description

The `createvvset` command defines a new set of Virtual Volumes (VV) and provides the option of assigning one or more existing VVs to that set. The command also allows the addition of VVs to an existing set by use of the -add option.

Options

`-add`

Specifies that the VVs listed should be added to an existing set. At least one VV must be specified.

`-cnt <num>`

Adds a sequence of <num> VVs starting with "vvname". vvname should be of the format <basename>.<int> For each VV in the sequence, the .<int> suffix of the vvname is incremented by 1.

`-comment <comment>`

Specifies any comment or additional information for the set. The comment can be up to 255 characters long. Unprintable characters are not allowed.

`-domain <domain>`

Creates the VV set in the specified domain. For an empty set the default is to create it in the current domain, or no domain if the current domain is not set. A VV set must be in the same domain as its members; if VVs are specified as part of the creation then the set will be created in their domain. A domain cannot be specified when adding a VV to an existing set with the -add option.

Specifiers

`<setname>`

Specifies the name of the vv set to create/add to, using up to 27 characters in length.

`<VV | pattern>...`

Optional list of VVs or glob-style pattern matching VVs which are to be included in the VV set. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.) If no <VV> or <pattern> is specified a VV set with no VVs is created. If the -cnt option is used then a VV must be specified.

Authority

- Super, Edit
- Any role granted the `vvset_create` right can create a volume set
- Any role granted the `vvset_set` right can add volumes to a volume set

Usage

None.

Examples

To create an empty vvset:

```
cli% createvvset vvset
```

To add a VV to the set:

```
cli% createvvset -add vvset vv1
```

To create a VV set with a comment and a collection of VVs in it:

```
cli% createvvset -comment "Our Oracle VVs" -cnt 10 oravv oravv.0
```

To create a VV set with VVs in it:

```
cli% createvvset vvset vv1 vv2
```

Or:

```
cli% createvvset set:vvset vv1 vv2
```

Dismiss Commands

dismisspd

Syntax

```
dismisspd <PD_ID>...
```

Description

The `dismisspd` command removes PD definitions from system use.

Options

None.

Specifiers

<PD_ID>

Specifies the PD(s), identified by integers, to be removed from system use.

Authority

- Super, Service
- Any role granted the `pd_dismiss` right

Usage

- Access to all domains is required to run this command.
- A PD that is in use cannot be removed.
- Verify the removal of a PD by issuing the **showpd** command.

Examples

The following example removes a PD with ID 1:

```
cli% dismisspd 1
```

dismissrcopylink

Syntax

Syntax for remote copy over IP (RCIP) is as follows:

```
dismissrcopylink <target_name> <node:slot:port:IP_address>...
```

Syntax for remote copy over FC (RCFC) is as follows:

```
dismissrcopylink <target_name> <node:slot:port:WWN>...
```

Description

The `dismissrcopylink` command removes one or more links (connections) created with the `admitrcopylink` command to a target system.

Options

None.

Specifiers

`<target_name>`

The target name, as specified with the `creatercopytarget` command.

`<node:slot:port:IP_address>...`

Specifies the node, slot, and port of the Ethernet port on the local system and an IP address of the peer port on the target system.

`<node:slot:port:WWN>...`

Specifies the node, slot, and port of the Fibre Channel port on the local system and World Wide Name (WWN) of the peer port on the target system.

Authority

- Super, Edit
- Any role granted the `rcopylink_dismiss` right

Usage

- This command requires the HPE 3PAR Remote Copy Software license. Contact your local service provider for more information.
- Access to all domains is required to run this command.
- This command should only be used to remove sending links.
- This command cannot be used to remove the last link of a target system with started groups.
- This command terminates with a list of one or more links to be dismissed.
- For IP targets, this list is made up of pairs composed of the node containing the Ethernet port on the primary system and an IP address on the backup system.
- For Fibre Channel targets, this list is made up of sets with the node, slot, and port of the Fibre Channel adaptor port on the primary system and a WWN address on the backup system.
- IP targets are made up of pairs composed of the node, slot and port of the ethernet port on the local system and an IP address of the peer port on the target system.
- FC targets are made up of sets with the node, slot, and port of the FC port on the local system and WWN of the peer port on the target system.

Examples

The following example removes the link from the Ethernet port located at 2:2:2 of System2:

```
dismissrcopylink System2 2:2:2:193.1.2.11
```

dismissrcopytarget

Syntax

`dismissrcopytarget [options] <target_name> <group_name>`

Description

The `dismissrcopytarget` command removes a remote copy target from a remote copy volume group.

Options

-f

Suppresses user interaction with the command. Without this option, the command requires confirmation before proceeding with its operation.

Specifiers

<target_name>

The name of the target to remove.

<group_name>

The name of the group that currently includes the target.

Authority

- Super, Edit
- Any role granted the `rcopytarget_dismiss` right

Usage

- This command requires the HPE 3PAR Remote Copy Software license. Contact your local service provider for more information.
- The `dismissrcopytarget` command does not allow the removal of the last target in a group.

Examples

The following example removes target Target1 from Group1:

```
cli% dismissrcopytarget Target1 Group1
```

dismissrcopyvv

Syntax

`dismissrcopyvv [options] {<pattern> | <VV_name>} <group_name>`

Description

Removes a virtual volume from a Remote Copy volume group.

Options

-f

Suppresses user interaction with the command, and does not ask for confirmation.

-pat

Treats patters as glob-style patterns and dismisses all Remote Copy volumes that match the specified pattern from the Remote Copy group. Required if using the *<pattern>* specifier.

-keepsnap

Specifies retaining the local volume resync snapshot. The retained snapshot reflects the state of the secondary volume. If you readmit the volume to a Remote Copy group, you can use the retained snapshot as the starting snapshot. The snapshot name begins with `sv.rcpy`.

-removevv

Removes remote side volumes.

-keepalua

Keeps the ALUA state of local volume from changing. Required during Peer Motion migrations of Remote Copy Groups. See *Volume migration in a Remote Copy Primary group* in the *HPE 3PAR Peer Motion User Guide* for details.

Specifiers

<VV_name>

Specifies the name of the volume to remove. Volumes are added to a group with the **admitrcopyvv** command.

<pattern>

Specifies the name of the group that currently includes the virtual volume.

<group_name>

Specifies a glob-style pattern. You can repeat this specifier to dismiss multiple Remote Copy volumes. If you do not use this specifier, you must use *<VV_name>*.

Authority

- Super, Edit
- Any role granted the `rcopyvv_dismiss` right

Usage

- This command requires the HPE 3PAR Remote Copy Software license. Contact your local service provider for more information.
- Command requires access to all domains.
- The `-keepsnap` and `-removevv` options are mutually exclusive.
- The `dismissrcopyvv` command removes any Remote Copy synchronization snapshots affiliated with the removed volume.
- If the target of a group has the `mirror_config` policy set, and the group is a primary group, this command is mirrored to that target and the volume is removed from the corresponding secondary group. If the policy is set and the group is a secondary, this command fails.
- If you remove a group that contains an automatically created volume set, specified when creating its virtual volumes (`createvlun set:<RCP_group_name>...`), dismissing a volume from the group also removes its VLUN.

Examples

The following example dismisses virtual volume vv1 from Group1:

```
cli% dismissrcopyvv vv1 Group1
```

The following example dismisses all virtual volumes that start with the name testvv from Group1:

```
cli% dismissrcopyvv -pat testvv* Group1
```

The following example dismisses volume vv1 from Group1 and removes the corresponding volumes of vv1 on all the target systems of Group1:

```
cli% dismissrcopyvv -removevv vv1 Group1
```

The following example dismisses volume vv2 from Group2 and retains the resync snapshot associated with vv2 for this group:

```
cli% dismissrcopyvv -keepsnap vv2 Group2
```

Free Command

freespace

Syntax

```
freespace [options] <VV_name>...|<pattern>...
```

Description

The `freespace` command frees snapshot administration and snapshot data spaces from a Virtual Volume (VV) if they are not in use.

Options

`-pat`

Removes the snapshot administration and snapshot data spaces from all the virtual volumes that match any of the specified glob-style patterns.

`-f`

Suppresses the prompt for confirmation before removing the snapshot administration and snapshot data space of each volume.

Specifiers

`<VV_name>...`

Specifies the virtual volume name, using up to 31 characters.

`<pattern>...`

Specifies a glob-style pattern. This specifier can be repeated to compact multiple volumes. If this specifier is not used, the VV_name specifier must be used.

Authority

- Super, Edit
- Any role granted the `space_free` right

Usage

This command fails if the virtual volume is an old-style thin provisioned VV (created on a 2.2.4 release or earlier) or it has snapshots.

Examples

The following example demonstrates how to remove SA and SD space from virtual volume testd:

```
cli% freespace testd
```

Grow Commands

growfpg

Syntax

```
growfpg <fpgname> <size>{t|T|g|G}
growfpg -recover_storage
```

Description

The `growfpg` command grows a file provisioning group of the given name by the size specified, within the CPG associated with the base FPG.

Options

`-nowwn`

For use with `-rcopy`. Do not use the same VV WWN for secondary volumes.

`-recover_storage`

Rescans all VVs associated with FPGs and ensures all storage is available for use.

Specifiers

`<fpgname>`

The name of the filesystem to be grown.

`<size>`

The size of the filesystem to be grown. The current FPG size plus `<size>` cannot exceed 128T.

Authority

- Super, Edit
- Any role granted `fpg_grow` right

Usage

- Access to all domains is required to run this command.
- For each grow undertaken, VVs will be grown more evenly, with priority given to those which are considered full by the 3PAR storage system.
- The grow operation will not be visible in the **showfpg** output until the FPG is activated.
- The `growfpg` command examines the existing FPG VVs to see if there is a remote copy relationship. If one is found, the new volumes will be added to the appropriate group and target automatically. If an error occurs, during this process the grow will not be undone. Secondary volumes are automatically created and have the base VV name with "_sec" appended.
- An FPG must be grown by a minimum of 100GiB.

Examples

The following example grows examplefpg by 8T:

```
cli% growfpg examplefpg 8T
```

growvv

Syntax

```
growvv [options] <VV_name> size[g|G|t|T]
```

Description

The `growvv` command increases the size of a virtual volume.

Options

`-f`

Suppresses the requested confirmation before growing a virtual volume size from under 2T to over 2T.

Specifiers

`<VV_name>`

Specifies the name of the virtual volume to grow.

`<size>[g|G|t|T]`

Specifies the size in MB to be added to the volume user space. The size must be an integer in the range from 1 MB to 128 TB. The volume size is rounded up to the next multiple of 256 MB. If the value is followed by a g or G (no whitespace before the suffix), the size is in GB; if by a t or T, the size is in TB.

Authority

- Super, Edit
- Any role granted the `vv_grow` right

Usage

- Access to all domains is required to run this command with logical disks specified.
- For this command MB = 1048576 bytes, GB = 1024MB, and TB = 1024GB.
- If the volume to be grown is in a Remote Copy group, the group must be stopped before the grow operation is permitted.
- The grow operations of Remote Copy volumes are coordinated between the primary and secondary targets. A coordinated grow can be started from either the primary or secondary target. Volumes on remote targets are grown to the intended size of the local volume. If a target cannot be contacted or Remote Copy is not started, only the local volume will be grown. A check is first made to see if the local and remote volumes can be grown by the requested amount. If this succeeds, all the volumes are grown; however, if other volumes are grown during this process, it is possible that remote volumes are grown and the local volume is not. If this occurs, remote volumes must be grown manually before the group can be started.
- If the volume to be grown is a member of a File Provisioning Group, it cannot be grown.

Examples

The following example displays the enlarging of the virtual volume vv0 by 10G:

```
cli% growvv vv0 10g
```

Hist Commands

histch

Syntax

```
histch [options <arg>]
```

Description

The `histch` command displays a histogram of service times in a timed loop for individual chunklets.

Options

`-ld <LD_name>`

Specifies the logical disk name from which chunklet statistics are sampled.

`-ch <chunklet_num>`

Specifies that statistics are limited to only the specified chunklet, identified by number.

`-metric both|time|size`

Selects which metric to display. Metrics can be one of the following:

- both—Display both I/O time and I/O size histograms (default).
- time—Display only the I/O time histogram.
- size—Display only the I/O size histogram.

`-timecols <fcol> <lcol>`

For the I/O time histogram, shows the columns from the first column `<fcol>` through last column `<lcol>`. The available columns range from 0 through 31.

The first column (`<fcol>`) must be a value greater than or equal to 0, but less than the value of the last column (`<lcol>`).

The last column (`<lcol>`) must be less than or equal to 31.

The first column includes all data accumulated for columns less than the first column and the last column includes accumulated data for all columns greater than the last column.

The default value of `<fcol>` is 6.

The default value of `<lcol>` is 15.

`-sizecols <fcol> <lcol>`

For the I/O size histogram, shows the columns from the first column (`<fcol>`) through the last column (`<lcol>`). Available columns range from 0 through 15.

The first column (`<fcol>`) must be a value greater than or equal to 0, but less than the value of the last column (`<lcol>`) (default value of 3).

The last column (`<lcol>`) must be less than or equal to 15 (default value of 11).

The default value of `<fcol>` is 3.

The default value of `<lcol>` is 11.

-pct

Shows the access count in each bucket as a percentage. If this option is not specified, the histogram shows the access counts.

-prev | -begin

Histogram displays data either from a previous sample (`-prev`) or from when the system was last started (`-begin`). If no option is specified, the histogram shows data from the beginning of the command's execution.

-rw

Specifies that the display includes separate read and write data. If not specified, the total is displayed.

-d <secs>

Specifies the interval in seconds that statistics are sampled from using an integer from 1 through 2147483. If no count is specified, the command defaults to 2 seconds.

-iter <number>

Specifies that the histogram is to stop after the indicated number of iterations using an integer from 1 through 2147483647.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (`<col>`). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (`<dir>`) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

-filt <fspec>

Specifies that histograms below the threshold specified by the `<fspec>` argument are not displayed. The `<fspec>` argument is specified in the syntax of `<op>,<val_ms>, <count>`, where:

- `<op>`

  The `<op>` argument can be specified as one of the following:

  - r—Specifies read statistics.
  - w—Specifies write statistics.
  - t—Specifies total statistics.
  - rw—Specifies total read and write statistics.

-ni

<val_ms>

Specifies the threshold service time in milliseconds.

<count>

Specifies the minimum number of access above the threshold service time. When filtering is done, the `<count>` is compared with the sum of all columns starting with the one which corresponds to the threshold service time. For example, `-t,8,100` means to only display the rows where the 8ms column and all columns to the right adds up to more than 100.

Specifies that histograms for only non-idle devices are displayed. This option is shorthand for the option `-filt t,0,0`.

Specifiers

None.

Authority

Any role in the system

Usage

- For the following options, by default the histogram shows data from the start of the command.
- The **`setstatch`** command must be issued to enable statistics collection on chunklets before issuing the `histch` command.
- Each service time histogram column shows the number of accesses with service times between the point in time shown in the column's heading to the point in time shown in the heading of the column to its right (or infinity for the last column). Unless all columns are shown, the first column shows the number of accesses with service times less than the heading of the second column.
- Each I/O size histogram column shows the number of accesses with I/O size between the size shown in the column's heading to the size shown in the heading of the column to its right (or infinity for the last column).
- For all the `<name_or_pattern>` fields, the patterns are glob-style (shell-style) patterns. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.)

Examples

The following example displays one iteration of a histogram of service times for system chunklets:

```
cli% histch -iter 1
12:42:57 10/20/15 ------------Time (millisec)------------
Ldid      Ldname LdCh Pdid PdCh 0.26 0.53 1.05 2.1 4.2 8.4 17 34 67 135
   2 tp-0-sa-0.0    1   21    0    0    0    0   0   0   0  0  0  0   0
   2 tp-0-sa-0.0    0   45    0    0    0    0   0   0   0  0  0  0   0
   1 tp-0-sa-0.1    1   20    0    0    0    0   0   0   1  0  0  0   0
   1 tp-0-sa-0.1    0   42    0    0    0    0   0   0   1  0  0  0   0
   0 admin.usr.0    1   22    3    0    0    0   1   0   3  2  0  0   0
   0 admin.usr.0    0   44    3    0    0    0   0   1   3  2  0  0   0
-------------------------------------------------------------------------
           total                  0    0    0   1   1   8  4  0  0   0
```

For the previous example, before the `histch` command was issued, the `setstatch start` command was issued for chunklets 0 and 1 on logical disks admin.usr.0, tp-0-sa-0.1, and tp-0-sa-0.0.

histld

Syntax

```
histld [options <arg>] [<LD_name_or_pattern> ...]
```

Description

The `histld` command displays a histogram of service times for LDs in a timed loop.

Options

-vv {<VV_name>|<pattern>}...

Shows only logical disks that are mapped to virtual volumes with names matching any of the names or patterns specified. Multiple volumes or patterns can be repeated using a comma separated list.

-domain {<domain_name>|<pattern>}...

Shows only logical disks that are in domains with names matching any of the names or patterns specified. Multiple domain names or patterns can be repeated using a comma separated list.

-metric both|time|size

Selects which metric to display. Metrics can be one of the following:

- both—Display both I/O time and I/O size histograms (default).
- time—Display only the I/O time histogram.
- size—Display only the I/O size histogram.

-timecols <fcol> <lcol>

For the I/O time histogram, shows the columns from the first column <fcol> through last column <lcol>. The available columns range from 0 through 31.

The first column (<fcol>) requires a value greater than or equal to 0, and less than the value of the last column (<lcol>).

The last column (<lcol>) must be less than or equal to 31.

The first column includes all data accumulated for columns less than the first column and the last column includes accumulated data for all columns greater than the last column.

The default value of <fcol> is 6.

The default value of <lcol> is 15.

-sizecols <fcol> <lcol>

For the I/O size histogram, shows the columns from the first column (<fcol>) through the last column (<lcol>). Available columns range from 0 through 15.

The first column (<fcol>) must be a value greater than or equal to 0, but less than the value of the last column (<lcol>) (default value of 3).

The last column (<lcol>) must be less than or equal to 15 (default value of 11).

The default value of <fcol> is 3.

The default value of <lcol> is 11.

-pct

Shows the access count in each bucket as a percentage. If this option is not specified, the histogram shows the access counts.

-prev | -begin

Histogram displays data either from a previous sample (-prev) or from when the system was last started (-begin). If no option is specified, the histogram shows data from the beginning of the command's execution.

-rw

Specifies that the display includes separate read and write data. If not specified, the total is displayed.

-d <secs>

Specifies the interval in seconds that statistics are sampled from using an integer from 1 through 2147483. If no count is specified, the command defaults to 2 seconds.

-iter <number>

Specifies that the histogram is to stop after the indicated number of iterations using an integer from 1 through 2147483647.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

-filt <fspec>

Specifies that histograms below the threshold specified by the <fspec> argument are not displayed. The <fspec> argument is specified in the syntax of <op>,<val_ms>, <count>, where:

<op>

The <op> argument can be specified as one of the following:

- r—Specifies read statistics.
- w—Specifies write statistics.
- t—Specifies total statistics.
- rw—Specifies total read and write statistics.

<val_ms>

Specifies the threshold service time in milliseconds.

<count>

Specifies the minimum number of access above the threshold service time. When filtering is done, the <count> is compared with the sum of all columns starting with the one which corresponds to the threshold service time. For example, -t,8,100 means to only display the rows where the 8ms column and all columns to the right adds up to more than 100.

-ni

Specifies that histograms for only non-idle devices are displayed. This option is shorthand for the option -filt t,0,0.

Specifiers

[<LD_name_or_pattern>]

Specifies the LD(s) or pattern(s) for which the histogram data is collected.

Authority

Any role in the system

Usage

- Each service time histogram column shows the number of accesses with service times between the point in time shown in the column's heading to the point in time shown in the heading of the column to its right (or infinity for the last column). Unless all columns are shown, the first column shows the number of accesses with service times less than the heading of the second column.
- Each I/O size histogram column shows the number of accesses with I/O size between the size shown in the column's heading to the size shown in the heading of the column to its right (or infinity for the last column).
- For all the `<name_or_pattern>` fields, the patterns are glob-style (shell-style) patterns. (For more information, enter the `clihelp sub,globpat` command, or see **Glob-style patterns**.)
- If the `<LD_name_or_pattern>` specifier is used, then logical disks with names that match any of the patterns are listed, otherwise all logical disks are listed. These patterns are glob-style patterns. (For more information, enter the `clihelp sub,globpat` command, or see, **Glob-style patterns**.)
- Patterns are specified as regular expressions. (For more information on `regexps`, enter the `clihelp sub,regexpat` command, or see, **Regular expressions and patterns**.) Issuing `histld -n LD_name.*` displays histogram data for all LDs with a name beginning with LD_name.

Examples

The following example displays one iteration of a histogram of service times for all LDs:

```
cli% histld -iter 1
08:38:19 06/04/2015 ----------Time (millisec)---------- --------------Size (bytes)--------------
             Ldname  0.50   1  2  4  8 16 32 64 128 256    4k   8k 16k 32k 64k 128k 256k 512k 1m
        admin.usr.0     0   0  0  0  2  0  0  0   0   0     2    0   0   0   0    0    0    0  0
      .srdata.usr.0     0   0  0  0  0  0  0  0   0   0     0    0   0   0   0    0    0    0  0
               log0.0     0   0  0  0  0  0  0  0   0   0     0    0   0   0   0    0    0    0  0
            pdsld0.0     0   0  0  0  0  0  0  0   0   0     0    0   0   0   0    0    0    0  0
```

histpd

Syntax

`histpd [options <arg>]`

Description

The `histpd` command displays a histogram of service times for Physical Disks.

Options

`-w <WWN>`

Specifies the world wide name of the PD for which service times are displayed.

`-nodes <node_list>`

Specifies that the display is limited to specified nodes and physical disks connected to those nodes. The node list is specified as a series of integers separated by commas (e.g. 1,2,3). The list can also consist of a single integer. If the node list is not specified, all disks on all nodes are displayed.

`-slots <slot_list>`

Specifies that the display is limited to specified PCI slots and physical disks connected to those PCI slots. The slot list is specified as a series of integers separated by commas (e.g. 1,2,3). The list can also consist of a single integer. If the slot list is not specified, all disks on all slots are displayed.

-ports <port_list>

Specifies that the display is limited to specified ports and physical disks connected to those ports. The port list is specified as a series of integers separated by commas (e.g. 1,2,3). The list can also consist of a single integer. If the port list is not specified, all disks on all ports are displayed.

-devinfo

Indicates the device disk type and speed.

-metric both|time|size

Selects which metric to display. Metrics can be one of the following:

- both—Display both I/O time and I/O size histograms (default).
- time—Display only the I/O time histogram.
- size—Display only the I/O size histogram.

-timecols <fcol> <lcol>

For the I/O time histogram, shows the columns from the first column <fcol> through last column <lcol>. The available columns range from 0 through 31.

The first column (<fcol>) must be a value greater than or equal to 0, but less than the value of the last column (<lcol>).

The last column (<lcol>) must be less than or equal to 31.

The first column includes all data accumulated for columns less than the first column and the last column includes accumulated data for all columns greater than the last column.

The default value of <fcol> is 6.

The default value of <lcol> is 15.

-sizecols <fcol> <lcol>

For the I/O size histogram, shows the columns from the first column (<fcol>) through the last column (<lcol>). Available columns range from 0 through 15.

The first column (<fcol>) must be a value greater than or equal to 0, but less than the value of the last column (<lcol>) (default value of 3).

The last column (<lcol>) must be less than or equal to 15 (default value of 11).

The default value of <fcol> is 3.

The default value of <lcol> is 11.

-pct

Shows the access count in each bucket as a percentage. If this option is not specified, the histogram shows the access counts.

-prev | -begin

Histogram displays data either from a previous sample (-prev) or from when the system was last started(-begin). If no option is specified, the histogram shows data from the beginning of the command's execution.

-rw

Specifies that the display includes separate read and write data. If not specified, the total is displayed.

-d <secs>

Specifies the interval in seconds that statistics are sampled from using an integer from 1 through 2147483. If no count is specified, the command defaults to 2 seconds.

-iter <number>

Specifies that the histogram is to stop after the indicated number of iterations using an integer from 1 through 2147483647.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

-filt <fspec>

Specifies that histograms below the threshold specified by the <fspec> argument are not displayed. The <fspec> argument is specified in the syntax of <op>,<val_ms>,<count>, where:

- <op>

  The <op> argument can be specified as one of the following:

  - r—Specifies read statistics.
  - w—Specifies write statistics.
  - t—Specifies total statistics.
  - rw—Specifies total read and write statistics.

<val_ms>

Specifies the threshold service time in milliseconds.

<count>

Specifies the minimum number of access above the threshold service time. When filtering is done, the <count> is compared with the sum of all columns starting with the one which corresponds to the threshold service time. For example, -t,8,100 means to only display the rows where the 8ms column and all columns to the right adds up to more than 100.

-ni

Specifies that histograms for only non-idle devices are displayed. This option is shorthand for the option -filt t,0,0.

-devsvtime

Specifies the time spent in the drive, not including queue time. Shows the average and current time in ms.

-p <pattern>

Specifies the pattern of PDs displayed in the histogram. Patterns are used to filter and select the disks displayed in the histogram. If specified multiple times, each instance of the specified pattern adds additional candidate disks that match the pattern.

The following arguments can be specified as patterns for this option:

An item is specified as an integer, a comma-separated list of integers, or a range of integers specified from low to high.

-nd <item>

Specifies one or more nodes. Nodes are identified by one or more integers (item). Multiple nodes are separated with a single comma (e.g. 1,2,3). A range of nodes is separated with a hyphen (e.g. 0- 7). The primary path of the disks must be on the specified node(s).

-st <item>

Specifies one or more PCI slots. Slots are identified by one or more integers (item). Multiple slots are separated with a single comma (e.g. 1,2,3). A range of slots is separated with a hyphen (e.g. 0- 7). The primary path of the disks must be on the specified PCI slot(s).

-pt <item>

Specifies one or more ports. Ports are identified by one or more integers (item). Multiple ports are separated with a single comma (e.g. 1,2,3). A range of ports is separated with a hyphen (e.g. 0- 4). The primary path of the disks must be on the specified port(s).

-cg <item>

Specifies one or more drive cages. Drive cages are identified by one or more integers (item). Multiple drive cages are separated with a single comma (e.g. 1,2,3). A range of drive cages is separated with a hyphen (e.g. 0-3). The specified drive cage(s) must contain disks.

-mg <item>

Specifies one or more drive magazines. The "1." or "0." displayed in the CagePos column of `showpd` output indicating the side of the cage is omitted when using the `-mg` option. Drive magazines are identified by one or more integers (item). Multiple drive magazines are separated with a single comma (e.g. 1,2,3). A range of drive magazines is separated with a hyphen (e.g. 0-7). The specified drive magazine(s) must contain disks.

-pn <item>

Specifies one or more disk positions within a drive magazine. Disk positions are identified by one or more integers (item). Multiple disk positions are separated with a single comma (e.g. 1,2,3). A range of disk positions is separated with a hyphen (e.g. 0-3). The specified position(s) must contain disks.

-dk <item>

Specifies one or more physical disks. Disks are identified by one or more integers (item). Multiple disks are separated with a single comma (e.g. 1,2,3). A range of disks is separated with a hyphen (e.g. 0-3). Disks must match the specified ID(s).

-tc_gt <number>

Specifies that physical disks with total chunklets greater than the number specified be selected.

-tc_lt <number>

Specifies that physical disks with total chunklets less than the number specified be selected.

-fc_gt <number>

Specifies that physical disks with free chunklets greater than the number specified be selected.

-fc_lt <number>

Specifies that physical disks with free chunklets less than the number specified be selected.

-devid <model>

Specifies that physical disks identified by their models be selected. Models can be specified in a comma-separated list. Models can be displayed by issuing the `showpd -i` command.

-devtype <type>

Specifies that physical disks must have the specified device type (FC for Fast Class, NL for Nearline, SSD for Solid State Drive, SCM for Storage Class Memory) to be used. Device types can be displayed by issuing the `showpd` command.

-rpm <number>

Drives must be of the specified relative performance metric, as shown in the "RPM" column of the `showpd` command. For FC and NL drives, the number corresponds to both a performance measure and actual rotational speed. For SSD & SCM drives, this field is "N/A".

Disks that satisfy all of the specified characteristics are used.

For example, `-p -fc_gt 60 -fc_lt 230 -nd 2` specifies all the disks that have greater than 60 and less than 230 free chunklets and that are connected to node 2 through their primary path.

Specifiers

None.

Authority

Any role in the system

Usage

- Each service time histogram column shows the number of accesses with service times between the point in time shown in the column's heading to the point in time shown in the heading of the column to its right (or infinity for the last column). Unless all columns are shown, the first column shows the number of accesses with service times less than the heading of the second column.
- Each I/O size histogram column shows the number of accesses with I/O size between the size shown in the column's heading to the size shown in the heading of the column to its right (or infinity for the last column).
- For all the `<name_or_pattern>` fields, the patterns are glob-style (shell-style) patterns. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**).

Examples

The following example displays a histogram of service times for all PDs:

```
cli% histpd
```

The following example displays a histogram of service times for all PDs and shows the columns from 8ms to 8192 ms. Only PDs with more than 500 accesses with service time equal or greater than 16ms are displayed:

```
cli% histpd -begin -metric time -timecols 10 20 -filt t,16,500
23:01:12 08/09/10 ------------------Time (millisec)------------------
      ID     Port     8    16   32 64 128 256 512 1024 2048 4096 8192
       0    0:0:1  3647   550   29  1   0   0   0    0    0    0    0
       1    1:0:1  2742   969  111  7   0   0   0    0    0    0    0
       3    1:0:1  4181  2786  292  0   0   0   0    0    0    0    0
       5    1:0:1  3397  1818  228  3   0   0   0    0    0    0    0
       7    1:0:1  2793  1281   77  0   1   0   0    0    0    0    0
       8    0:0:1  1449   801   23  0   0   0   0    0    0    0    0
       9    1:0:1  2425  2779  287  3   0   0   0    0    0    0    0
      11    1:0:1  3526   860  157  1   1   0   0    0    0    0    0
      12    0:0:1  6049  2171   58  3   0   0  13    6    0    0    0
      13    1:0:1  4444  2942  261  9   0   0   0    0    0    0    0
      16    0:0:2  3630   555   35  0   0   0   0    0    0    0    0
      17    1:0:2  2769   958  105  0   0   0   0    0    0    0    0
      19    1:0:2  4227  2773  263  0   0   0   0    0    0    0    0
      21    1:0:2  2757  2353  339  4   0   0   0    0    0    0    0
      23    1:0:2  2735  1321   94  2   0   0   0    0    0    0    0
```

```
      24    0:0:2  1360   577   15  2   0   2   3    0    0    0    0
      25    1:0:2  2484  2420  286  0   0   0   0    0    0    0    0
      27    1:0:2  4504   633  122  1   0   0   0    0    0    0    0
      28    0:0:2  6059  2177   65  3   0   1   1    2    0    0    0
      29    1:0:2  4245  3113  300  1   0   0   0    0    0    0    0
------------------------------------------------------------------------
   total          69423 33837 3147 40   2   3  17    8    0    0    0
Press the enter key to stop...
```

histport

Syntax

`histport [options <arg>]`

Description

The `histport` command displays a histogram of service times for ports within the system.

Options

`-both|-ctl|-data`

Specifies that both control and data transfers are displayed (`-both`), only control transfers are displayed (`-ctl`), or only data transfers are displayed (`-data`). If this option is not specified, only data transfers are displayed.

`-nodes <node_list>`

Specifies that the display is limited to specified nodes and physical disks connected to those nodes. The node list is specified as a series of integers separated by commas (e.g. 1,2,3). The list can also consist of a single integer. If the node list is not specified, all disks on all nodes are displayed.

`-slots <slot_list>`

Specifies that the display is limited to specified PCI slots and physical disks connected to those PCI slots. The slot list is specified as a series of integers separated by commas (e.g. 1,2,3). The list can also consist of a single integer. If the slot list is not specified, all disks on all slots are displayed.

`-ports <port_list>`

Specifies that the display is limited to specified ports and physical disks connected to those ports. The port list is specified as a series of integers separated by commas (e.g. 1,2,3). The list can also consist of a single integer. If the port list is not specified, all disks on all ports are displayed.

`-host|-disk|-rcfc|-peer`

Specifies to display only host ports (target ports), only disk ports (initiator ports), only Fibre Channel Remote Copy configured ports, or only Fibre Channel ports for Data Migration.

If no option is specified, all ports are displayed.

`-metric both|time|size`

Selects which metric to display. Metrics can be one of the following:

- both—Display both I/O time and I/O size histograms (default).
- time—Display only the I/O time histogram.
- size—Display only the I/O size histogram.

`-timecols <fcol> <lcol>`

For the I/O time histogram, shows the columns from the first column `<fcol>` through last column `<lcol>`. The available columns range from 0 through 31.

The first column (<fcol>) must be a value greater than or equal to 0, but less than the value of the last column (<lcol>).

The last column (<lcol>) must be less than or equal to 31.

The first column includes all data accumulated for columns less than the first column and the last column includes accumulated data for all columns greater than the last column.

The default value of <fcol> is 6.

The default value of <lcol> is 15.

-sizecols <fcol> <lcol>

For the I/O size histogram, shows the columns from the first column (<fcol>) through the last column (<lcol>). Available columns range from 0 through 15.

The first column (<fcol>) must be a value greater than or equal to 0, but less than the value of the last column (<lcol>) (default value of 3).

The last column (<lcol>) must be less than or equal to 15 (default value of 11).

The default value of <fcol> is 3.

The default value of <lcol> is 11.

-pct

Shows the access count in each bucket as a percentage. If this option is not specified, the histogram shows the access counts.

-prev|-begin

Histogram displays data either from a previous sample (-prev) or from when the system was last started (-begin). If no option is specified, the histogram shows data from the beginning of the command's execution.

-rw

Specifies that the display includes separate read and write data. If not specified, the total is displayed.

-d <secs>

Specifies the interval in seconds that statistics are sampled from using an integer from 1 through 2147483. If no count is specified, the command defaults to 2 seconds.

-iter <number>

Specifies that the histogram is to stop after the indicated number of iterations using an integer from 1 through 2147483647.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

-filt <fspec>

Specifies that histograms below the threshold specified by the <fspec> argument are not displayed. The <fspec> argument is specified in the syntax of <op>,<val_ms>, <count>, where:

- <op>

  The <op> argument can be specified as one of the following:

  - r—Specifies read statistics.
  - w—Specifies write statistics.
  - t—Specifies total statistics.
  - rw—Specifies total read and write statistics.

<val_ms>

Specifies the threshold service time in milliseconds.

<count>

Specifies the minimum number of access above the threshold service time. When filtering is done, the <count> is compared with the sum of all columns starting with the one which corresponds to the threshold service time. For example, -t,8,100 means to only display the rows where the 8ms column and all columns to the right adds up to more than 100.

-ni

Specifies that histograms for only non-idle devices are displayed. This option is shorthand for the option -filt t,0,0.

Specifiers

None.

Authority

Any role in the system

Usage

- Each service time histogram column shows the number of accesses with service times between the point in time shown in the column's heading to the point in time shown in the heading of the column to its right (or infinity for the last column). Unless all columns are shown, the first column shows the number of accesses with service times less than the heading of the second column.
- Each I/O size histogram column shows the number of accesses with I/O size between the size shown in the column's heading to the size shown in the heading of the column to its right (or infinity for the last column).
- For all the <name_or_pattern> fields, the patterns are glob-style (shell-style) patterns. (For more information, enter clihelp sub,globpat or see **Glob-style patterns**).

Examples

The following example displays a histogram of service times for reads and writes to ports:

```
cli% histport
12:35:24 10/20/04 ------------Time (millisec)-------------
    Port      D/C 0.26 0.53 1.05 2.1 4.2 8.4  17 34 67 135
   0:0:1     data    0    0    0   0   0   0   0  0  0   0
   0:0:2     data  315  778    2   0   0   0   0  0  0   0
   0:1:1     data    0    0    0   5  24  51  25  4  0   0
   0:1:2     data    0    0    0   5  27  53  23  1  0   0
```

```
   1:5:1     data    0    0    0   2  19  38  28 11  0   0
   1:5:2     data    0    0    0   5  20  36  29  7  0   0
------------------------------------------------------------
   total     data  315  778    2  17  90 178 105 23  0   0
Press the enter key to stop...
```

histrcvv

Syntax

```
histrcvv [options <arg>] [<VV_name>|<pattern>]...
```

Description

The `histrcvv` command shows a histogram of total remote copy service times and backup remote copy service times in a timed loop.

Options

`-d <secs>`

Specifies the interval in seconds that statistics are sampled from using an integer from 1 through 2147483. If no count is specified, the command defaults to 2 seconds.

`-iter <number>`

Specifies that the statistics are to stop after the indicated number of iterations using an integer from 1 through 2147483647.

`-t <target_name>...|<pattern>...`

Shows only volumes whose group is copied to the specified target name or pattern. Multiple target names or patterns may be specified using a comma-separated list.

`-port <port_NSP>...|<pattern>...`

Shows only volumes which are copied over the specified port or pattern.

Multiple ports or patterns may be specified using a comma-separated list.

`-g <group_name>...|<pattern>...`

Shows only volumes whose group matches the specified group name or pattern. Multiple group names or patterns may be specified using a comma-separated list.

`-async`

Shows only volumes which are being copied in asynchronous mode.

`-sync`

Shows only volumes which are being copied in synchronous mode.

`-periodic`

Shows only volumes which are being copied in periodic mode.

`-primary`

Shows only volumes which are in the primary role.

`-secondary`

Shows only volumes which are in the secondary role.

`-targetsum`

Specifies that the sums for all volumes of a target are displayed.

-portsum

Specifies that the sums for all volumes on a port are displayed.

-groupsum

Specifies that the sums for all volumes of a group are displayed.

-vvsum

Specifies that the sums for all targets and links of a volume are displayed.

-domainsum

Specifies that the sums for all volumes of a domain are displayed.

-domain <domainname>...|<pattern>...

Shows only the virtual volumes that are in domains with names that match the specified domain name(s) or pattern(s).

-timecols <fcol> <lcol>

For the I/O time histogram, shows the columns from the first column <fcol> through last column <lcol>. The available columns range from 0 through 31.

0 <= fcol < lcol <= 31 must be satisfied.

The default value of <fcol> is 6.

The default value of <lcol> is 15.

-pct

Shows the access count in each bucket as a percentage. If this option is not specified, the histogram shows the access counts.

-prev

Specifies that the histogram displays data from a previous sample.

If no option is specified, the histogram shows data from the beginning of the command's execution.

-d <secs>

Specifies the interval in seconds that statistics are sampled from using an integer from 1 through 2147483. If no count is specified, the command defaults to 2 seconds.

-iter <number>

Specifies that the histogram is to stop after the indicated number of iterations using an integer from 1 through 2147483647.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

-filt <fspec>

Specifies that histograms below the threshold specified by the <fspec> argument are not displayed. The <fspec> argument is specified in the syntax of <op>,<val_ms>, <count>, where:

- <op>

  The <op> argument can be specified as one of the following:

  - r—Specifies read statistics.
  - w—Specifies write statistics.
  - t—Specifies total statistics.
  - rw—Specifies total read and write statistics.

<val_ms>

Specifies the threshold service time in milliseconds.

<count>

Specifies the minimum number of access above the threshold service time. When filtering is done, the <count> is compared with the sum of all columns starting with the one which corresponds to the threshold service time. For example, -t,8,100 means to only display the rows where the 8ms column and all columns to the right adds up to more than 100.

-ni

Specifies that histograms for only non-idle devices are displayed. This option is shorthand for the option -filt t,0,0.

Specifiers

<VV_name|pattern>

Displays statistics only for the specified virtual volume or volume name pattern. Multiple volumes or patterns can be repeated (for example, <VV_name> <VV_name>). If not specified, all virtual volumes that are configured for remote copy are listed.

Authority

Any role in the system

Usage

This command requires the HPE 3PAR Remote Copy Software license. Contact your local service provider for more information.

Examples

The following example shows histograms sums for all synchronous Remote Copy volumes:

```
cli% histrcvv -vvsum -sync
           09:50:40 03/30/10 ----------Svt (millisec)---------- -------Rmt Svt (millisec)--------
  VVname RCGroup Target Mode 0.50 1 2   4   8  16 32 64 128 256 0.50    1  2 4 8 16 32 64 128 256
testvv.2 multi.1   amp2 Sync    0 0 0 163  61  16  3  0   0   0    0  224 19 0 0  0  0  0   0   0
testvv.3 multi.1   amp2 Sync    0 0 0 138  62  37  4  2   0   0    0  224 19 0 0  0  0  0   0   0
testvv.8  sync.2   amp3 Sync    0 0 0   0   0   4 13 11   5   4    0    0  0 0 4 12 12  7   1   1
testvv.4 multi.1   amp2 Sync    0 0 0  54 132  52 17  2   0   0    0  252  5 0 0  0  0  0   0   0
testvv.0 multi.1   amp2 Sync    0 0 0  59 138  57 17  2   0   0    0  258 15 0 0  0  0  0   0   0
testvv.9  sync.2   amp3 Sync    0 0 0   0   1   0 14  8   6   3    0    0  0 0 2  7 14  4   4   1
testvv.1 multi.1   amp2 Sync    0 0 0 130  71  40  3  0   0   0    0  225 19 0 0  0  0  0   0   0
------------------------------------------------------------------------------------------------
   total                       0 0 0 544 465 206 71 25  11   7    0 1183 77 0 6 19 26 11   5   2
Press the enter key to stop...
```

Use the following to get a description of each column in the above examples:

```
clihelp -col histrcvv
```

histvlun

Syntax

```
histvlun [options <arg>]
```

Description

The `histvlun` command displays VLUN service time histograms.

Options

`-domain <domain_name>...|<pattern>...`

Shows only VLUNs whose virtual volumes are in domains with names that match one or more of the specified domain names or patterns. Multiple domain names or patterns can be repeated using a comma-separated list.

`-host <hostname>...|<pattern>...`

Shows only VLUNs exported to the specified host(s) or pattern(s).

Multiple host names or patterns can be repeated using a comma-separated list.

`-v <VV_name>...|<pattern>...`

Requests that only logical disks mapped to virtual volumes that match any of the specified names or patterns be displayed. Multiple volume names or patterns can be repeated using a comma-separated list.

`-l <LUN>...|<pattern>...`

Specifies that VLUNs with LUNs matching the specified LUN(s) or pattern(s) are displayed. Multiple LUNs or patterns can be repeated using a comma-separated list.

`-nodes <node_list>`

Specifies that the display is limited to specified nodes and physical disks connected to those nodes. The node list is specified as a series of integers separated by commas (e.g. 1,2,3). The list can also consist of a single integer. If the node list is not specified, all disks on all nodes are displayed.

`-slots <slot_list>`

Specifies that the display is limited to specified PCI slots and physical disks connected to those PCI slots. The slot list is specified as a series of integers separated by commas (e.g. 1,2,3). The list can also consist of a single integer. If the slot list is not specified, all disks on all slots are displayed.

`-ports <port_list>`

Specifies that the display is limited to specified ports and physical disks connected to those ports. The port list is specified as a series of integers separated by commas (e.g. 1,2,3). The list can also consist of a single integer. If the port list is not specified, all disks on all ports are displayed.

`-lw`

Lists the host's WWN or iSCSI name. This is especially useful when multiple WWNs or iSCSI names belonging to the same host are visible on the same port.

`-domainsum`

Specifies that sums for VLUNs are grouped by domain in the display. All VLUNs to unnamed hosts are added and displayed as a single set of data with a "-" host name.

`-vvsum`

Specifies that sums for VLUNs of the same virtual volume are displayed.

-hostsum

Specifies that sums for VLUNs are grouped by host in the display. All VLUNs to unnamed hosts are added and displayed as a single set of data with a nameless host.

-metric both|time|size

Selects which metric to display. Metrics can be one of the following:

- both—Display both I/O time and I/O size histograms (default).
- time—Display only the I/O time histogram.
- size—Display only the I/O size histogram.

-timecols <fcol> <lcol>

For the I/O time histogram, shows the columns from the first column <fcol> through last column <lcol>. The available columns range from 0 through 31.

The first column (<fcol>) must be a value greater than or equal to 0, but less than the value of the last column (<lcol>).

The last column (<lcol>) must be less than or equal to 31.

The first column includes all data accumulated for columns less than the first column and the last column includes accumulated data for all columns greater than the last column.

The default value of <fcol> is 6.

The default value of <lcol> is 15.

-sizecols <fcol> <lcol>

For the I/O size histogram, shows the columns from the first column (<fcol>) through the last column (<lcol>). Available columns range from 0 through 15.

The first column (<fcol>) must be a value greater than or equal to 0, but less than the value of the last column (<lcol>) (default value of 3).

The last column (<lcol>) must be less than or equal to 15 (default value of 11).

The default value of <fcol> is 3.

The default value of <lcol> is 11.

-pct

Shows the access count in each bucket as a percentage. If this option is not specified, the histogram shows the access counts.

-prev | -begin

Histogram displays data either from a previous sample (-prev) or from when the system was last started (-begin). If no option is specified, the histogram shows data from the beginning of the command's execution.

-rw

Specifies that the display includes separate read and write data. If not specified, the total is displayed.

-d <secs>

Specifies the interval in seconds that statistics are sampled from using an integer from 1 through 2147483. If no count is specified, the command defaults to 2 seconds.

-iter <number>

Specifies that the histogram is to stop after the indicated number of iterations using an integer from 1 through 2147483647.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

-filt <fspec>

Specifies that histograms below the threshold specified by the <fspec> argument are not displayed. The <fspec> argument is specified in the syntax of <op>,<val_ms>, <count>, where:

- <op>

  The <op> argument can be specified as one of the following:

  - r—Specifies read statistics.
  - w—Specifies write statistics.
  - t—Specifies total statistics.
  - rw—Specifies total read and write statistics.

<val_ms>

Specifies the threshold service time in milliseconds.

<count>

Specifies the minimum number of access above the threshold service time. When filtering is done, the <count> is compared with the sum of all columns starting with the one which corresponds to the threshold service time. For example, -t,8,100 means to only display the rows where the 8ms column and all columns to the right adds up to more than 100.

-ni

Specifies that histograms for only non-idle devices are displayed. This option is shorthand for the option -filt t,0,0.

The VV filtering option includes:

-p <pattern>

Patterns for matching VVs to show. (see below for description of <pattern>) If the -p option is specified multiple times, each instance of <pattern> adds additional candidate VVs that match that pattern.

A <pattern> is one or more of the following:

-vmname {<vmname>|<pattern>}[,{<vmname>|<pattern>}...]

Shows only VVs whose vmname matches one or more of the vvname patterns.

-vmid <vmid>[,<vmid>...]

Shows only VVs whose vmid matches one or more of the vmids.

-vmhost {<vmhost>|<pattern>}[,{<vmhost>|<pattern>}...]

Shows only VVs whose vmhost matches one or more of the vmhost patterns.

-vvolstate {bound|unbound}

Shows only VVs whose vvolstate matches the specified state - bound or unbound.

-vvolsc {<vvset>|<pattern>}[,{<vvset>|<pattern>}...]

Shows only VVs whose storage container (vvset) name matches one or more of the vvset patterns.

Specifiers

None.

Authority

Any role in the system

Usage

- Each service time histogram column shows the number of accesses with service times between the point in time shown in the column's heading to the point in time shown in the heading of the column to its right (or infinity for the last column). Unless all columns are shown, the first column shows the number of accesses with service times less than the heading of the second column.
- Each I/O size histogram column shows the number of accesses with I/O size between the size shown in the column's heading to the size shown in the heading of the column to its right (or infinity for the last column).
- For all the `<name_or_pattern>` fields, the patterns are glob-style (shell-style) patterns. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**).

Examples

The following example displays two iterations of a histogram of service times for all VLUNs:

```
cli% histvlun -iter 2
```

histvv

Syntax

```
histvv [options <arg>] [<VV_name>|<pattern>]...
```

Description

The `histvv` command displays virtual volume service time histograms in a timed loop.

Options

-domain <domainname>...|<pattern>...

Shows only the virtual volumes that are in domains with names that match the specified domain name(s) or pattern(s).

-metric both|time|size

Selects which metric to display. Metrics can be one of the following:

- both—Display both I/O time and I/O size histograms (default).
- time—Display only the I/O time histogram.
- size—Display only the I/O size histogram.

-timecols <fcol> <lcol>

For the I/O time histogram, shows the columns from the first column <fcol> through last column <lcol>. The available columns range from 0 through 31.

The first column (<fcol>) must be a value greater than or equal to 0, but less than the value of the last column (<lcol>).

The last column (<lcol>) must be less than or equal to 31.

The first column includes all data accumulated for columns less than the first column and the last column includes accumulated data for all columns greater than the last column.

The default value of <fcol> is 6.

The default value of <lcol> is 15.

-sizecols <fcol> <lcol>

For the I/O size histogram, shows the columns from the first column (<fcol>) through the last column (<lcol>). Available columns range from 0 through 15.

The first column (<fcol>) must be a value greater than or equal to 0, but less than the value of the last column (<lcol>) (default value of 3).

The last column (<lcol>) must be less than or equal to 15 (default value of 11).

The default value of <fcol> is 3.

The default value of <lcol> is 11.

-pct

Shows the access count in each bucket as a percentage. If this option is not specified, the histogram shows the access counts.

-prev

Specifies that the histogram displays data from a previous sample.

If no option is specified, the histogram shows data from the beginning of the command's execution.

-rw

Specifies that the display includes separate read and write data. If not specified, the total is displayed.

-d <secs>

Specifies the interval in seconds that statistics are sampled from using an integer from 1 through 2147483. If no count is specified, the command defaults to 2 seconds.

-iter <number>

Specifies that the histogram is to stop after the indicated number of iterations using an integer from 1 through 2147483647.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

-filt <fspec>

Specifies that histograms below the threshold specified by the <fspec> argument are not displayed. The <fspec> argument is specified in the syntax of <op>,<val_ms>, <count>, where:

- <op>

  The <op> argument can be specified as one of the following:

  - r—Specifies read statistics.
  - w—Specifies write statistics.
  - t—Specifies total statistics.
  - rw—Specifies total read and write statistics.

<val_ms>

Specifies the threshold service time in milliseconds.

<count>

Specifies the minimum number of access above the threshold service time. When filtering is done, the <count> is compared with the sum of all columns starting with the one which corresponds to the threshold service time. For example, -t,8,100 means to only display the rows where the 8ms column and all columns to the right adds up to more than 100.

-ni

Specifies that histograms for only non-idle devices are displayed. This option is shorthand for the option -filt t,0,0.

The VV filtering option include:

-p <pattern>

Patterns for matching VVs to show (see below for description of <pattern>). If the -p option is specified multiple times, each instance of <pattern> adds additional candidate VVs that match that pattern.

A <pattern> is one or more of the following:

-vmname {<vmname>|<pattern>}[,{<vmname>|<pattern>}...]

Shows only VVs whose vmname matches one or more of the vvname patterns.

-vmid <vmid>[,<vmid>...]

Shows only VVs whose vmid matches one or more of the vmids.

-vmhost {<vmhost>|<pattern>}[,{<vmhost>|<pattern>}...]

Shows only VVs whose vmhost matches one or more of the vmhost patterns.

-vvolstate {bound|unbound}

Shows only VVs whose vvolstate matches the specified state — bound or unbound.

-vvolsc {<vvset>|<pattern>}[,{<vvset>|<pattern>}...]

Shows only VVs whose storage container (vvset) name matches one or more of the vvset patterns.

Specifiers

None.

Authority

Any role in the system

Usage

- Each service time histogram column shows the number of accesses with service times between the point in time shown in the column's heading to the point in time shown in the heading of the column to its right (or infinity for the last column). Unless all columns are shown, the first column shows the number of accesses with service times less than the heading of the second column.
- Each I/O size histogram column shows the number of accesses with I/O size between the size shown in the column's heading to the size shown in the heading of the column to its right (or infinity for the last column).

  For all the `<name_or_pattern>` fields, the patterns are glob-style (shell-style) patterns. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.)
- If a `<VV_name>` or `<pattern>` are specified, then virtual volumes with names matching any of the patterns are listed. Otherwise all virtual volumes are listed. These patterns are glob-style patterns. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.)
- Virtual volumes may be accessed externally by hosts and internally by the prefetcher. Virtual volume data measured by this command include accesses by the prefetcher.
- In addition to external accesses by hosts, virtual volumes can be read internally by the system read-ahead prefetcher. The `histvv` data includes read-ahead accesses from the prefetcher that can cause the read data to appear more than seen by the hosts. Use the `histvlun -vvsum` command to see data for only accesses from the host.
- Volumes for VMware virtual machines (vVols), which are hidden by default, may be included by using the **`setclienv`** command to set matchbulkobjs to "1".
- Additionally, vVols can displayed if the CLI was started with the `-matchbulkobjs` option, or if the CLI was started with the TPDMATCHBULKOBJS environment variable set to "1". Run "`cli -h`" and "`setclienv -h`" for details of the environment variables. In addition, using the pattern matching options `-vmname`, `-vmid`, `-vmhost`, `-vvolstate`, or `-vvolsc` will display vVols.

Examples

The following example displays two iterations of a histogram of service times for all virtual volumes:

```
%cli histvv -iter 2
12:53:03 10/20/14 ------------Time (millisec)------------
VVname 0.26 0.53 1.05 2.1 4.2 8.4 17 34 67 135
admin     0    0    0   0   0   0  0  0  0   0
tpvv     29    0    0   0   0   0  0  0  0   0
-----------------------------------------------------------
total    29    0    0   0   0   0  0  0  0   0

12:53:05 10/20/14 ------------Time (millisec)------------
VVname 0.26 0.53 1.05 2.1 4.2 8.4 17 34 67 135
admin     0    0    0   0   0   0  0  0  0   0
tpvv    105    0    0   0   0   0  0  0  0   0
-----------------------------------------------------------
total   105    0    0   0   0   0  0  0  0   0
```

Import Commands

importcert

Syntax

```
importcert <SSL_service> [-f] <service_cert> [<CA_bundle>]
importcert <SSL_service> [-f] -ca <CA_bundle>
```

Description

The `importcert` command allows a user to import certificates for a given service. The user can import a CA bundle containing the intermediate and/or root CAs prior to importing the service certificate. The CA bundle can also be imported alongside the service certificate.

Options

`-f`

Imports a certificate without prompting the user.

`-ca <CA_bundle>`

Allows the import of a CA bundle without importing a service certificate. Note the filename "stdin" can be used to paste the CA bundle into the CLI.

Specifiers

`<SSL_service>`

Valid service names are `cim`, `cli`, `ekm-client`, `ekm-server`, `ldap`, `qw-client`, `qw-server`, `syslog-gen-client`, `syslog-gen-server`, `syslog-sec-client`, `syslog-sec-server`, `wsapi`, `vasa`, and `unified-server`.

Authority

Super

Usage

- The `unified-server` establishes a common certificate among CIM, CLI, and WSAPI. Also, the CIM and WSAPI services are restarted when a self-signed certificate is generated.
- Data at rest encryption using an external key manager employs the `ekm-client` and `ekm-server` services. The external key manager services are not covered by the unified-server service.
- The filename "stdin" can be used to paste the CA bundle and or service certificate into the CLI.
- Use the **`createcert`** command to create a CSR and use the **`showcert`** command to display the certificates.

Examples

The following example shows how to import a signed service certificate with the supporting CA for the cli service:

```
cli% importcert cli cli-service.pem ca.pem
```

The following example shows how to import just the supporting CAs for the wsapi service without importing the service certificate itself:

```
cli% importcert wsapi -ca ca-bundle.pem
```

Now that the CA bundle has been imported, the service certificate can be imported:

```
cli% importcert wsapi wsapi-service.pem
```

importvv

Syntax

```
importvv [options] <usrcpg> <VV_name|pattern|VV_set> ...
```

Description

The `importvv` command starts migrating the data from a remote LUN to the local HPE 3PAR Storage system. The remote LUN should have been prepared using the **admitvv** command.

Options

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

`-snap <snapname>`

Creates a snapshot of the volume at the end of the import phase. While the data is getting imported the local volume and the remote volume are kept in sync (in other words all the data is written while data is getting imported is mirrored to the remote array). Once the volume import is complete data is only written to the volume on the array that was importing the data, and the volume on the old array will no longer be updated. The `-snap` parameter will create a snapshot at the last phase of the import and the content of the created snapshot will be identical to the state of volume on the old array when the import task is completed.

If multiple VVs are being imported in a single command, `<snapname>` should be a VV pattern. (For more information, issue "clihelp sub,vvnamepat".)

`-snp_cpg <snp_cpg>`

Specifies the name of the CPG from which the snapshot space will be allocated. This option needs to be specified if the `-snap` option is specified.

`-nocons`

Any VV sets specified will not be imported as consistent groups. Allows multiple VV sets to be specified.

If the VV set contains any VV members that in a previous import attempt were imported consistently, they will continue to get imported consistently.

`-pri <high|med|low|auto>`

Specifies the priority of migration of a volume or a volume set. If this option is not specified, the default priority will be auto. The volumes with priority set to high will migrate faster than other volumes with medium and low priority. If the priority is set to auto, then the array controls the speed of the operation based on the load in the array for all flash models. Auto priority behaves like medium priority for the converged flash models of the arrays containing Nearline (NL) and Fast Class (FC) drives.

`-jobid <Job_ID>`

Specifies the Job ID up to 511 characters for the volume. The Job ID will be tagged in the events that are posted during volume migration. Use `-jobid ""` to remove the Job ID.

`-notask`

Performs import related pre-processing which results in transitioning the volume to exclusive state and setting up of the "consistent" flag on the volume if importing consistently. The import task will not be created, and hence volume migration will not happen. The "`importvv`" command should be rerun on the volume at a later point of time without specifying the `-notask` option to initiate the actual migration of the volume. With the `-notask` option, other options (namely `-tpvv`, `-dedup`, `-snp_cpg`, `-snap`, `-clrsrc`, `-jobid` and `-pri`) cannot be specified.

-clrsrc

Performs cleanup on source array after successful migration of the volume. As part of the cleanup, any exports of the source volume will be removed, the source volume will be removed from all of the VV sets it is member of, the VV sets will be removed if the source volume is their only member, all of the snapshots of source volume will be removed, and finally the source volume itself will be removed. The `-clrsrc` option is valid only when the source array is running HPE 3PAR OS release 3.2.2 or higher. The cleanup will not be performed if the source volume has any snapshots that have VLUN exports.

The following options can be used to select the provisioning type of the imported volumes:

-tpvv

Imports the VV into a thinly provisioned space in the CPG specified in the command line. The import will enable zero detect for the duration of import so that the data blocks containing zero do not occupy space on the new array.

-tdvv

This option is deprecated, see -dedup

-dedup

Import the VV into a thinly provisioned space in the CPG specified in the command line. This volume will share logical disk space with other instances of this volume type created from the same CPG to store identical data blocks for space saving.

-compr

Import the VV into a compressed virtual volume in the CPG specified in the command line.

If no provisioning type is specified, the default is `-tpvv`.

Specifiers

<usrcpg>

Specifies the name of the CPG from which the volume user space will be allocated. If the `-tpvv` option is specified, the volume is thinly provisioned. Otherwise, the volume is fully provisioned.

<VV_name|pattern|VV_set> ...

Specifies the VVs with the specified name (up to 31 characters in length), matching the glob-style pattern or that are members of the supplied VV set will be imported onto local storage. The VV set name must start with "set:". This specifier can be repeated to import multiple VVs at the same time.

If a VV set is specified then by default all the VV members of the set will be imported consistently. Only one VV set can be specified to import consistently in a single invocation of the `importvv` command. The import parameters such as user and snapshot space CPGs, snapshot name, and provision type will by default be common to all the VV members of the set. To specify different values for VV members of the set, invoke the `importvv` command with the set name and specify individual VV members or patterns and their own user and snapshot space CPGs, snapshot name, and provision type (see Examples).

Note that the import of a set will not complete until imports have been initiated and completed for all of the VV members of the set.

To override the default behavior and to import VV members of the set individually, use the "`-nocons`" option. There is no limit on the number of VV sets or individual VVs specified with "`-nocons`" option.

Authority

- Super, Edit
- Any role granted the `vv_import` right

Usage

- You need access to all domains in order to run this command.
- The `importvv` command can be run manually, but it can also be scheduled to run automatically at a specific time using the `createsched -run_once` command.
- A volume is imported online but the volume can only be accessed using the storage system that is importing the volume.
- For all the `<name_or_pattern>` fields, the patterns are glob-style (shell-style) patterns. (For more information, issue `clihelp sub,globpat`, or see, **Glob-style patterns**.)
- The imported volumes can have different volume characteristics and is independent from legacy volume characteristics.

Examples

The following example imports vv1 previously admitted using the **admitvv** command. :

```
cli% importvv cpg1 vv1
```

The following example imports vv1 as a thin provisioned volume:

```
cli% importvv -tpvv cpg1 vv1
```

The following example imports multiple volumes as thin provisioned volumes and creates read only snapshots:

```
cli% importvv -tpvv -snap @vvname@.ro -snp_cpg cpg2 cpg1 vv*
```

The following example imports vv1 with high priority:

```
cli% importvv -pri high cpg1 vv1
```

The following example imports vv1 as a thin provisioned volume that shares logical disk space with other instances of the same volume type using the same CPG as vv1:

```
cli% importvv -dedup cpg1 vv1
```

The following example performs import related pre-processing including transitioning the volume to exclusive state, without creating an import task:

```
cli% importvv -notask cpg1 vv1
```

The following example cleans up volume vv1 on source array after successful migration of the volume:

```
cli% importvv -tpvv -clrsrc cpg1 vv1
```

The following example imports VV set vvset1 as a consistent group using thinly deduplicated and compressed provisioning:

```
cli% importvv -clrsrc cpg1 vv1
```

The following example imports VV set vvset1 with members vv1, vv2, and vv3 as a consistent group, with vv1 and vv2 imported using thin provisioning, and vv3 imported using full provisioning:

```
cli% importvv -tpvv cpg1 set:vvset1 vv1 vv2
cli% importvv cpg2 set:vvset1 vv3
```

The following example imports VV sets vvset1 and vvset2 without using consistent groups, and additionally imports vv7:

```
cli% importvv -nocons cpg1 set:vvset1 set:vvset2 vv7
```

Locate Commands

locatecage

Syntax

To locate an entire drive cage:

```
locatecage [option <arg>] <cage_name>
```

To locate a drive magazine:

```
locatecage [option <arg>] <cage_name> <mag>
```

To locate a port:

```
locatecage [option <arg>] <cage_name> <port_name>
```

To locate a module:

```
locatecage [option <arg>] <cage_name> <module_name> <module_number>
```

To locate a module:

```
locatecage [option <arg>] <cage_name> <module_name> <module_number>
```

Description

Allows system administrators to locate a drive cage, drive magazine, fan, power supply, io card, or port in the system using the locate LEDs of the device.

Options

`-t <sec>`

Specifies the length of time to blink the LEDs in seconds. For HPE 3PAR 7000 and HPE 3PAR 8000 storage systems, defaults to 15 minutes with a maximum time of one hour. For HPE 3PAR 20000 system, defaults to 1 hour with a maximum time of 7 days.

For all other systems, defaults to 60 seconds with a maximum time of 255 seconds.

Specifiers

`<cage_name>`

Specifies the drive cage name as shown in the **Name** column of `showcage` command output.

`<port_name>`

Indicates the port specifiers. Accepted values are A0|B0|A1|B1|A2|B2|A3|B3. The <port_name> specifier is not supported for `DC3, DCS1, DCS2, DCS3, DCS4, DCS5, DCS6, DCS9, DCS10, DCN1, DCN3` and `DCN4` drive cages.

`<mag>`

Indicates the drive magazine by number.

For DC1 drive cages, accepted values are 0 through 4.

For DC2 and DC4 drive cages, accepted values are 0 through 9.

For DC3 drive cages, accepted values are 0 through 15.

For DCS3, DCS5, DCS9 and DCN3 drive cages, accepted values are 0 through 11.

For DCS4, DCS6, DCS10 and DCN4 drive cages, accepted values are 0 through 23.

<module_name>

Indicates the module name to locate. Accepted values are `enclosure|fan|powersupply|battery|iocard|disk|magazine`. The `<module_name>` specifier is only supported for DCS3, DCS4, DCS5, DCS6, DCS7, DCS8, DCS9, DCS10, DCN2, DCN3 and DCN4 cages.

<module_number>

Indicates the module number to locate. The `<module_number>` specifier is only supported for DCS3, DCS4, DCS5, DCS6, DCS7, DCS8, DCS9, DCS10, DCN2, DCN3 and DCN4 cages.

The allowed module numbers per module type are:

0 for enclosure.

0 .. 1 for fan (0 .. 3 on DCS7).

0 .. 1 for powersupply (0 .. 3 on DCS7).

0 .. 1 for battery (0 .. 3 on DCS7).

0 .. 1 for iocard.

0 .. 23 for disk or magazine.

Authority

- Super, Service
- Any role granted the `cage_locate` right

Usage

- Requires access to all domains.
- To list cage names, use the **showcage** command.

Examples

Locate `cage0` for 20 seconds:

```
cli% locatecage -t 20 cage0
```

locatenode

Syntax

```
locatenode [options] <nodeid>
```

Description

The `locatenode` command helps locate a particular node or its components by illuminating LEDs on the node.

Options

`-t <sec>`

Specifies the number of seconds to illuminate the LEDs. For HPE 3PAR 7000 and HPE 3PAR 8000 storage systems, the default time to illuminate the LEDs is 15 minutes with a maximum time of one hour. For STR (Safe to Remove) systems, the default time is one hour with a maximum time of one week. For all other systems, the default time is 60 seconds with a maximum time of 255 seconds. Issuing "locatenode -t 0 <nodeid>" will turn off LEDs immediately.

-ps <psid>

Only the service LED for the specified power supply will blink.

Accepted values for <psid> are 0 and 1.

-pci <slot>

Only the service LED corresponding to the PCI card in the specified slot will blink. Accepted values for <slot> are 3 through 5.

-fan <fanid>

Only the service LED on the specified node fan module will blink.

Accepted values for <fanid> are 0 and 1 for HPE 3PAR 10000 systems.

Accepted values for <fanid> are 0, 1 and 2 for HPE 3PAR 20000 systems.

-drive

Only the service LED corresponding to the node's internal drive will blink.

-bat

Only the service LED on the battery backup unit will blink.

Specifiers

<nodeid>

Indicates which node the `locatenode` operation will act on. Accepted values are 0 through 7.

Authority

- Super, Service
- Any role granted the `node_locate` right

Usage

- Access to all domains is required to run this command.
- On HPE 3PAR 7000 and HPE 3PAR 8000 series systems only the node LED can be illuminated.

Examples

In the following example, power supply 1 on node 3 is identified by blinking its hotplug LED for 90 seconds.

```
cli% locatenode -t 90 -ps 1 3
```

locatesys

Syntax

```
locatesys [options <arg>]
```

Description

The `locatesys` command helps locate a storage system by blinking the node status LEDs on all nodes of the storage system alternating amber and green. By default, the LEDs in all connected cages are also set amber or oscillate (depending on the cage).

Options

-t <sec>

Specifies the number of seconds to blink the LEDs. For HPE 3PAR 7000 and HPE 3PAR 8000 storage systems, the default time to blink the LEDs is 15 minutes with a maximum time of one hour. For all other systems, the default time to blink the LEDs is 60 seconds with a maximum time of 255 seconds.

-nodes <nodelist>

Specifies a comma-separated list of nodes on which to blink LEDs. The default is all nodes.

-nocage

Specifies that LEDs on the drive cages should not blink. The default is to blink LEDs for all cages in the system.

Specifiers

None.

Authority

- Super, Service
- Any role granted the `sys_locate` right

Usage

Access to all domains is required to run this command.

Examples

In the following example, a storage system is identified by blinking the LEDs on all drive cages in the system for 90 seconds:

```
cli% locatesys -t 90
```

Move Commands

movech

Syntax

```
movech [options] <fd:fp-td:tp> ...
```

Description

The `movech` command moves a list of chunklets from one physical disk to another.

Options

`-nowait`

Specifies that the command returns before the operation is completed.

`-dr`

Specifies that the operation is a dry run. No chunklets are actually moved.

`-devtype`

Permits the moves to happen to different device types.

`-perm`

Specifies that chunklets are permanently moved and the chunklets' original locations are not remembered. If the -perm option is not specified, the chunklets' original locations are retained, thereby allowing the chunklets to be returned to their original locations through the `moverelocpd` and `servicemag` resume commands.

If this option is used on a non-spare chunklet and it is moved to itself, the chunklet's original location will be cleared without actually moving the data. The same thing will occur, if the destination is left blank and the system chooses a destination on the same disk.

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

`-ovrd`

Permits the moves to happen to a destination even when there will be a loss of quality because of the move. This option is only necessary when the target of the move is not specified and the -perm flag is used.

Specifiers

`<fd:fp>[<-td:tp>]...`

Specifies that the chunklet located at the specified disk (<fd>) and the chunklet's position on that disk (<fp>) be moved to either the specified destination disk (<td>) and chunklet position (<tp>), or a location determined by the system if a destination (<-td:tp>) is not specified. This specifier must be used at least once on the command line. Repeated use of this specifier allows multiple chunklets to be moved.

Authority

- Super, Service, Edit
- Any role granted the `ch_move` right

Usage

- Access to all domains is required to run this command.
- Chunklets moved through the `movech` command are only moved temporarily.
- Issuing either the **`moverelocpd`** or `servicemag resume` command (see the **`servicemag`** command) can move the chunklet back to its original position.
- The -dr option can be used to see if the specified moves succeed and what the results (quality) of the moves are.

Examples

The following example moves the chunklet in position 0 on disk 24, to position 50 on disk 64 and chunklet in position 0 on disk 25, to position 1 on disk 27.

```
cli% movech 24:0-64:50 25:0-27:1
Are you sure you want to move the chunklets?
select q=quit y=yes n=no: y
Move       -State- --Detailed_State---
24:0-64:50 normal   normal
25:0-27:1  degraded disks_reused_in_row
```

The following example doesn't move the chunklet in position 0 on disk 10, to itself, but clears the chunklet's original location.

```
cli% movech -perm 10:0-10:0
Are you sure you want to move the chunklets?
select q=quit y=yes n=no: y
Move        -State-  -Detailed_State-
10:0-10:0   normal    no_data_move,normal
```

Use the following to get a description of each column in the above examples:

```
clihelp -col movech
```

movechtospare

Syntax

```
movechtospare [options] <fd:fp>
```

Description

The `movechtospare` command moves data from specified physical disks to a temporary location selected by the system.

Options

`-nowait`

Specifies that the command returns before the operation is completed.

`-dr`

Specifies that the operation is a dry run. No chunklets are actually moved.

`-devtype`

Permits the moves to happen to different device types.

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

Specifiers

<fd:fp>

Indicates that the move takes place from the specified physical disk (<fd>) and chunklet position (<fp>).

Authority

- Super, Service, Edit
- Any role granted the `ch_movetospare` right

Usage

Access to all domains is required to run this command.

Examples

In the following example, chunklet 0 from physical disk 66 is moved to spare:

```
cli% movechtospare 66:0
Are you sure you want to move the chunklet to spare?
select q=quit y=yes n=no: y
Move       -State- -Detailed_State-
66:0-70:50 normal   normal
```

Use the following to get a description of each column in the above examples:

```
clihelp -col movechtospare
```

movepd

Syntax

```
movepd [options] <PD_ID>...
```

Description

The `movepd` command moves data from specified physical disks to a temporary location selected by the system.

Options

-nowait

Specifies that the command returns before the operation is completed.

-dr

Specifies that the operation is a dry run, and no physical disks are actually moved.

-devtype

Permits the moves to happen to different device types.

-perm

Makes the moves permanent, removes source tags after relocation.

-f

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

Specifiers

<PD_ID>...

Specifies the physical disk ID. This specifier can be repeated to move multiple physical disks.

Authority

- Super, Service, Edit
- Any role granted the `pd_move` right

Usage

- Access to all domains is required to run this command.
- The destination physical disks do not need to be specified as the system automatically determines the spare locations.
- Specifying the -dr option can be used to see if the specified moves succeed and the results (quality) of the moves.

Examples

The following example displays a dry run of moving the data on physical disk 0 to free or spare space:

```
cli% movepd -dr 0
Are you sure you want to move the pd to spare?
select q=quit y=yes n=no: y
Move        -State- -Detailed_State-
0:0-66:50   normal   normal
0:1-62:51   normal   normal
0:2-92:51   normal   normal
0:3-88:52   normal   normal
0:4-50:53   normal   normal
0:5-86:54   normal   normal
0:6-82:55   normal   normal
```

Use the following to get a description of each column in the above examples:

```
clihelp -col movepd
```

movepdtospare

Syntax

```
movepdtospare [options] <PD_ID>...
```

Description

The `movepdtospare` command moves data from specified physical disks to a temporary location selected by the system.

Options

`-nowait`

Specifies that the command returns before the operation is completed.

`-dr`

Specifies that the operation is a dry run. No physical disks are moved.

`-devtype`

Permits the moves to happen to different device types.

`-vacate`

Deprecates, use -perm instead.

`-perm`

Makes the moves permanent, removes source tags after relocation.

-ovrd

Permits the moves to happen to a destination even when there will be a loss of quality because of the move. This option is only necessary when the target of the move is not specified and the -perm flag is used.

-f

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

Specifiers

<PD_ID>...

Specifies the physical disk ID. This specifier can be repeated to move multiple physical disks.

Authority

- Super, Service, Edit
- Any role granted the `pd_movetospare` right

Usage

- Access to all domains is required to run this command.
- The destination physical disks do not need to be specified as the system automatically determines the spare locations.
- Specifying the -dr option can be used to see if the specified moves succeed and the results (quality) of the moves.

Examples

The following example displays a dry run of moving the data on physical disk 0 to free or spare space:

```
cli% movepdtospare -dr 0
Are you sure you want to move the pd to spare?
select q=quit y=yes n=no: y
Move        -State- -Detailed_State-
0:0-66:50   normal   normal
0:1-62:51   normal   normal
0:2-92:51   normal   normal
0:3-88:52   normal   normal
0:4-50:53   normal   normal
0:5-86:54   normal   normal
0:6-82:55   normal   normal
```

Use the following to get a description of each column in the above examples:

```
clihelp -col movepdtospare
```

moverelocpd

Syntax

```
moverelocpd [options] [<fd>[-<td>]...]
```

Description

The `moverelocpd` command moves chunklets that were on a physical disk to the target of relocation.

Options

-nowait

Specifies that the command returns before the operation is completed.

-dr

Specifies that the operation is a dry run. No physical disks are moved.

-partial

Moves as many chunklets as possible. If this option is not specified, the command fails if not all specified chunklets can be moved.

-f

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

-p <pattern>

Specifies a pattern to select <fd> disks. The <td> specifier cannot be used with this -p option. If no `<fd>` are specified, then all disks that match the pattern are selected as the <fd> disks. If `<fd>` disks are specified along with `-p`, then only disks that match the pattern are selected as `<fd>` disk.

The following arguments can be specified as patterns for this option:

An item is specified as an integer, a comma-separated list of integers, or a range of integers specified from low to high.

-nd <item>

Specifies one or more nodes. Nodes are identified by one or more integers (item). Multiple nodes are separated with a single comma (e.g. 1,2,3). A range of nodes is separated with a hyphen (e.g. 0-7). The primary path of the disks must be on the specified node(s).

-st <item>

Specifies one or more PCI slots. Slots are identified by one or more integers (item). Multiple slots are separated with a single comma (e.g. 1,2,3). A range of slots is separated with a hyphen (e.g. 0-7). The primary path of the disks must be on the specified PCI slot(s).

-pt <item>

Specifies one or more ports. Ports are identified by one or more integers (item). Multiple ports are separated with a single comma (e.g. 1,2,3). A range of ports is separated with a hyphen (e.g. 0-4). The primary path of the disks must be on the specified port(s).

-cg <item>

Specifies one or more drive cages. Drive cages are identified by one or more integers (item). Multiple drive cages are separated with a single comma (e.g. 1,2,3). A range of drive cages is separated with a hyphen (e.g. 0-3). The specified drive cage(s) must contain disks.

-mg <item>

Specifies one or more drive magazines. The "1." or "0." displayed in the CagePos column of showpd output indicating the side of the cage is omitted when using the -mg option. Drive magazines are identified by one or more integers (item). Multiple drive magazines are separated with a single comma (e.g. 1,2,3). A range of drive magazines is separated with a hyphen (e.g. 0-7). The specified drive magazine(s) must contain disks.

-pn <item>

Specifies one or more disk positions within a drive magazine. Disk positions are identified by one or more integers (item). Multiple disk positions are separated with a single comma (e.g. 1,2,3). A range of disk positions is separated with a hyphen (e.g. 0-3). The specified position(s) must contain disks.

-dk <item>

Specifies one or more physical disks. Disks are identified by one or more integers (item). Multiple disks are separated with a single comma (e.g. 1,2,3). A range of disks is separated with a hyphen (e.g. 0-3). Disks must match the specified ID(s).

-tc_gt <number>

Specifies that physical disks with total chunklets greater than the number specified be selected.

-tc_lt <number>

Specifies that physical disks with total chunklets less than the number specified be selected.

-fc_gt <number>

Specifies that physical disks with free chunklets greater than the number specified be selected.

-fc_lt <number>

Specifies that physical disks with free chunklets less than the number specified be selected.

-devid <model>

Specifies that physical disks identified by their models be selected. Models can be specified in a comma-separated list. Models can be displayed by issuing the `showpd -i` command.

-devtype <type>

Specifies that physical disks must have the specified device type (FC for Fast Class, NL for Nearline, SSD for Solid State Drive, SCM for Storage Class Memory) to be used. Device types can be displayed by issuing the `showpd` command.

-rpm <number>

Drives must be of the specified relative performance metric, as shown in the "RPM" column of the `showpd` command. For FC and NL drives, the number corresponds to both a performance measure and actual rotational speed. For SSD & SCM drives, this field is "N/A".

Disks that satisfy all the specified characteristics are used. For example `-p -fc_gt 60 -fc_lt 230 -nd 2` specifies all the disks that have greater than 60 and less than 230 free chunklets and that are connected to node 2 through their primary path.

Specifiers

<fd>[-<td>]...

Specifies that the chunklets that were relocated from specified disk (<fd>), are moved to the specified destination disk (<td>). If destination disk (<td>) is not specified, then the chunklets are moved back to original disk (<fd>), see NOTES. The <fd> specifier is not needed if -p option is used, otherwise it must be used at least once on the command line. If this specifier is repeated, then the operation is performed on multiple disks.

Authority

- Super, Service
- Any role granted the `pd_movereloc` right

Usage

- Access to all domains is required to run this command.
- Chunklets moved from physical disk <fd> are treated as if they originated on disk <td>. Disk <td> can be the same as disk <fd>.
- Specifying the -dr option can be used to see if the specified moves succeed and what the results (quality) of the moves are.
- If the -partial option is used, the command relocates as many chunklets as possible and prints messages for the chunklets it could not move.
- If the <fd> move would reduce the availability of an LD and the <td> is not specified, then the system will find the best available destination.

Examples

The following example moves chunklets that were on physical disk 8 that were relocated to another position, back to physical disk 8:

```
cli% moverelocpd 8
Are you sure you want to move the chunklets ?
select q=quit y=yes n=no: y
Move        -State- -Detailed_State-
50:53-8:4   normal   normal
62:51-8:1   normal   normal
65:51-8:7   normal   normal
66:51-8:0   normal   normal
82:55-8:6   normal   normal
86:54-8:5   normal   normal
88:52-8:3   normal   normal
92:51-8:2   normal   normal
```

Use the following to get a description of each column in the above examples:

```
clihelp -col moverelocpd
```

movetodomain

Syntax

```
movetodomain [options] <object_name> <domain_name | domain_set>
```

Description

Moves objects from one domain to another.

Options

-vv

Specifies that the object is a virtual volume.

-cpg

Specifies that the object is a common provisioning group (CPG).

-host

Specifies that the object is a host.

-f

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

Specifiers

<object_name>

Specifies the name of the object to be moved.

<domain_name | domain_set>

Specifies the domain or domain set to which the specified object is moved. The domain set name must start with "set:".

To remove an object from any domain, specify the string "-unset" for the domain name or domain set specifier.

Authority

- Super, Edit
- Any role granted the `domain_moveto` right

Usage

- Access to all domains is required to run this command.
- The `movetodomain` command moves all objects that are directly or indirectly related to the specified object into the specified domain. Possible relationships include, but are not limited to, VLUNs between hosts and virtual volumes, virtual volumes using a CPG for snapshot space, and two virtual volumes sharing the same logical disc. If the -f option is not used, a list of all objects that will be modified is shown before the confirmation prompt.
- Only hosts are permitted to be members of domain sets.
- When moving a host to a domain set all objects directly or indirectly related to that host must be members of a domain contained in that domain set. If this is not the case then the command will fail.
- If the -cpg option is selected and the selected CPG is part of an Adaptive Optimization (AO) configuration, all CPGs in that AO configuration will be moved to the specified domain.

Examples

The following example displays the movement of virtual volume vv1 to domain SampleDomain:

```
cli% movetodomain -vv  vv1 SampleDomain
The following volumes will have their domain modified:
vv1
The following hosts will have their domain modified:
testhost
The following CPGs will have their domain modified:
SampleCPG
14 associated LDs will also have their domain changed.
Do you want to proceed with moving the above to domain SampleDomain?
select y=yes n=no: y
```

The following example displays the removal of host testhost from any domain:

```
cli% movetodomain -host testhost -unset
The following hosts will have their domain modified:

Id Name     Persona -WWN/iSCSI_Name- Port
2 testhost  Generic

Do you want to proceed with removing the domain of the above?
select y=yes n=no: y
```

Promote Commands

promotegroupsv

Syntax

```
promotegroupsv [options] <virtual_copy_name>[:<target_vvname>]...
```

Description

The promotegroupsv command copies the differences of snapshots back to their base volumes, allowing you to revert the base volumes to an earlier point in time.

Options

-rcp

Allows the promote operation to proceed even if the RW parent volume is currently in a Remote Copy volume group, if that group has not been started. If the Remote Copy group has been started, this command fails. This option cannot be used in conjunction with the -halt option.

-pri <high|med|low|auto>

Specifies the priority of the copy operation when it is started. This option allows the user to control the overall speed of a particular task. If this option is not specified, the promotegroupsv operation is started with default priority of medium. Auto priority can be used only if -online option is specified and this operation is started with auto as the default priority. High priority indicates that the operation will complete faster. Low priority indicates that the operation will run slower than the default priority task. If the priority is set to auto, then the array controls the speed of the operation based on the load in the array for all flash models. Auto priority behaves like medium priority for the converged flash models of the arrays containing Nearline (NL) and Fast Class (FC) drives. This option cannot be used with -halt option.

-halt

Cancels ongoing snapshot promotions. Marks the RW parent volumes with the "cpf" status that can be cleaned up using the **promotevvcopy** command or by issuing a new instance of the **promotesv** or promotegroupsv command. This option cannot be used in conjunction with any other option.

-online

Indicates that the promote operation will be executed while the target volumes have VLUN exports. The hosts should take the target LUNs offline to initiate the promote command, but can be brought online and used during the background tasks. Each specified virtual copy and its base volume must be the same size. The base volume is the only possible target of online promote, and is the default. To halt a promote started with the online option, use the **canceltask** command. The -halt and -target options cannot be combined with the -online option.

Specifiers

<virtual_copy_name>[:<target_vvname>]...

Specifies the names of the virtual copy volumes to be promoted, using up to 31 characters in length. Currently, vvset is not supported. If a RW target_vvname is supplied, then the target vv will be treated as the RW parent volume to be rolled back. If the target_vvname is not supplied, then the virtual copy volume will be promoted to its base. If the -online option is used, then no target names can be specified, the virtual volume(s) will be promoted to the base.

Authority

- Super, Edit
- Any role granted the svgrp_promote right

Usage

- A promote to a parent of larger size will include zeroing the additional capacity of the parent in order to correctly represent the data from the smaller virtual copy. Such zeroing activity will proportionally increase the duration of the promote task.
- Issue the **showvv** command to verify that differences in the snapshot volume are promoted to its target volume.
- Issue the `showvv -d` command to display the number of remaining blocks to be copied.
- Without the `-online` option, the RW virtual copy and the target of the promote must not be exported. With the `-online` option, the base (target) may be exported, but no snapshots in the family tree may be exported.
- Only one promote operation is allowed at a time within a virtual volume family tree.

Examples

The following example illustrates how virtual volumes vv1 and vv2 can be promoted with the differences from their base volume:

```
cli% promotegroupsv vv1 vv2
```

promotesv

Syntax

```
promotesv [options] <virtual_copy_name | VV_set>
```

Description

The `promotesv` command copies the differences of a snapshot back to its base volume, allowing you to revert the base volume to an earlier point in time.

Options

`-target <vvname|VV_set>`

Copies the differences of the virtual copy to the specified RW parent in the same virtual volume family tree. The default is to copy the differences to the base volume. This option cannot be used in conjunction with the `-halt` or `-online` options. When target is a set, the snapshot source must be a set as well. The order of the snapshot set must correspond to the order of the target set.

`-rcp`

Allows the promote operation to proceed even if the RW parent volume is currently in a Remote Copy volume group, if that group has not been started. If the Remote Copy group has been started, this command fails. This option cannot be used in conjunction with the `-halt` option.

`-halt`

Cancels an ongoing snapshot promotion. Marks the RW parent volume with the "cpf" status that can be cleaned up using the **promotevvcopy** command or by issuing a new instance of the `promotesv` command. This option cannot be used in conjunction with any other option.

`-pri <high|med|low|auto>`

Specifies the priority of the copy operation when it is started. This option allows the user to control the overall speed of a particular task. If this option is not specified, the `promotesv` operation is started with default priority of medium. Auto priority can be used only if `-online` option is specified and this operation is started with auto as the default priority. High priority indicates that the operation will complete faster. Low priority indicates that the operation will run slower than the default priority task. If the priority is set to auto, then the array controls the speed of the operation based on the load in the array for all flash models. Auto priority behaves like medium priority for the

converged flash models of the arrays containing Nearline (NL) and Fast Class (FC) drives. This option cannot be used with `-halt` option.

`-online`

Indicates that the promote operation will be executed while the target volume has VLUN exports. The host should take the target LUN offline to initiate the promote command, but can bring it online and use it during the background task. The specified virtual copy and its base volume must be the same size. The base volume is the only possible target of online promote, and is the default. To halt a promote started with the online option, use the `canceltask` command. The `-halt` and `-target` options cannot be combined with the `-online` option.

Specifiers

`<virtual_copy_name | VV_set>`

Specifies the name of the virtual copy volume, or set of virtual copy volumes, to be promoted, using up to 31 characters in length. The vv set name must start with "set:".

Authority

- Super, Edit
- Any role granted the `sv_promote` right

Usage

- A promote to a parent of larger size will include zeroing the additional capacity of the parent in order to correctly represent the data from the smaller virtual copy. Such zeroing activity will proportionally increase the duration of the promote task.
- Issue the `showvv` command to verify that differences in the snapshot volume are promoted to its target volume.
- Issue the `showvv -d` command to display the number of remaining blocks to be copied.
- Without the `-online` option, the RW virtual copy and the target of the promote must not be exported. With the `-online` option the base (target) may be exported, but no snapshots in the family tree may be exported.
- Only one promote operation is allowed at a time within a virtual volume family tree.

Examples

The following example illustrates how virtual volume vv1 can be promoted with the differences from its base volume:

```
cli% promotesv vv1
```

The following will promote all snapshots in a VV set:

```
cli% promotesv set:vv_set
```

The following will promote snapshot snap1 to parent vv1:

```
cli% promotesv -target vv1 snap1
```

The following will promote all snapshots in snapset to their parents in targetset:

```
cli% promotesv -target set:targetset set:snapset
```

promotevvcopy

Syntax

```
promotevvcopy <physical_copy_name | VV_set>
```

Description

The `promotevvcopy` command promotes a physical copy back to a regular base volume.

Options

None.

Specifiers

<physical_copy_name | VV_set>

Specifies the name of the physical copy, or set of physical copies, to be promoted, using up to 31 characters in length. The vv set name must start with "set:".

Authority

- Super, Edit
- Any role granted the `vvcopy_promote` right

Usage

- The physical copy must have completed the copy from the base volume.
- The saved snapshot of the parent of <physical_copy_name> is also removed.
- The `promotevvcopy` command can also be used to clean up a failed physical copy.
- Issue the **showvv** command to verify that promoted volume is a base volume.
- After a physical copy has been promoted, the association between it and its parent volume is broken; the physical copy and base volume can no longer resync. The saved snapshot of the parent of physical copy is also removed.

Examples

The following example promotes virtual volume vv1 to a base volume:

```
cli% promotevvcopy vv1
```

The following will promote all VV copies in a VV set:

```
cli% promotevvcopy set:vv_set
```

Remove Commands

removealert

Syntax

```
removealert [option] -a|<alert_ID>...
```

Description

The `removealert` command removes one or more alerts from the system.

Options

`-a`

Specifies all alerts from the system and prompts removal for each alert.

If this option is not used, then the <alert_ID> specifier must be used.

`-f`

Specifies that the command is forced. If this option is not used and there are alerts in the "new" state, the command requires confirmation before proceeding with the operation.

Specifiers

`<alert_ID>...`

Indicates a specific alert to be removed from the system. This specifier can be repeated to remove multiple alerts. If this specifier is not used, the -a option must be used.

Authority

- Super, Service
- Any role granted the `alert_remove` right

Usage

CAUTION: Use care when removing alerts. Alerts that have not been fixed or acknowledged should NOT be removed.

Access to all domains is required to run this command.

Examples

The following example displays all alerts from the system with the option to remove individual alerts:

```
cli% removealert -a

Id 120 - New
 Occurred 4 times, last at Tue May 03 22:45:47 PDT 2005
 Message code: 196609
 Tue May 03 22:23:17 PDT 2015
 Node: 0 Severity: Minor
 Firmware coredump event
 Firmware COREDUMP: recovered file /var/core/hba/fwcore.n00.s02.p01.20050503.224547

Alert 120 is marked as "New".
Are you sure you want to remove it?
select q=quit y=yes n=no: y
...
```

removeaocfg

Syntax

```
removeaocfg [options] {<AOCFG_name>|<pattern>}...
```

Description

The `removeaocfg` command removes specified Adaptive Optimization configurations from the system.

Options

`-pat <pattern>`

Specifies that specified patterns are treated as glob-style patterns and that all AO configurations matching the specified pattern are removed. By default, confirmation is required to proceed with the command unless the -f option is specified. This option must be used if the pattern specifier is used.

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

Specifiers

`<AOCFG_name>`

Specifies the name of the AO configuration to be removed

`<pattern>`

Specifies a glob-style pattern. This specifier can be repeated to remove multiple AO configurations. If this specifier is not used, the <AOCFG_name> specifier must be used. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.)

Authority

Super, Edit

Usage

None.

Examples

The following example removes AO configurations that start with test.

```
cli% removeaocfg -f -pat test*
```

removebanner

Syntax

```
removebanner [-cli|-ssh|-all]
```

Description

Removes the banner that was set using the **`setbanner`** command. The banner is text displayed to a user before login.

Options

`-cli`

Removes CLI banner.

`-ssh`

Removes SSH banner (default).

-all

Removes all banners.

Specifiers

None.

Authority

- Super
- Any role granted the `sshbanner_set` right

Usage

None.

Examples

The following example will remove the SSH banner:

```
cli% removebanner
    Banner(s) were removed.
```

removecert

Syntax

```
removecert {all|<SSL_service_name>} [-type <typename>]
```

Description

The `removecert` command is used to remove certificates that are no longer trusted. In most cases, it is better to overwrite the offending certificate with **importcert**. The user specifies which service to have its certificates removed. The removal can be limited to a specific type.

Options

-f

Skips the prompt warning the user of which certificates will be removed and which services will be restarted.

-type <typename>

Allows the user to limit the removal to a specific type. Note that types are cascading. For example, intca will cause the service certificate to also be removed.

Valid types are csr, cert, intca, and rootca.

Specifiers

<SSL_service>

Valid service names are cim, cli, ekm-client, ekm-server, ldap, qw-client, qw-server, syslog-gen-client, syslog-gen-server, syslog-sec-client, syslog-sec-server, wsapi, vasa, and unified-server.

The user may also specify all, which will remove certificates for all services.

Authority

Super

Usage

Use of the `removecert` command may be confusing. The SSL server services (cim, cli, and wsapi) must always have a certificate, and so a self-signed certificate will be automatically generated when their certificate is removed.

Examples

The following example shows how to remove all certificates for the cli:

```
cli% removecert cli
```

The following example shows how to remove just the root Certificate Authority for the cli:

```
cli% removecert cli -type rootca
```

removecorequest

Syntax

```
removecorequest [-f] -state {executed|cancelled|rejected}
removecorequest [-f] <requestid>
```

Description

The `removecorequest` command removes specified requests from the queue.

Options

`-f`

Suppresses the warning message.

`-state {executed|cancelled|rejected}`

Specifies removal of all commands from queue whose state is canceled, executed, or rejected.

Specifiers

`<requestid>`

Request ID that is to be removed from queue.

Authority

- CO
- Any role granted the `corequest_remove` right

Usage

None.

Examples

The following example shows how to remove a specific request with Id 12345:

```
cli% removecorequest 12345
This command will remove the request with id 12345.
Do you wish to continue?
select y=yes n=no:
```

The following example shows how to remove a specific request with Id 12345 forcefully:

```
cli% removecorequest -f 12345
```

The following example shows how to remove all executed requests:

```
cli% removecorequest -state executed
```

The following example shows how to remove all rejected requests:

```
cli% removecorequest -state rejected
```

The following example shows how to remove all canceled requests:

```
cli% removecorequest -state cancelled
```

removecpg

Syntax

```
removecpg [options <arg>] {<CPG_name>|<pattern>}...
```

Description

The `removecpg` command removes common provisioning groups (CPGs) from the system or removes specific logical disks from common provisioning groups.

Options

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

`-sa <LD_name>`

Specifies that the logical disk, as identified with the <LD_name> argument, used for snapshot administration space allocation is removed.

The <LD_name> argument can be repeated to specify multiple logical disks.

This option is deprecated and will be removed in a subsequent release.

`-sd <LD_name>`

Specifies that the logical disk, as identified with the <LD_name> argument, used for snapshot data space allocation is removed. The <LD_name> argument can be repeated to specify multiple logical disks.

This option is deprecated and will be removed in a subsequent release.

`-pat`

The specified patterns are treated as glob-style patterns and that all common provisioning groups matching the specified pattern are removed.

By default, confirmation is required to proceed with the command unless the -f option is specified. This option must be used if the pattern specifier is used.

Specifiers

`<CPG_name>`

Specifies the name of the common provisioning group that is either being removed or losing logical disks.

`<pattern>`

Specifies a glob-style pattern. This specifier can be repeated to remove multiple common provisioning groups. If this specifier is not used, the <CPG_name> specifier must be used. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.)

Authority

- Super, Edit
- Any role granted the `cpg_remove` right

Usage

- Access to all domains is required to run this command.
- The `removecpg` command fails if any of the logical disks, or the entire common provisioning group, is in use by a thinly provisioned virtual volume.
- If neither the -sa or -sd options are specified, the entire common provisioning group is removed, including all logical disks.
- The operation fails if any of the logical disks are in use.

Examples

The following example displays the removal of common provisioning group cpg1:

```
cli% removecpg cpg1
```

removedomain

Syntax

```
removedomain [option] <domain_name|pattern>...
```

Description

The `removedomain` command removes an existing domain from the system.

Options

-f

When using this option, the command does not ask for confirmation before removing the domain.

-pat

Specifies that names will be treated as glob-style patterns and that all domains matching the specified pattern are removed.

Specifiers

<domain_name|pattern>

Specifies the domain that is removed. If the -pat option is specified the domain_name will be treated as a glob-style pattern, and multiple domains will be considered.

Authority

- Super
- Any role granted the `domain_remove` right

Usage

- Access to all domains is required to run this command.
- The patterns are glob-style (shell-style) patterns. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.)
- Issuing the `removedomain` command with invalid domain names causes the command to exit without removing any domains.

Examples

The following example removes the domain named sample_domain from the system:

```
cli% removedomain -f sample_domain
```

The following example removes the domains that start with test:

```
cli% removedomain -f -pat test*
```

The following example removes any domains that start with test or sample:

```
cli% removedomain -f -pat test* sample*
```

removedomainset

Syntax

```
removedomainset [options] <setname|pattern> [<domain|pattern>...]
```

Description

The `removedomainset` command removes a domain set or removes domains from an existing set.

Options

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

`-pat`

Specifies that both the set name and domains will be treated as glob-style patterns.

Specifiers

`<setname|pattern>`

Specifies the name of the domain set. If the -pat option is specified the setname will be treated as a glob-style pattern, and multiple domain sets will be considered.

`<domain|pattern>`

Optional list of domain names that are members of the set. If no <domain>s are specified, the domain set is removed, otherwise the specified <domain>s are removed from the domain set. If the -pat option is specified the domain will be treated as a glob-style pattern, and multiple domains will be considered.

Authority

- Super, Edit
- Any role granted the `domainset_set` right can remove a domain set or domains from a domain set

Usage

- Access to all domains is required to run this command.
- The patterns are glob-style (shell-style) patterns. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.)

Examples

To remove a domain set:

```
cli% removedomainset domainset
```

To remove a single domain from a set:

```
cli% removedomainset domainset domain1
```

The following example removes the domain sets that start with test:

```
cli% removedomainset -f -pat test*
```

The following example removes any domains that start with test or sample from all domain sets:

```
cli% removedomainset -f -pat * test* sample*
```

removeflashcache

Syntax

```
removeflashcache [option]
```

Description

The `removeflashcache` command removes the flash cache from the cluster and will stop use of the extended cache.

Options

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

Specifiers

None.

Authority

- Super, Edit
- Any role granted the `flashcache_remove` right

Usage

None.

Examples

The following example removes the flash cache from the cluster, and stops use of the extended cache:

```
removeflashcache
```

removefpg

Syntax

```
removefpg [options] <fpgname|pattern> ...
```

Description

The `removefpg` command removes a file provisioning group and its underlying components from the system.

Options

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

`-forget`

Removes the specified file provisioning group which is involved in Remote DR, keeping the virtual volume intact.

`-wait`

Waits until the associated task is completed before proceeding.

This option will produce verbose task information.

-pat

The specified patterns are treated as glob-style patterns and that all file provisioning groups matching the specified pattern are removed.

By default, confirmation is required to proceed with the command unless the -f option is specified. This option must be used if the pattern specifier is used.

Specifiers

<fpgname>

The name of the file provisioning group to be removed.

This specifier can be repeated to remove multiple common provisioning groups.

<pattern>

Specifies a glob-style pattern. This specifier can be repeated to remove multiple common provisioning groups. If this specifier is not used, the <fpgname> specifier must be used. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.)

Authority

- Super, Edit.
- Any role granted `fpg_remove` right

Usage

- Access to all domains is required to run this command.
- It is necessary to remove any shares on the file provisioning group before removing the file provisioning group itself.
- When using the -forget option, the VVs must be specified when the recover action is performed, e.g. `createfpg -recover <vv>`.
- FPG cannot be removed if it contains a compliance enabled VFS or File Store. It can be removed once all associated compliance enabled VFSs and File Stores have been removed.

Examples

The following example removes a file provisioning group:

```
cli% removefpg examplefpg
```

removefsarchive

Syntax

The syntax for the `removefsarchive` command can be one of the following:

- `removefsarchive files [-basepath <basepath>] {-files <filepath> [,<filepath>]... | -importfile <source_path> | -inputfile <pathoffile>} -fstore <fstorename> [-fpg <fpgname>]<vfs>`
- `removefsarchive retention [-basepath <basepath>] {-files <filepath> [,<filepath>]... | -importfile <source_path> | -inputfile <pathoffile>} -fstore <fstorename> [-fpg <fpgname>] <vfs>`

- `removefsarchive scan [-fpg <fpgname>] -fstore <fstorename> <vfs> <jobid>`
- `removefsarchive auditlogs [-f] [-log <logname>] [-fpg <fpgname>] <vfs>`

Description

The `removefsarchive` commands are used to delete the WORM/WORM-retained file(s), remove the retention period, or delete the data validation scan on the fstore.

Options

`-f`

Suppresses the confirmation from user before removing the specified audit log or all compressed audit logs.

`-basepath <basepath>`

Base path of files for which retention settings are to be deleted.

`-files <filepath>[,<filepath>]...`

Relative path of files from the above base path to remove retention period or specified files. If basepath is not specified, it should be the absolute path of files.

`-inputfile <pathoffile>`

Path of input file containing list of files for which retention settings are to be deleted. The input file should be placed in file persona namespace (.admin file store). It should be absolute path of the input file, which is visible to the file persona namespace.

`-importfile <source_path>`

Path of import file containing list of files for which retention settings to be removed or files to be deleted. The input file should be present on client machine. The max size of import file can be up to 2GiB.

`-fpg <fpgname>`

Specifies the FPG that contains the VFS. This is required if VFS exists under multiple FPGs.

`-fstore <fstorename>`

Name of the file store under VFS.

`-log <logname>`

Name of the log file to be removed. If not mentioned, it will delete all the compressed logs at VFS level.

Specifiers

`<vfs>`

Specifies the name of the VFS.

`<jobid>`

Job Id of the data validation scan process.

Subcommands

`files`

Deletes the specified WORM/WORM-retained file(s).

`retention`

Removes the retention period for the file(s) specified.

`scan`

Removes the retention validation scan job.

auditlogs

Removes either all or a specific audit log from VFS.

Authority

- Super, Edit
- Any role granted the `fsarchive_remove` right
- Super for `-importfile <source_path>`

Usage

- Access to all domains is required to run this command.
- The option `-importfile <source_path>` is not available with SSH access to CLI.
- The following operations on compliance enabled File Stores are not allowed. - Removal of retained WORMed files. WORMed files with expired retention can be removed from a client accessing the associated share (e.g. NFS/CIFS). - Removal of the retention of retained WORMed files.

Examples

The following example deletes the specified WORM/WORM-retained files at the given file path:

```
cli% removefsarchive files -basepath /fpg1/vfs1 -files fstore1/abcd.txt,/fstore2/pqr.txt -fpg fpg1 -fstore fstore1 vfs1
cli% removefsarchive files -basepath /fpg1/vfs1 -importfile /home/worm/abcd.txt -fpg fpg1 -fstore fstore1 vfs1
```

The following example removes the retention period for the specified file:

```
cli% removefsarchive retention -basepath /fpg1/vfs1 -inputfile /fpg1/vfs1/.admin/abcd.txt -fstore fstore1 -fpg fpg1 vfs1
cli% removefsarchive retention -basepath /fpg1/vfs1 -fstore fstore1 -importfile /home/worm/abcd.txt -fpg fpg1 vfs1
```

The following example removes data validation scan process:

```
cli% removefsarchive scan -fstore fstore1 -fpg fpg1 vfs1 10
```

The following example removes the audit logs:

```
cli% removefsarchive auditlogs -f -log audit_fpg1_vfs1_2016-03-02-04-32-45.log.zip vfs1
cli% removefsarchive auditlogs vfs1
```

removefsgroup

Syntax

```
removefsgroup [options] <groupname|GID>
```

Description

The `removefsgroup` command removes a local group account associated with File Persona.

Options

-f

Specifies that the operation is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

Specifiers

<groupname|GID>

Specifies the local group name using up to 31 characters or a numeric GID. Valid characters are alphanumeric characters, periods, dashes (except first character), and underscores.

Authority

- Super
- Any role granted the `fsgroup_remove` right

Usage

Access to all domains is required to run this command.

Examples

The following example removes a local group named accounting:

```
cli% removefsgroup accounting
```

removefshare

Syntax

```
removefshare {smb|nfs|obj|ftp} [options <arg>] <vfs> <sharename>
```

Description

The `removefshare` command removes a file share for supported protocols. The folder will be deleted when all the following conditions are met:

- The share folder is empty.
- The share folder is not the root of the file store.
- The share folder does not have any other associated file shares.

The cleanup is repeated for each parent directory.

Options

The following options are for all subcommands:

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

`-fpg <fpgname>`

Specifies the file provisioning group (FPG) that <vfs> belongs.

If not specified, the command will find out the FPG based on the specified <vfs>. However, if <vfs> exists under multiple FPGs, -fpg must be specified.

`-fstore <fstore>`

Specifies the file store that the file share to be removed belongs.

If not specified, the <sharename> will be used as <fstore>.

Specifiers

`<vfs>`

Specifies the virtual file system (VFS) name.

`<sharename>`

The name of the share to be removed.

Subcommands

smb

Removes an SMB file share.

nfs

Removes an NFS file share.

obj

Removes an Object file share.

ftp

Removes an FTP file share.

Authority

- Super, Edit
- Any role granted the `fshare_remove` right.

Usage

Access to all domains is required to run this command.

Examples

The following example removes the SMB file share "myshare" without prompting for confirmation:

```
cli% removefshare smb -f -fstore myfstore myvfs myshare
```

The following example removes the NFS file share "myshare":

```
cli% removefshare nfs -fpg myfpg -fstore myfstore myvfs myshare
```

The following example removes the Object share "myshare":

```
cli% removefshare obj -fstore myfstore myvfs myshare
```

The following example removes the FTP share "myshare":

```
cli% removefshare ftp -fstore myfstore myvfs myshare
```

removefsip

Syntax

```
removefsip [options] <vfs> <id|ip>
```

Description

The `removefsip` command removes the network config of a virtual file server (VFS).

Options

-fpg <fpgname>

Specifies the file provisioning group (FPG) in which the VFS was created.

-f

Specifies that the operation is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

Specifiers

<id|ip>

Specifies the ID/IP for the network config.

<vfs>

Specifies the VFS which is to have its network config removed.

Authority

- Super, Edit
- Any role granted the `fsip_remove` right

Usage

Access to all domains is required to run this command.

Examples

The following example removes the network configuration of virtual file server vfs_1:

```
cli% removefsip -f vfs_1 12345678
cli% removefsip -f vfs_1 10.10.10.1
```

removefsnap

Syntax

```
removefsnap [options <arg>] <vfs> <fstore>
```

Description

The `removefsnap` command removes file store snapshots for File Persona.

Options

-f

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

-fpg <fpgname>

Specifies the file provisioning group (FPG) that <vfs> belongs.

If this is not specified, the command will find out the FPG based on the specified <vfs>. However, if <vfs> exists under multiple FPGs, -fpg must be specified.

-snapname <name>

Specifies the name of the snapshot to be removed. If this is not specified, all snapshots of the file store specified by <fstore> will be removed.

Specifiers

<vfs>

Specifies the virtual file server (VFS) name.

<fstore>

Specifies the file store name.

Authority

- Super, Edit
- Any role granted the `fsnap_remove` right

Usage

Access to all domains is required to run this command.

Examples

The following example removes all snapshots of "myfstore":

```
cli% removefsnap -fpg myfpg myvfs myfstore
```

removefspn

Syntax

```
removefspn [options] <vfs> <user> <spnname>
removefspn -force [options] <vfs> <spnname>
```

Description

The `removefspn` command removes SPN associated with a VFS to disable Kerberos authentication for File Persona.

Options

The following options are for all subcommands:

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

`-force`

Removes Kerberos keys from File Persona.

`-fpg <fpgname>`

Specifies the file provisioning group (FPG) that `<vfs>` belongs.

If not specified, the command will find out the FPG based on the specified `<vfs>`. However, if `<vfs>` exists under multiple FPGs, `-fpg` must be specified.

`-passwd <password>`

Specifies the password of the authorized user. If not specified, the user will be prompted for a cleartext password.

Specifiers

`<vfs>`

Specifies the virtual file system (VFS) name.

`<user>`

Specifies the authorized user having access to Active Directory domain controller.

`<spnname>`

Specifies the name of SPN to be removed.

Authority

- Super, Edit
- Any role granted the `fspn_delete` right.

Usage

Access to all domains is required to run this command.

Examples

The following example removes SPN associated with the specified VFS:

```
cli% removefspn -fpg fpg1 -passwd password vfs1 aduser spn1
```

The following example removes Kerberos keys associated with the specified VFS:

```
cli% removefspn -force vfs1 spn1
```

removefsroute

Syntax

```
removefsroute [-f] {<targetaddr>,{<subnetmask>|<prefixlen>},<vlantag> |
<routeidentifier>}
```

Description

`removefsroute` removes a route for a target identified with either route id or a unique combination of target address, vlantag and subnet mask. A warning is displayed to user before removing the route.

Options

`-f`

Suppresses confirmation from user before removing the route.

Specifiers

`<targetaddr>`

The target IPv4/IPv6 address for which the route is to be removed.

`<subnetmask>|<prefixlen>`

The subnet mask or prefix length for the target IP Address.

`<vlantag>`

The VLAN tag associated with route which needs to be removed.

`<routeidentifier>`

Instead of providing a combination of targetaddr,subnetmask|prefixlen and vlantag - a route identifier can be provided. Obtain the route identifier from the "`showfsroute -d`" command.

Authority

- Edit
- Any role granted `fsroute_set` right.

Usage

Access to all domains is required to run this command.

Examples

The following example removes a route with following details:

```
cli% removefsroute 10.16.23.101,16,8
```

The following example deletes a route without confirmation:

```
cli% removefsroute -f 10.16.23.101,255.255.192.0,12
```

The following example deletes a route with specified route ID:

```
cli% removefsroute staticRoute4e426403-98a7-4025-9b11-ca59afd9971c
```

removefstore

Syntax

```
removefstore [options] <vfs> <fstore>
```

Description

The `removefstore` command removes a file store and its underlying components from the system.

Options

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

`-fpg <fpgname>`

The name of the parent file provisioning group (FPG).

Specifiers

`<vfs>`

The name of the containing virtual file server (VFS).

`<fstore>`

The name of the file store to be removed.

Authority

- Super, Edit
- Any role granted `fstore_remove` right

Usage

- Access to all domains is required to run this command.
- The command will remove all underlying shares.

Examples

The following example removes a file store:

```
cli% removefstore -f -fpg fpg1 vfs1 examplefstore.
```

removefsuser

Syntax

```
removefsuser [options] <username|UID>
```

Description

The `removefsuser` command removes a local user account.

Options

`-f`

Specifies that the operation is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

Specifiers

<username|UID>

Specifies the local user name using up to 31 characters or a numeric UID.

Valid characters are alphanumeric characters, periods, dashes (except first character), and underscores.

Authority

- Super
- Any role granted the `fsuser_remove` right

Usage

Access to all domains is required to run this command.

Examples

The following example removes a local user named jdoe:

```
cli% removefsuser -f jdoe
```

removehost

Syntax

```
removehost [options] <hostname|pattern> [<WWN>... | <iscsi_name> ...]
```

Description

The `removehost` command removes a system host or paths to a host. If one or more paths are specified, the command removes only those paths, otherwise the entire host definition is removed.

Options

-rvl

Removes WWN(s) or iSCSI name(s) even if there are VLUNs exported to the host. This option cannot be used if the entire host definition is being removed.

-iscsi

Specifies that the paths are iSCSI names. If this option is not specified, the paths are WWNs.

-pat

Specifies that host name will be treated as a glob-style pattern and that all hosts matching the specified pattern are removed. This option cannot be used if the WWN or ISCSI_name specifiers are used.

-port <node:slot:port>...|<pattern>...

Specifies the NSP(s) for the zones, from which the specified WWN will be removed in the target driven zoning. Multiple array ports can be specified by either using a pattern or a comma-separated list. At least one WWN must be specified with this option.

Specifiers

<hostname|pattern>

Specifies the host name, using up to 31 characters. If the -pat option is specified the hostname will be treated as a glob-style pattern, and multiple hosts will be considered.

<WWN>

Only the specified WWN(s) path to the specified host is removed. This specifier is not required on the command line. If a WWN is not specified, the entire host definition is removed.

<ISCSI_name>

Specifies the host iSCSI name to be removed from the specified host. If no ISCSI name is specified, the entire host definition is removed.

Authority

- Super, Edit
- Any role granted the `host_remove` right

Usage

- Access to all domains is required to run this command.
- A host that has one or more VLUNs exported on it cannot be removed.
- Removing an entire host definition by issuing the `removehost <hostname>` command is not equivalent to removing all the paths associated with a host. The latter leaves a host definition with no paths associated to it, whereas the former removes the entire host definition.
- Verify the removal of hosts by issuing the **showhost** command.
- The patterns are glob-style (shell-style) patterns. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.)
- Issuing the `removehost` command with invalid host names causes the command to exit without removing any hosts.

Examples

The following example removes host test01:

```
cli% removehost test01
```

The following example removes the WWN 1000A0B3CC1C68BE for a host associated with the target driven zone for port 1:2:1. This does not remove the WWN from the host. If you want to remove the WWN from the host AND from the zone, then use "`removehost <host> <WWN>`".

```
cli% removehost -port 1:2:1 lynx 1000A0B3CC1C68BE
```

The following example removes the hosts that start with test:

```
cli% removehost -pat test*
```

removehostset

Syntax

```
removehostset [options] <setname|pattern> [<host|pattern>...]
```

Description

The `removehostset` command removes a host set or removes hosts from an existing set.

Options

-f

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

-pat

Specifies that both the set name and hosts will be treated as glob-style patterns.

Specifiers

<setname|pattern>

Name of the host set to remove. If the -pat option is specified the setname will be treated as a glob-style pattern, and multiple host sets will be considered.

<host|pattern>

Optional list of host names that are members of the set.

If no <host>s are specified, the host set is removed, otherwise the specified <host>s are removed from the host set. If the -pat option is specified the host will be treated as a glob-style pattern, and multiple hosts will be considered.

Authority

- Super, Edit
- Any role granted the `hostset_set` right can remove a host set or hosts from a host set

Usage

The patterns are glob-style (shell-style) patterns. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.)

Examples

The following example removes a host set:

```
cli% removehostset hostset
```

The following example removes a single host from a set:

```
cli% removehostset hostset host1
```

The following example removes the host sets that start with test:

```
cli% removehostset -f -pat test*
```

The following example removes any hosts that start with test or sample from all host sets:

```
cli% removehostset -f -pat * test* sample*
```

removeld

Syntax

```
removeld [options] {<LD_name>|<pattern>}...
```

Description

The `removeld` command removes a specified LD from the system service group.

Options

-pat

Specifies glob-style patterns. All LDs matching the specified pattern are removed. By default, confirmation is required to proceed with the command unless the -f option is specified. This option must be used if the pattern specifier is used.

-dr

Specifies that the operation is a dry run and no LDs are removed.

-f

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

-rmsys

Specifies that system resource LDs such as logging LDs and preserved data LDs are removed.

-unused

Specifies the command to remove non-system LDs. This option cannot be used with the -rmsys option.

Specifiers

<LD_name>

Specifies the LD name, using up to 31 characters. Multiple LDs can be specified.

<pattern>

Specifies a glob-style pattern. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.) Multiple LDs can be specified. If this specifier is not used, the <LD_name> specifier must be used.

Authority

- Super, Service, Edit
- Any role granted the `ld_remove` right

Usage

CAUTION: System resource LDs are required for correct operation of the HPE 3PAR Storage system. Removal of system resource LDs should be performed by qualified service personnel. Incorrect use of the -rmsys option can result in data loss.

- If the LD is mapped to a virtual volume, the LD is not removed and the command stops at that LD.
- Issuing the `removeld` command with the -pat option specified returns a request for confirmation to remove LDs, unless the -f option is specified.
- Do not issue the `removeld` command while a **movech**, **movechtospare**, **movepd**, **movepdtospare**, or **moverelocpd** operation is being executed. Issue the **showldch** command to view operations currently running on the system.
- Verify the removal of LDs by issuing the **showld** command.

Examples

The following example removes LD nf_vv_1.0.usr.1:

```
cli% removeld -f nf_vv_1.0.usr.1
```

removercopygroup

Syntax

`removercopygroup [options] {<group_name> ... | <pattern> ...}`

Description

Removes a remote copy volume group or multiple remote copy groups that match a given pattern.

Options

-f

Suppresses user interaction with the command. Without this option, the command requires confirmation before proceeding with its operation.

If the Remote Copy group is not empty and you use the `-removevv` option, the command requires confirmation even if you specify `-f`.

-removevv

Removes volumes from the remote sides.

-keepsnap

Specifies retaining the resync snapshot of local volume. The retained snapshot reflects the state of the secondary volume. If you readmit the volume to a remote copy group, you can use the retained snapshot as the starting snapshot. The snapshot begins with `sv.rcpy`.

-pat

Treats patterns as glob-style patterns and all remote copy groups matching the specified pattern are removed. By default, requires confirmation to proceed with the command unless you specify the `-f` option. Required when using the `<pattern>` specifier.

-keepalua

Keeps the ALUA state of local volume from changing. Required during Peer Motion migrations of Remote Copy Groups. See *Volume migration in a Remote Copy Primary group* in the *HPE 3PAR Peer Motion User Guide* for details.

Specifiers

<group_name>

The name of the group to remove.

<pattern>

Specifies a glob-style pattern. You can repeat this specifier to remove multiple groups. With no *<pattern>* specified, you must specify the *<group_name>*.

Authority

- Super, Edit
- Any role granted the `rcopygroup_remove` right

Usage

- Requires the HPE 3PAR Remote Copy Software license.
- Contact your local service provider for more information.
- The system generates an error if you use this command while remote copy is in progress. The group must be stopped using the **stoprcopygroup** command.
- You must have access to the domain of the specified group (*<group_name>*) to run this command.
- If the `mirror_config` policy is set for the target system of this group, and the group is a primary group, this command is mirrored to the target and the corresponding secondary group is also removed. If the policy is set and the group is a secondary group, this command fails.

- The `-keepsnap` and `-removevv` options are mutually exclusive.
- The `removercopygroup` command removes all the associations configured in the specified group and removes the group name and any remote copy synchronization snapshots affiliated with volumes in the group.

Examples

The following example removes Group1 from a remote copy system:

```
removercopygroup Group1
```

The following example removes remote copy groups that start with the name testgroup:

```
removercopygroup -pat testgroup*
```

The following example removes the remote copy group (Group1) and retains the resync snapshots associated with each volume:

```
removercopygroup -keepsnap Group1
```

The following example removes Group1 and associated volumes from the local system and removes the corresponding volumes on all the target systems of Group1:

```
removercopygroup -removevv Group1
```

removercopytarget

Syntax

```
removercopytarget [options] <target_name>
```

Description

The `removercopytarget` command removes target designation from a remote copy system and removes all links affiliated with that target definition.

Options

`-cleargroups`

Removes all groups that have no other targets or dismiss this target from groups with additional targets. All groups that contain this target must be stopped before the -cleargroups command can be issued.

`-f`

Suppresses user interaction with the command. No confirmation requested when specifying `-cleargroups`.

`-keepalua`

Keeps the ALUA state of local volume from changing. Used only with the `-cleargroups` option. Required during Peer Motion migrations of Remote Copy Groups. See *Volume migration in a Remote Copy Primary group* in the *HPE 3PAR Peer Motion User Guide* for details.

Specifiers

`<target_name>`

The target name for the target definition to remove.

Authority

- Super, Edit
- Any role granted the `rcopytarget_remove` right

Usage

- Access to all domains is required to run this command.
- This command requires the HPE 3PAR Remote Copy Software license. Contact your local service provider for more information.
- A target definition cannot be deleted if it is being used by any group.

Examples

The following example removes targets System1 and System2 from the primary and backup systems in a remote copy pair:

```
removercopytarget System1
removercopytarget System2
```

removesched

Syntax

```
removesched [options] {<schedname>|<pattern>}...
```

Description

The `removesched` command removes a scheduled task from the system.

Options

`-pat`

Specifies that certain patterns are treated as glob-style patterns and that all scheduled tasks matching the specified pattern are removed. This option must be used if the pattern specifier is used.

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

Specifiers

`<schedname>`

Specifies the scheduled task name. This specifier can be repeated to remove multiple scheduled tasks. If this specifier is not used, the pattern specifier must be used. Can be at most 31 characters in length.

`<pattern>`

Specifies a glob-style pattern. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.) This specifier can be repeated to remove multiple scheduled tasks. If this specifier is not used, the <schedname> specifier must be used.

Can be at most 31 characters in length.

Authority

- Super, Service, Edit
- Any role granted the `sched_setany` right can remove any user scheduled tasks
- Any role granted the `sched_setown` right can only remove their own scheduled tasks

Usage

None.

Examples

None.

removesnmpmgr

Syntax

```
removesnmpmgr [option <arg>] <manager_IP>
```

Description

The `removesnmpmgr` command removes previously registered SNMP managers from receiving traps.

Options

`-p <port_number>`

Specifies the port number where the manager receives traps. If not specified, the port number defaults to 162.

Specifiers

`<manager_IP>`

Specifies the IPv4 or IPv6 address of the host where the manager runs.

IPv6 address is in hexadecimal, case insensitive, and separated by colons. An example would be:

`5def:2008:0:0:abcd:0:0:161a`

In addition, a double colon (::) can be used once in an address to replace multiple fields of zeros. For example: `5def:2008:0:0:abcd::161a`

Authority

- Super
- Any role granted the `snmpmgr_remove` right

Usage

- Access to all domains is required to run this command.
- SNMP managers are registered by issuing the **`addsnmpmgr`** command.
- Verify the removal of SNMP managers by issuing the **`showsnmpmgr`** command.

Examples

The following example displays the removal of a manager with the IPv4 address from the list of registered managers:

```
cli% removesnmpmgr 123.45.67.89
```

The following example displays the removal of a manager with the IPv6 address from the list of registered managers:

```
cli% removesnmpmgr 5def:2008:abcd::161a
```

removesnmppw

Syntax

```
removesnmppw [options]
```

Description

The `removesnmppw` command allows a user to remove SNMP access community string passwords.

Options

-rw|-r|-w

Removes the read-write (-rw), read-only (-r) or write-only (-w) password.

If not specified, the read-write community string password is removed.

-f

Forces the operation so that the command does not require confirmation before proceeding.

Specifiers

None.

Authority

- Super
- Any role granted the `snmppw_remove` right

Usage

- Access to all domains is required to run this command.
- The following codes are returned indicating success or failure:
  - 0 indicates that the password was removed and the command was successful.
  - 1 indicates that the command failed.
  - 2 indicates that the password does not exist.
- After a password has been removed, the system manager can no longer use that password to send requests to the SNMP agent.
- Verify the removal of SNMP passwords by issuing the **showsnmppw** command.

Examples

The following example displays the removal of the read-write SNMP access password:

```
cli% removesnmppw -f
```

removesnmpuser

Syntax

```
removesnmpuser [options] <username>
```

Description

The `removesnmpuser` command allows a user to remove SNMPv3 user.

Options

-f

Forces the operation so that the command does not require confirmation before proceeding.

Specifiers

<username>

Specifies the SNMPv3 user name, using up to 31 characters.

Authority

- Super
- Any role granted the `snmpuser_remove` right

Usage

- The following codes are returned indicating success or failure:
  - 0 indicates that the password was removed and the command was successful.
  - 1 indicates that the command failed.
  - 2 indicates that the user does not exist.
- After a user has been removed, an SNMPv3 manager can no longer use that username to send requests to the SNMP agent. The `removesnmpuser` command has no impact on local user since it only removes the user from SNMPv3.
- Verify the removal of SNMPv3 users by issuing the **`showsnmpuser`** command.

Examples

The following example displays the removal of the an SNMPv3 user:

```
cli% removesnmpuser -f joe
```

removespare

Syntax

```
removespare [options] <chunklet_specifier>...
```

Description

The `removespare` command removes chunklets from the spare chunklet list.

Options

`-f`

Specifies that the operation is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

`-p`

Specifies that partial completion of the command is acceptable.

Chunklets specified that are not on the current service group spare list are ignored.

Specifiers

`<chunklet_specifier>...`

The chunklet specifier can be issued in the following formats:

`<PD_ID:chunklet_num>...`

Specifies the identification of the physical disk (PD_ID) and the position number of the chunklet on the disk (chunklet_num). This specifier can be repeated.

`<PD_ID>:a...`

Specifies the identification of the physical disk (PD_ID) and all (a) chunklets on the disk. This specifier can be repeated.

a:<chunklet_num>...

Specifies the chunklet number on all physical disks. This specifier can be repeated.

a:a

Remove all spare chunklets that are not currently used.

-pos <cage:mag:disk:chunklet_num>

Specifies the position of a specific chunklet identified by its position in a drive cage, drive magazine, physical disk, and chunklet number. For example -pos 1:0.2:3:121, where 1 is the drive cage, 0.2 is the drive magazine, 3 is the physical disk, and 121 is the chunklet number.

-pos <cage:mag:disk:a>

Specifies that all chunklets on a physical disk, identified by drive cage number, drive magazine number, and disk number, are marked to be removed.

Authority

- Super, Service
- Any role granted the `spare_remove` right

Usage

- Access to all domains is required to run this command.
- Verify the removal of spare chunklets by issuing the **showspare** command.
- If a wildcard ("a") is used or the -p flag is specified, prints the number of spares removed. Otherwise, if all the explicitly specified spares could not be removed, prints an error message.

Examples

The following example removes a spare chunklet from position 3 on physical disk 1:

```
cli% removespare 1:3
```

removespcredential

Syntax

```
removespcredential
```

Description

The `removespcredential` command removes the SP credentials on the array and sets a random password for 3parsvc.

CAUTION: It is a nontrivial process to reverse this action. It SHOULD NOT be run unless you are removing the SP completely for security reasons.

Options

None.

Specifiers

None.

Authority

- Super
- Any role granted the `spcred_remove` right

Usage

Once the command is executed, the 3paredit, 3parbrowse and 3parservice users are removed and a new random password is set for 3parsvc. This command also removes the SSH key for 3parsvc.

Examples

The following example removes all the 3parservice, 3paredit and 3parbrowse users and sets a new random password for 3parsvc user:

```
cli% removespcredential
```

removesralertcrit

Syntax

```
removesralertcrit [options <arg>] <name>|<pattern>
```

Description

The `removesralertcrit` command removes a criterion that System Reporter evaluates to determine if a performance alert should be generated.

Options

`-pat`

Specifies that certain patterns are treated as glob-style patterns and that all criteria matching the specified pattern are removed. By default, confirmation is required to proceed with the command unless the -f option is specified. This option must be used if the pattern specifier is used.

`-f`

Do not ask for confirmation before removing this criterion.

Specifiers

`<name>`

Specifies the name of the criterion to remove.

`<pattern>`

Specifies a glob-style pattern. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.) This specifier can be repeated to remove multiple criteria tasks. If this specifier is not used, the <name> specifier must be used.

Authority

- Super, Service
- Any role granted the `sralertcrit_remove` right

Usage

Removing a criterion also automatically fixes all alerts associated with that criterion.

Examples

The following example removes the criterion named busy_port:

```
cli% removesralertcrit busy_port
```

removesshkey

Syntax

```
removesshkey [options]
```

Description

The `removesshkey` command removes SSH public key to disable key authentication.

Options

`-i key_ID`

Only removes the key with the specified ID. The ID can be found in the output of `showsshkey`.

Specifiers

None.

Authority

- Super, Service, Edit, Browse
- Any role granted the `sshkey_remove` right

Usage

- After removing an SSH public key on the storage system, the user it belongs to cannot use the associated private key to log in. If all SSH keys for a user are removed, the user must use name and password to log in.
- The `showuser -k` command can be used to display users that have SSH keys (see **showuser**).
- Only a user can remove their keys.

Examples

The following example removes all ssh keys for the current user:

```
cli% removesshkey
```

removetask

Syntax

```
removetask [options <arg>] -a
removetask [options <arg>] -t <hours>
removetask [options <arg>] <task_ID>...
```

Description

The `removetask` command removes information about one or more completed tasks and their details.

Options

`-a`

Removes all tasks including details.

`-d`

Removes task details only.

-f

Specifies that the command is to be forced. You are not prompted for confirmation before the task is removed.

-t <hours>

Removes tasks that have not been active within the past <hours>, where <hours> is an integer from 1 through 99999.

Specifiers

<task_ID>

Allows you to specify tasks to be removed using their task IDs.

Authority

- Super, Edit
- Any role granted the `task_remove` right

Usage

With this command, the specified task ID and any information associated with it are removed from the system. However, task IDs are not recycled, so the next task started on the system uses the next whole integer that has not already been used. Task IDs roll over at 29999. The system stores information for the most recent 2,000 tasks.

Examples

The following example shows how to remove a task based on the task ID:

```
cli% removetask 2
Remove the following tasks?
2
select q=quit y=yes n=no: y
```

The following example shows how to remove all tasks, including details:

```
cli% removetask -a
Remove all tasks?
select q=quit y=yes n=no: y
```

removetemplate

Syntax

```
removetemplate [options] {<template_name>|<pattern>} ...
```

Description

The `removetemplate` command removes one or more virtual volume (VV), logical disk (LD), and common provisioning group (CPG) templates.

Options

-f

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

-pat

The specified patterns are treated as glob-style patterns and that all templates matching the specified pattern are removed. By default, confirmation is required to proceed with the command unless the -f

option is specified. This option must be used if the pattern specifier is used.

Specifiers

<template_name>

Specifies the name of the template to be deleted, using up to 31 characters. This specifier can be repeated to remove multiple templates.

<pattern>

Specifies a glob-style pattern. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.) This specifier can be repeated to remove multiple templates. If this specifier is not used, the <template_name> specifier must be used.

Authority

- Super
- Any role granted the `template_remove` right

Usage

Access to all domains is required to run this command.

Examples

The following example displays the forced removal of template vv1:

```
cli% removetemplate -f vv1
```

removeuser

Syntax

```
removeuser [option] <username>...
```

Description

The `removeuser` command removes a user account from the system.

Options

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

Specifiers

`<username>...`

Specifies a login name using any combination of letters and numbers.

This argument can be repeated to specify multiple user names.

Authority

- Super
- Any role granted the `user_remove` right

Usage

- Access to all domains is required to run this command.
- Do not remove users 3parsvc and 3paradm.
- A user cannot remove oneself. The last user on the system cannot be removed.

- Verify the removal of users by issuing the **showuser** command.
- All open connections for a given user will be removed when that user is removed.
- When a user's connection is removed any running commands will be aborted. Use **showuserconn** to examine open connections.
- When successful, `removeuser` removes a user's SSH key.
- Before attempting to remove a user connected through the WSAPI server, all sessions associated with that WSAPI user must first be removed via the **removewsapisession** command.
- Use the **showwsapisession** command to view the WSAPI session information details.

Examples

The following example displays the forced removal of user1 from the system:

```
cli% removeuser -f user1
User removed.
```

removeuserconn

Syntax

```
removeuserconn [options] <user_ID> <user_name> <IP_address>
```

Description

The `removeuserconn` command removes user connections to the current system.

Options

-pat

Specifies that the <user_ID>, <user_name>, and <IP_address> specifiers are treated as glob-style (shell-style) patterns and all user connections matching those patterns are removed. By default, confirmation is required to proceed with removing each connection unless the -f option is specified.

-dr

Specifies that the operation is a dry run and no connections are removed.

-f

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

Specifiers

<user_ID>

Specifies the ID of the user to be removed.

<user_name>

Specifies the name of the user to be removed.

<IP_address>

Specifies the IP address of the user to be removed.

Authority

- Super
- Any role granted the `userconn_remove` right

Usage

- Access to all domains is required to run this command.
- Because user connections can disappear from the time they are listed and the time they are removed, the `removeuserconn` command continues past errors while removing individual connections if the -pat option is specified.
- Connections associated with the WSAPI server can only be removed via the **removewsapisession** command.
- Use the **showwsapisession** command to view the WSAPI session information details.

Examples

The following example displays the forced removal of user user1 at IP address 127.0.0.1:

```
cli% removeuserconn -f 2315 user1 127.0.0.1
Removing user connection Id:2315 Name:user1 Addr:127.0.0.1
```

removevfs

Syntax

```
removevfs [options] <vfs>
```

Description

The `removevfs` command removes a virtual file server (VFS) and its underlying components from the system.

Options

-f

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

-fpg <fpgname>

Names of the file provisioning group (FPG) containing the VFS.

Specifiers

<vfs>

The name of the VFS to be removed.

Authority

- Super, Edit
- Any role granted `vfs_remove` right

Usage

- Access to all domains is required to run this command.
- All underlying elements will be removed from the system.

Examples

The following example removes a VFS examplevfs, which is contained in examplefpg:

```
cli% removevfs examplevfs
```

removevlun

Syntax

```
removevlun [options] <VV|VVSet> <LUN> <n:s:p>
removevlun [options] <VV|VVSet> <LUN> <host|hostset>
removevlun [options] <VV|VVSet> <LUN> <host|hostset> <n:s:p>
removevlun [options] <VV|VVSet> <LUN> <n:s:p> <host|hostset>
```

Description

Removes a virtual volume SCSI LUN export definition from the system.

Options

-novcn

Prevents the issuance of a VLUN Change Notification (VCN) after removal of the VLUN.

For direct connect or loop configurations, a VCN consists of a Fibre Channel Loop Initialization Primitive (LIP).

For fabric configurations, a VCN consists of a Registered State Change Notification (RSCN) that is sent to the fabric controller.

-pat

Treats *<VV_name>*, *<LUN>*, *<node:slot:port>*, and *<host_name>* specifiers glob-style patterns, and removes all VLUNs matching the specified pattern. By default, confirmation is required to proceed with the command unless you specify the `-f` option.

-dr

Specifies that the operation is a dry run and does not remove VLUNs.

-f

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

-set

Removes all the VLUNs connected with the VVs in a vvset used as a VMware vVol Storage Container.

Specifiers

<VV|VVSet>

Specifies the volume or volume set name, using up to 31 characters in length. The volume set name must start with "set:".

<LUN>

Specifies the LUN id to remove.

<n:s:p>

Specifies that exports to the specified port are removed. If this specifier is not used, the `<host|hostset>` specifier must be used.

n

Specifies the system node.

s

Specifies the PCI bus slot in the node.

p

Specifies the FCS port number of the card in PCI bus slot.

<host|hostset>

Specifies requests that exports to the specified host or host set, named using up to 31 characters, be removed. The host set name must start with "set:". If this specifier is not used, the `node:slot:port` specifier must be used.

Authority

- Super, Edit
- Any role granted the `vlun_remove` right

Usage

- To remove a specific VLUN, you must supply the same specifiers and options that you used when creating that VLUN. Use the `showvlun -t` command to view all created VLUN specifiers and options.
- If the VLUN to be removed is a matched-set, specify both the port and the hostname.
- Entering the `removevlun` command with the `-pat` option specified returns a request for confirmation to remove VLUNs, unless you also specify the `-f` option.
- To verify the removal of VLUNs, enter the **showvlun** command.

Examples

The following example deletes VLUNs for volume vv0, LUN 0, host host1:

```
cli% removevlun -f vv0 0 host1
```

The following will remove exported VLUNs for vv0, LUN 1, host set host_set:

```
cli% removevlun -f vv0 1 set:host_set
```

The following will remove the VLUNs for VV set vv_set, starting with LUN 1, host host1:

```
cli% removevlun -f set:vv_set 1 host1
```

The following will remove the VLUNs for VV set vv_set, starting with LUN 1, host set host_set:

```
cli% removevlun -f set:vv_set 1 set:host_set
```

removevv

Syntax

```
removevv [options] <VV_name>|<pattern>...
removevv -expired [options] [<VV_name>|<pattern>...]
```

Description

The `removevv` command removes Virtual Volumes (VVs) from the system.

Options

-pat

Specifies that specified patterns are treated as glob-style patterns and that all VVs matching the specified pattern are removed. By default, confirmation is required to proceed with the command unless the -f option is specified. This option must be used if the pattern specifier is used.

-f

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

-dr

Specifies that the operation is a dry run and no VVs are removed.

-stale

Specifies that all stale VVs can be removed.

This option cannot be used together with -expired.

-expired

Removes specified expired volumes.

This option cannot be used together with -stale.

-snaponly

Removes the snapshot copies only.

-cascade

Removes specified volumes and their descendent volumes as long as none has an active VLUN. It will remove any VLUN templates as long as there were no active VLUNs. It will remove the volumes from all the volume sets.

If -expired is specified, all expired volumes and their descendent volumes will be removed regardless if they are expired or not.

If -stale is specified, all stale volumes and their descendent volumes will be removed regardless if they are stale or not.

-nowait

Prevents command blocking that is normally in effect until the vv is removed. Removal of the vv is performed in the background. Only snapshots can be removed with this option. An attempt to remove a base VV results in an error.

Specifiers

<VV_name>

Specifies the virtual volume name, using up to 31 characters. This specifier can be repeated to remove multiple VVs. If this specifier is not used, the pattern specifier must be used.

<pattern>

Specifies a glob-style pattern. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.) This specifier can be repeated to remove multiple VVs. If this specifier is not used, the <VV_name> specifier must be used.

Authority

- Super, Edit
- Any role granted the `vv_remove` right can remove base volumes
- Any role granted the `vvcopy_remove` right can remove physical volumes
- Any role granted the `sv_remove` right can remove virtual copies

Usage

- Issuing the `removevv` command with invalid virtual volume names causes the command to exit without removing any virtual volumes.
- Any virtual volumes exported as VLUNs are not removed and the command stops at that virtual volume.

- Any virtual volume that contains snapshots cannot be removed.
- Verify the removal of VVs by issuing the **showvv** command.

Examples

The following example removes virtual volume vv0:

```
cli% removevv -f vv0
```

The following example removes all the expired volumes:

```
cli% removevv -f -expired
```

The following example removes the volumes that start with test and are snapshot copies:

```
cli% removevv -f -snaponly -pat test*
```

The following example removes vv1_snap if it is a snapshot and all its descendents:

```
cli% removevv -f -snaponly -cascade vv1_snap
```

removevvset

Syntax

```
removevvset [options] <setname|pattern> [<VV|pattern>...]
```

Description

The `removevvset` command removes a VV set or removes VVs from an existing set. Exported VV sets cannot be removed.

Options

-f

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

-pat

Specifies that both the set name and VVs will be treated as glob-style patterns.

Specifiers

<setname|pattern>

Specifies the name of the VV set to remove. If the -pat option is specified the setname will be treated as a glob-style pattern, and multiple VV sets will be considered.

<VV|pattern>

Optional list of VV names that are members of the set.

If no <VV>s are specified, the vv set is removed, otherwise the specified <VV>s are removed from the VV set. If the -pat option is specified each VV will be treated as a glob-style pattern, and multiple VVs will be considered.

Authority

- Super, Edit
- Any role granted the `vvset_set` right can remove a volume set or volumes from a volume set

Usage

The patterns are glob-style (shell-style) patterns. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.)

Examples

The following example removes a VV set:

```
cli% removevvset vvset
```

The following example removes a single VV from a set:

```
cli% removevvset vvset vv1
```

Trying to remove an exported VV set results in an error:

```
cli% removevvset -f set:1
    Unable to delete vv set 1 exported to lun 1, host h1
```

The following example removes the vvsets that start with test:

```
cli% removevvset -f -pat test*
```

The following example removes any VVs that start with test or sample from all vvsets:

```
cli% removevvset -f -pat * test* sample*
```

removewsapisession

Syntax

```
removewsapisession [options] <id> <user_name> <IP_address>
```

Description

Removes the WSAPI user connections from the current system.

Options

-pat

Treats the *<id>*, *<user_name>* and *<IP_address>* as glob-style (shell-style) patterns and removes all WSAPI user connections matching those patterns. Requires confirmation to proceed with removing each connection unless you specify the `-f` option.

-dr

Specifies that the operation is a dry run and no connections are removed.

-f

Specifies a forced command. Without this option, the command requires confirmation before proceeding with the operation.

-close_sse

Specifies closing of the Server Sent Event (SSE) connection channel. Does not remove the WSAPI session credential for SSE.

Specifiers

<id>

Specifies the Id of the WSAPI session connection to be removed.

<user_name>

Specifies the name of the WSAPI user to be removed.

<IP_address>

Specifies the IP address of the WSAPI user to be removed.

Authority

- Super
- Any role granted the `wsapisession_remove` right

Usage

- Access to all domains is required to run this command.
- Because user connections can disappear from the time they are listed and the time they are removed, the `removewsapisession` command continues past errors while removing individual connections if the -pat option is specified.

Examples

The following example displays the forced removal of all sessions for user user1.

```
cli% removewsapisession -f -pat * user1 *
Removing user WSAPI connection id:1383876394258938 Name:user1 Ip_Addr:16.94.225.96
Removing user WSAPI connection id:1383876393423508 Name:user1 Ip_Addr:16.94.225.96
```

Restore Command

restorefsconf

Syntax

```
restorefsconf [-fpg <fpg>] <vfs>
```

Description

The `restorefsconf` command restores a configuration backup for a VFS. The FPG/VFS names of the source and destination need to match before the restore can be done. Configuration backup artifact created by `backupfsconf` needs to be kept in the .admin file store of the destination VFS in the directory configbackup.

Options

`-fpg <fpgname>`

Specifies the FPG for the VFS.

Specifiers

`<vfs>`

Specifies the name of the VFS.

Authority

- Super, Service
- Any role granted `fsconf_restore` right

Usage

- Access to all domains is required to run this command.
- Hewlett Packard Enterprise recommends that you perform the configuration restore before restoring a data backup.
- When performing a configuration restore, the process maintains a restoration order of the file stores, VFS IP, Share information of NFS, CIFS, and Object/Certificates, AV configuration and Quota information.
- Manually configure the cluster-wide AV VSE data before doing a configuration restore.
- The following actions must be observed when creating a configuration restore:
  - Manually set up the FPG and the VFS.
  - Create a folder called "configbackup" under the VFS in the .admin file store, and manually copy the backup artifact into that folder.
  - The NDMP data restore creates folders for each of the CIFS and NFS shares.
  - The backed-up VFS IP is recreated during restore. It will fail if there is an IP conflict because the IP already taken.

  See "**`backupfsconf`**" for additional information on configuration backup and restore.

Examples

The following example restores a configuration backup on file provisioning group named "samplefpg" in virtual server named "samplevfs":

```
cli% restorefsconf -fpg samplefpg samplevfs
```

servicecage

Syntax

```
servicecage remove [-force] <cagename|wwn>
servicecage startfc|unstartfc|endfc|resetfc|hresetfc|clearlog [-f] [-ovrd]
[a|b]|[<0|1>] <cagename>
```

Description

The `servicecage` command is necessary when executing removal and replacement actions for a drive cage interface card module. The `startfc` or `unstartfc` subcommands are used to initiate service on a cage, and the `endfc` subcommand is used to indicate that service is completed.

Options

`-f`

Forces the operation. When this option is not used, the subcommands `reset` and `hreset` require confirmation before proceeding. This option will be ignored if specified with subcommand `start` or `end`.

`-force`

Only valid with `remove` subcommand. Forces removal of the cage even if it has active ports in use. If drives within this cage are in use, and if sysmgr is restarted after the remove operation is complete, access to those PDs will be lost. HPE recommends not to use the `-force` operation unless absolutely necessary. Instead, remove the paths to the specified cage before running the `remove` subcommand.

`-pcm`

For DCS11 and DCS12, this specifies that the Power Cooling Module (PCM) will be serviced. For DCN5, this specifies the Power Cooling Battery Module (PCBM) will be serviced.

`-iom`

Specifies that the I/O module will be serviced. This option is not valid for DCN5 or DCN6 cages.

Specifiers

`0|1`

For subcommands `reset` and `hreset`, this specifies the interface card number of the cage to be reset. For subcommands `start` and `end`, this specifies the number of the module indicated by `-pcm` or `-iom` to be serviced.

`<cage_name>`

Specifies the name of the drive cage to be serviced.

Subcommands

`startfc`

Prepare an FC-AL module for removal.

`unstartfc`

Stop the `startfc` subcommand.

endfc

Indicates that service on the drive cage is completed and allows the system to resume use of the interface card module. Both `startfc` and `unstartfc` need this action to complete the service operation. Permitted for Edit, Super and Service roles.

resetfc

Initiates a soft reset of the interface card for DC1, DC3, DCS1, DCS2, and DCN1 drive cages. DC2 and DC4 drive cages will be reset and the interface card specifier is ignored.

hresetfc

Initiates a hard reset of the interface card for DC1, DC3, DCS1, DCS2, and DCN1 drive cages. DC2 and DC4 drive cages will be reset and the interface card specifier is ignored.

remove

Removes the indicated drive cage (indicated with the <cage_name> specifier) from the system. If this subcommand is used, the a|b or 0|1 port specifier is not required. This command fails when the cage has active ports or is in use.

clearlog

For DC2 and DC4 cages, clears the log in the cage. The Interface card module specifier is ignored. It is not supported on any cages other than DC1, DC2 and DC4.

Authority

Super, Service

Usage

- Access to all domains is required to run this command.
- Issuing the `servicecage` command results in chunklet relocation, causing a dip in throughput.
- The `unstartfc` subcommand is provided if a mistake was made when issuing the `servicecage` command. The `unstartfc` subcommand stops the original command.
- After issuing the `startfc` or `unstartfc` subcommands, the `endfc` subcommand must be issued to indicate that service is completed and to restore the drive cage to its normal state.

Examples

The following example displays the commencement of Fibre Channel hot-plugging for drive cage cage0:

```
cli% servicecage startfc -f -ovrd cage0
```

servicehost

Syntax

```
servicehost list
servicehost remove [-f] [<node:slot:port> [<WWN_or_iSCSI_name> ... ]]
servicehost copy [-f] <src_node:slot:port> <WWN_or_iSCSI_name_pattern>
<dest_node:slot:port>
```

Description

The `servicehost` command executes removal and replacement actions for a host connecting to a storage system port.

Options

-f

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

Specifiers

<node:slot:port>

Specifies the host to be removed which is visible to the specified port.

This specifier can only be used with the remove subcommand.

<WWN_or_iSCSI_name>

Specifies a host's World Wide Name(WWN). This specifier can only be used with the remove and copy subcommands. When used with the remove subcommand, this specifier can be repeated.

<WWN_or_iSCSI_name_pattern>

Specifies that the indicated WWN or iSCSI name is treated as a glob-style pattern.

<src_node:slot:port>

Specifies the source port when exporting VLUNs. This specifier can only be used with the copy subcommand.

<dest_node:slot:port>

Specifies the destination port when exporting VLUNs. This specifier can only be used with the copy subcommand.

Subcommands

list

Displays a list of all inactive hosts. Permitted for all users.

remove

Removes an inactive host, as specified with the <WWN_or_iSCSI_name> specifier, from the indicated port (<node:slot:port>) and its associated VLUNs. If the <WWN_or_iSCSI_name> specifier is not issued with the remove subcommand, all inactive hosts on the specified port and their LUNs are removed. If the <node:slot:port> specifier is not used with the remove subcommand, all inactive hosts in the system and their LUNs are removed.

copy

Copies all active VLUNs from the specified source port (as specified with <src_node:slot:port>) from host WWNs or iSCSI names matching the specified pattern (<WWN_or_iSCSI_name_pattern>) to the destination port (as specified with <dest_node:slot:port>). If necessary, the port is reset to target mode.

Authority

- Super, Service
- Any role granted the `host_service` right

Usage

- Access to all domains is required to run this command.
- If the <WWN_or_iSCSI_name> specifier is not issued with the remove subcommand, all inactive hosts on the specified port and their LUNs are removed.
- If the <node:slot:port> specifier is not used with the remove subcommand, all inactive hosts in the system and their LUNs are removed.

Examples

The following example displays the creation of a host on port 0:2:1 for the export of VLUNs from port 2:1:1:

```
cli% servicehost copy 2:1:1 20000200000CF790 0:2:1
Are you sure you want to run servicehost?
select q=quit y=yes n=no: y
```

The following example displays the removal of an inactive host from port 2:1:1:

```
cli% servicehost remove 2:1:1 20000200000CF790
Removing inactive host 20000200000CF790 on port 2:1:1
Are you sure?
select q=quit y=yes n=no: y
```

servicemag

Syntax

The syntax for the `servicemag` command can be one of the following:

```
servicemag start [options] <cage_ID> <magazine>
servicemag start [options] -pdid <PD_ID_0>...<PD_ID_3>
servicemag resume|unmark [options] <cage_ID> <magazine>
servicemag status [options] [<cage_ID> <magazine>]
servicemag clearstatus <cage_ID> <magazine>
```

Description

The `servicemag` command executes service on a drive magazine or disk.

Options

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

`-wait`

Specifies that the `servicemag` command must be completed before exiting.

If not specified, the `servicemag` command automatically exits before service begins and the remainder of the process proceeds in the background. This option is only valid for the `start`/`resume` operations.

`-log`

Specifies that write operations to chunklets of valid disks are written to a logging area when the magazine is out of service (or removed). When the disks return and you enter the `servicemag resume` option, the system writes the data from the logging logical disks to those disks. The system also relocates chunklets to free or spare space if chunklet failure might result in a RAID set becoming invalid (for example, if two disks would be missing from a RAID-5 logical disk). All other used chunklets are placed in the logging mode. This option is only valid for the start operation and should always specify `-nodisks` or `-pdid <PD_ID>`.

NOTE: This option can reduce the redundancy level of RAID sets until `servicemag resume` completes successfully.

`-nodisks`

Specifies that the serviced drive magazine's disk drives are valid and do not need to be replaced. This option can only be used with the `start` subcommand and `-log` option.

-partial

This option can only be used with the `resume` subcommand. Specifies that as many chunklets as possible are relocated. Error messages are printed for those chunklets that could not be relocated.

-pdid <PD_ID_0>...<PD_ID_3>

Specifies one to four physical disks (by physical disk ID) that need to be serviced or replaced. The PDs have to be in the same cage and magazine. If the `-log` option is also specified, logging will only apply to the remaining disks on the magazine and not the ones specified in this option. The disks specified in this option will be vacated out to be replaced and will not be logged. This option can only be used with the `start` subcommand, and cannot be used with the `<cage_ID>` or `<magazine>` specifiers.

-d

Displays detailed status of a `servicemag` operation. If the `-d` option is excluded, a summary of the status is displayed. This option is only valid for the `status` subcommand.

-dryrun

This option is used to estimate the length of time a `servicemag start` or `servicemag resume` will take to complete. If a `-dryrun` is started before an actual `servicemag start` or `resume`, the estimation will be based on empirical data with no I/O on the system. For a more accurate estimate, which will do runtime calculations based on the load on the system, please use '`servicemag status <cage_ID> <magazine>`' once '`servicemag start`' or '`servicemag resume`' has been initiated. This option is only valid for the `start` and `resume` operations.

Specifiers

<cage_ID>

Specifies the ID of the cage. Use the **showcage** command to determine the system's drive cage IDs.

<magazine>

Specifies the drive magazine within the specified drive cage to be serviced based on the side of the drive magazine and drive magazine slot number.

For drive chassis with a single drive cage (type DC2, DC4, and DC3), the valid syntax is `<position>` (the numeric position of the drive magazine).

Position values for DC2 drive cages can be from 0 to 9. Position values for DC3 drive cages can be from 0 to 15.

Subcommands

start

Specifies that the `servicemag` command informs the system manager to log or relocate disks on a drive magazine so that the drive magazine can be removed for service.

resume

Specifies that the `servicemag` command informs the system manager that a drive magazine is replaced and that data services can be resumed.

unmark

Specifies that the `servicemag` operation is stopped and its internal state is reset. Since the `servicemag` operation is a multi-step process, specifying unmark stops the `servicemag` operation at the completion of the current step. Relocation of chunklets is considered one step and can take from several minutes to several hours (depending on number of chunklets) to complete. If `servicemag unmark`' is issued during the relocation phase, the spin-down phase will indicate that the `servicemag -start` operation 'Failed'. This can be confirmed by issuing `servicemag status -d`'. If the intent was to prevent `servicemag` from completing, issue `servicemag resume`' to move the data back onto the drive(s) in this magazine once the `servicemag start`' operation reports a failure. The 'unmark' option should not be used without consulting with HPE engineering.

status

Specifies that the status of the servicemag command operations on a drive magazine are displayed. Permitted for Super, Service, Edit and Browse users. Summary status is given as default, and the -d option may be used to get detailed status.

clearstatus

Clears the log shown by the servicemag command status for the given cage and magazine.

Authority

- Super, Service
- Any role granted the mag_service right

Usage

- Access to all domains is required to run this command.
- Issuing the servicemag command results in chunklet relocation that causes a dip in throughput.
- When a servicemag command is issued with the -log option, all chunklets on the disks in the drive magazine being serviced are marked as "normal,smag".
- This state indicates an active servicemag operation on the disks.
- Any I/O on the chunklets marked "normal,smag" changes the states to "logging" and I/O is written to the logging logical disks.
- Issuing the servicemag resume command causes playback of the log. Any chunklets in the "logging" state enter "playback" state as their data is played back. After all the data is played back, the chunklets return to the "normal" state. Any chunklets in the "normal,smag" state return directly to the "normal" state.
- Chunklet states can be checked by issuing either the **showldch** or **showpdch** commands.
- By default, the servicemag command relocates all chunklets in the magazine to destinations starting first with local (such as on the owning node for the logical disk) spares, then local free space, then remote spare and finally remote free space.
- In the case when a drive needs to be replaced, the -log option should always be used in conjunction with the -pdid option.
- Replacing disks that have not had data completely relocated can lead to data loss. If the drives need to be replaced, the -log option should be used in conjunction with the -pdid option.
- For the servicemag start command only, instead of using the <cage_ID> <magazine> specifiers, the disk(s) to be serviced can also be specified using the -pdid option. The advantage is that "servicemag resume" will dismiss the old pdid when it successfully completes. However if "servicemag unmark <cage_ID> <magazine>" is run after servicemag start command with -pdid option, then the disk will not be dismissed.

Examples

The following example displays the suspension and resumption of data services on drive magazine 0 in drive cage 2:

```
cli% servicemag start -log -wait 2 0
Begin servicemag start log 2 0.0...
... disks in mag : 2 0.0
... valid disks: wwn [2000000087043098] id [20] diskpos [0]
..................... wwn [2000000087008150] id [21] diskpos [1]
... not valid disks:
... mark disk wwn [2000000087043098] id [20] as non usable for ld allocation
```

```
... mark disk wwn [2000000087008150] id [21] as non usable for ld allocation
... relocating chunklets to spare space
... relocating chunklets of fail sets after logging to spare space
... logging chunklets from pd wwn [2000000087043098] id [20]
... logging chunklets from pd wwn [2000000087008150] id [21]
... spinning down disk wwn [2000000087043098] id [20]
... spinning down disk wwn [2000000087008150] id [21]
... bypassing mag 2 0.0
... bypassed mag 2 0.0
servicemag start 2 0.0  Succeeded

cli% servicemag resume 2 0
Begin servicemag resume 2 0.0...
... onlooping mag 2 0.0
... checking for valid disks...
... disks in mag : 2 0.0
... valid disks: wwn [2000000087043098] id [20] diskpos [0]
..................... wwn [2000000087008150] id [21] diskpos [1]
..................... wwn [20000000870042F6] id [22] diskpos [2]
..................... wwn [2000000087007E6D] id [23] diskpos [3]
... not valid disks:
... playback chunklets from pd wwn [2000000087043098] id [20]
... playback chunklets from pd wwn [2000000087008150] id [21]
... playback chunklets from pd wwn [20000000870042F6] id [22]
... playback chunklets from pd wwn [2000000087007E6D] id [23]
... 74 chunklets still waiting to be played back or relocating...
.... 18 chunklets still waiting to be played back or relocating..
... All chunklets played back / relocated.
... no chunklets to move
... marking pd wwn [2000000087043098] id [20] as usable for ld allocation
... marking pd wwn [2000000087008150] id [21] as usable for ld allocation
... marking pd wwn [20000000870042F6] id [22] as usable for ld allocation
... marking pd wwn [2000000087007E6D] id [23] as usable for ld allocation
servicemag resume 2 0.0  Succeeded
```

servicenode

Syntax

```
servicenode <subcommand> [options] <nodeid>
```

Description

The `servicenode` command informs the system that a certain component will be replaced, and will cause the system to indicate the physical location of that component.

Options

`-ps <psid>`

Specifies which power supply will be placed into servicing-mode.

Accepted values for <psid> are 0 and 1.

`-pci <slot>`

Only the service LED corresponding to the PCI card in the specified slot will be illuminated. Accepted values for <slot> are 0 through 9.

`-fan <fanid>`

Specifies which node fan will be placed into servicing-mode. Accepted values for `<fanid>` are 0 and 1 for HPE 3PAR 10000 systems and 0, 1, and 2 for HPE 3PAR 20000 systems.

-drive

Only the service LED corresponding to the node's internal drive will be illuminated.

-bat

Specifies that the node's battery backup unit will be placed into servicing-mode.

Specifiers

<nodeid>

Indicates which node the `servicenode` operation will act on. Accepted values are 0 through 7.

Subcommands

start

Specifies the start of service on a node.

status

Displays the state of any active servicenode operations.

end

Specifies the end of service on a node.

Authority

- Super, Service
- Any role granted the `node_service` right

Usage

- Access to all domains is required to run this command.
- On HPE 3PAR 7000 and HPE 3PAR 8000 series systems only the node LED can be illuminated. On Safe-To-Remove platforms such as HPE 3PAR 20000 systems, if a FRU is found unsafe to remove, the command will return an error.

Examples

In the following example, fan 1 on node 3 is put into servicing-mode.

```
cli% servicenode start -fan 1 3
```

Set Commands

setalert

Syntax

```
setalert new|ack|fixed {<alert_ID>...|-a}
```

Description

The `setalert` command sets the status of system alerts.

Options

`-a`

Specifies that the status of all alerts be set. If not specified, the <alert_ID> specifier must be specified.

Specifiers

`<alert_ID>...`

Specifies that the status of a specific alert be set. This specifier can be repeated to indicate multiple specific alerts. Up to 99 alerts can be specified in one command. If not specified, the -a option must be specified on the command line.

`new|ack|fixed`

Specifies that the alert(s), as indicated with the <alert_ID> specifier or with option -a, be set as "New" (new), "Acknowledged" (ack), or "Fixed" (fixed).

Authority

- Super, Service
- Any role granted the `alert_set` right

Usage

- Verify the status of alerts by issuing the **`showalert`** command.
- Access to all domains is required to run this command.

Examples

The following example sets the status of all system alerts to "New":

```
cli% setalert new -a
```

setaocfg

Syntax

```
setaocfg [options <arg>] <aocfg_name>
```

Description

The `setaocfg` command updates an Adaptive Optimization configuration.

Options

`-t0cpg <cpgname>`

Specifies the Tier 0 CPG for this AO config.

-t1cpg <cpgname>

Specifies the Tier 1 CPG for this AO config.

-t2cpg <cpgname>

Specifies the Tier 2 CPG for this AO config.

-mode <mode>

Specifies the optimization bias for the AO config and can be one of the following:

- Performance: Move more regions towards higher performance tier.
- Balanced: Balanced between higher performance and lower cost.
- Cost: Move more regions towards lower cost tier.

-t0min <size>[g|G|t|T]
-t1min <size>[g|G|t|T]
-t2min <size>[g|G|t|T]

Specifies the minimum space utilization of the tier CPG for AO to maintain when optimizing regions between tiers. The size can be specified in MB (default) or GB (using g or G) or TB (using t or T).

Setting a minimum to 0 (default) indicates that no minimum space utilization will be enforced.

-t0max <size>[g|G|t|T]
-t1max <size>[g|G|t|T]
-t2max <size>[g|G|t|T]

Specifies the maximum space utilization of the tier CPG. AO will move regions into and out of the CPG based on their relative access rate history, but will not exceed this maximum size in the CPG.

The size can be specified in MB (default) or GB (using g or G) or TB (using t or T). Setting a max to 0 (default) indicates that AO will use other indicators to decide the maximum CPG space utilization:

either the CPG sdgl, sdgw, or maximum possible growth size.

-name <newname>

Specifies a new name for the AO configuration of up to 31 characters in length.

Specifiers

None.

Authority

Super, Service, Edit

Usage

- Two or more Tier CPGs must be defined.
- If domains are in use on the system, all CPGs must be in the same domain or not in any domain.
- A CPG can only belong to one AO configuration.
- A CPG can be removed from a tier by specifying a null string "". At least two tiers must remain. When a CPG is removed from a tier, the allocated space on that CPG is not removed and will remain in the CPG. So before removing a CPG from the AOCFG, all the space in the CPG should be moved to other tiers. This can be done by setting the max for the AOCFG tier to 1 which will force subsequent runs of `startao` command to move space out of the CPG to other tiers. After all the space has been moved out of the CPG, it can be removed from the AOCFG.
- A CPG can be moved and swapped between tiers within the same AO configuration.

- A CPG cannot be moved between AO configurations - it must first be removed from one and then added to the second.
- Active use of Adaptive Optimization requires an Adaptive Optimization license. Contact your local 3PAR representative for information.

Examples

The following example changes the optimization bias (mode) of an AO configuration called aocfg1 to Performance:

```
cli% setaocfg -mode Performance aocfg1
```

The following example sets the tier 1 CPG for AO configuration aocfg2 to R5FCcpg. R5FCcpg must not already be used in an AO configuration, and the tier 1 CPG for aocfg2 must not already be set:

```
cli% setaocfg -t1cpg R5FCcpg aocfg2
```

The following example swaps a CPG between tiers:

```
cli% setaocfg -t0cpg R1SSDcpg -t2cpg R6NLcpg aocfg3
```

The following example removes a CPG from a tier. This assumes that aocfg1 has 3 tiers initially:

```
cli% setaocfg -t0cpg "" aocfg1
```

setauthparam

Syntax

The syntax of the `setauthparam` command can be one of the following:

```
setauthparam [-f] <param> <value>
setauthparam [-f] <map-param> <map-value>...
setauthparam [-f] -clear <param>...
setauthparam [-f] -clearall
```

Description

The `setauthparam` command is used to set the authentication and authorization parameters.

Options

`-f`

Does not ask for a confirmation before performing the operation.

`-clearall`

Clears all the authentication parameters.

`-clear`

Clears only the specified authentication parameters.

Specifiers

Specifiers for the `setauthparam` command can be issued as the following:

Table 3: Values for <param> <value> Specifiers

<param>	<value>
account-name-attr	The attribute of an account object that holds the user's username.
account-obj	The objectClass attribute of an account object.
accounts-dn	Indicates the base of the subtree in the DIT to search for objects that hold account information. Multiple bases can be specified using multiple <value> arguments. It is mutually exclusively with the groups-dn variable.
allow-ssh-key	Set this value to 1 to allow LDAP users to set a public SSH key with the **setsshkey** command (default 0). Clearing or setting the variable to 0 disables the setting of new keys for LDAP users but any existing keys remain until they are removed with the **removesshkey** command. This parameter only affects LDAP users, not local users.
binding	The LDAP binding type must be one of the following: • simple—use simple binding with the server. • SASL—use a SASL mechanism that is expected by the server, with the mechanism set by the sasl-mechanism variable.
domain-name-attr	When set, the mapping of groups to domains is enabled. For a user that is a member of a group that maps to a role level, the value of domain-name-attr is used to look up an attribute in the group that holds the name of the domain. If the domain is too long or contains characters that are not allowed in a domain name, the name is truncated to the maximum length of a domain name and invalid characters are replaced with an underscore (_).
domain-name-prefix	When domain-name-prefix is set, the value of the attribute specified by domain-name-attr is a candidate domain name. The value of domain-name-prefix is a character string used to extract the domain name from the candidate. The value is an optional exclamation point (!) followed by a character string called the prefix. The exclamation point is a flag that means the presence of the prefix is required and is described more in the paragraphs that follow. The candidate domain name is searched for the presence of the prefix and if found, the domain name starts after the first occurrence of the prefix and stops before the first space or tab following it or at the end of the candidate domain name. If the prefix is not found, the behavior depends on the flag. If the exclamation point was not used (there is no flag), the candidate domain name becomes the domain name. If the flag is present, the candidate domain name is rejected and there is no domain name. As a last step, and as described for domain-name-attr, domain names can be truncated and have invalid characters replaced. Some examples of the effects of domain-name-prefix are shown in the following table, "Examples of Domain Name Prefix Effects".

Table Continued

<param>	<value>
group-name-attr	The attribute in the group object that holds the group's name.
group-obj	Indicates the objectClass attribute of a group object.
groups-dn	Indicates the base of the subtree in the DIT in which to search for objects that hold group information. It is mutually exclusively with the accounts-dn variable.
kerberos-realm	The Kerberos realm.
kerberos-server	Indicates the numeric IP address of the Kerberos server if different from the LDAP server.
ldap-2FA-cert-field	The name of the field in the two factor authentication certificate that will be used to uniquely identify the user in LDAP.
ldap-2FA-object-attr	The attribute in the LDAP object that corresponds to the two factor authentication certificate field value.
ldap-port	Indicates the port of the LDAP server (default: 389 for non-SSL, 636 for SSL).
ldap-reqcert	Indicates whether a valid server certificate should be required in order to proceed. (The default value is 0.) The value is only used when LDAP binding = sasl; ignored when binding = simple.
ldap-server	Numeric IP address of the LDAP server. If hostname, the value is the name of the LDAP server in its certificate or the value of the LDAP principal stored in the Kerberos database, and will usually be a fully qualified domain name. Multiple servers can be specified using multiple <value> arguments.
ldap-server-hn	Indicates the hostname of the LDAP server if ldap-server is an IP address. This value must be set when the ldap-reqcert option is set or the sasl-mechanism option is set to GSSAPI and ldap-server is set to an IP address.
ldap-service-account	The username of the service account used to communicate with LDAP to validate LDAP user security group membership.
ldap-service-account-password	The password that corresponds to the user set in ldap-service-account.
ldap-ssl	To use SSL when communicating with the LDAP server, set the value to 1. (The default value is 0.)
ldap-StartTLS	Set this parameter to one of the following: • no—(default) to not request the server use StartTLS. • try—to request the server use StartTLS but does not require it to proceed. • require—requests that the server uses StartTLS and continue only when it succeeds.

Table Continued

<param>	<value>
ldap-type	The type of LDAP server being used: • MSAD—Microsoft Active Directory. • OPEN—OpenLDAP. • RHDS—Red Hat Directory Server. • LOAD—Load Balancer.
member-attr	The attribute that holds the names of users in the group.
memberof-attr	The attribute that holds the name of a group of which the user is a member.
native-domain	The default LDAP domain to be prepended to the Kerberos realm when no other subdomain is supplied.
sasl-mechanism	When the binding is SASL, the SASL mechanism must be one supported by the LDAP server. The system allows the mechanisms of PLAIN, DIGEST-MD5, and GSSAPI.
user-attr	Indicates the attribute used to form a DN for simple binding. When the attribute ends with a backslash, the DN is the concatenation of the value of user-attr and the username. When the attribute does not end with a backslash, it is as described for the user-dn-base variable.
user-dn-base	When using simple binding, the authentication process attempts to bind the user to an entry in the server's Directory Information Tree (DIT). The Distinguished Name (DN) of the entry is a concatenation of the value of user-attr, "=", the username, ",", and the value of user-dn-base. If group-obj is set to posixGroup, the value of user-dn-base is also used as the base for searching for the user's posixAccount entry, regardless of binding type.

Table 4: Examples of Domain Name Prefix Effects

Candidate	Domain-name-prefix	Result
dom1	ISDom=	dom1
ISDom=dom2	ISDom=	dom2
ISDom=dom3	!ISDom=	dom3
dom4	!ISDom=	There is no resulting domain name because "ISDom=" does not appear in the candidate.

Table 5: Values for Specifiers <map-param> and <map-value>

<map-param>	<map-value>
super-map	A group name that grants the user the Super role level if the user is a member of that group. Multiple group names can be specified using multiple <map-value> arguments. A value of "*" matches any group name.
service-map	Same as super-map, but for the Service level.
edit-map	Same as super-map, but for the Edit level.
browse-map	Same as super-map, but for the Browse level.
create-map	Same as super-map, but for the Create level.
basic_edit-map	Same as super-map, but for the Basic_edit level.
3PAR_AO-map	Same as super-map, but for the 3PAR_AO role.
3PAR_RM-map	Same as super-map, but for the 3PAR_RM role.

The user will be mapped to highest permission level that matches.

Authority

- Super
- Any role granted the `authparam_set` right.

Usage

- When using Microsoft Active Directory, the minimum parameters that must be specified:
- ldap-type, kerberos-realm, and the necessary group permission maps.
- The array will use GSSAPI SASL binding with the default values for all other parameters, using all LDAP servers that have SRV records in DNS.
- Users who have been provided with a password that allows successful binding with the LDAP server will nevertheless be denied access if they are not members of any of the groups specified by the map parameters.
- The matching of a user's groups with the mapping rules is done in the order of the mapping parameters provided previously. When there are multiple matches, the first match determines the user's role.
- Domain names found with the use of domain-name-attr and domain-name-prefix are only potential domains and a user will only have roles in those if they are actually existing domains. The **showdomain** command will list existing domains.
- The **showauthparam** command displays authentication parameter settings and the **checkpassword** command can be used to see how the parameters are used to bind with an LDAP server and search for data to determine the user's role.
- When domains are enabled, Super or Service role provide access across all domains. Any other domain names are ignored for Super or Service role users. You can only have the Service role when no other domains match for roles other than Super or Service. If other such domains match, the Service role match is ignored.

Examples

None.

setbanner

Syntax

```
setbanner [-cli|-ssh|-all]
```

Description

Sets the banner that displays to users before login.

This command prompts for a new banner. After entering the banner text, press **ENTER** to save the new banner. To verify the banner, use the **showbanner** command.

Options

-cli

Sets CLI banner.

-ssh

Sets SSH banner (default).

-all

Sets all banners.

Specifiers

None.

Authority

- Super
- Any role granted the `sshbanner_set` right

Usage

None.

Examples

The following example will set all banners:

```
cli% setbanner -all
```

setbattery

Syntax

```
setbattery [options <arg>] <node_ID>
setbattery [options <arg>] <node_ID> <powersupply_ID> <battery_ID>
```

Description

The `setbattery` command may be used to set battery information such as the battery's expiration date, its recharging time, and its serial number. This information gives the system administrator a record or log of the battery age and battery charge status.

Options

-s <serial_number>

Specifies the serial number of the battery using a limit of 31 alphanumeric characters. This option is not supported on HPE 3PAR 10000 and 20000 systems.

-x <exp_date>

Specifies the expiration date of the battery (mm/dd/yyyy). The expiration date cannot extend beyond 2037.

-l

Specifies that the battery test log is reset and all previous test log entries are cleared.

-r

Specifies that the battery recharge time is reset and that 10 hours of charging time are required for the battery to be fully charged.

This option is deprecated.

Specifiers

<node_ID>

Specifies the node number where the battery is installed. Node_ID is an integer from 0 through 7.

<powersupply_ID>

Specifies the power supply number on the node using either 0 (left side from the rear of the node) or 1 (right side from the rear of the node).

<battery_ID>

Specifies the battery number on the power supply where 0 is the first battery.

Authority

- Super, Service
- Any role granted the `battery_set` right

Usage

- Access to all domains is required to run this command.
- To view battery status information, issue the **showbattery** command.
- In order to modify battery attributes on HPE 3PAR 10000 systems and 20000 systems, only the node_ID and the battery attributes of interest are required. Older platforms require the node_ID, powersupply_ID, and battery_ID in order to modify battery attributes.

Examples

The following example resets the battery test log and the recharging time for a newly installed battery on node 2, power supply 1, and battery 0, with an expiration date of July 4, 2006:

```
cli% setbattery -x 07/04/2006 2 1 0
```

setcage

Syntax

```
setcage [subcommand <arg>]... <cage_name>
```

Description

The `setcage` command enables service personnel to set or modify parameters for a drive cage.

Options

None.

Specifiers

<cage_name>

Indicates the name of the drive cage that is the object of the `setcage` operation.

Subcommands

position <position>

Sets a description for the position of the cage in the rack, where <position> is a description to be assigned by service personnel (for example, "left-top")

ps <model>

Sets the model of a cage power supply, where <model> is a model name to be assigned to the power supply by service personnel.

id <id>

Changes the id of the specified cage, where <id> is the new id to be assigned to the cage.

Authority

- Super, Service
- Any role granted the `cage_set` right.

Usage

- Access to all domains is required to run this command.
- The parameters specified by the `setcage` command appear in the `showcage -d` output (see the **showcage** command).
- The power supply model cannot be modified if the information is automatically retrieved from the system.
- The cage id cannot be changed if:
  - There is an ongoing `servicemag` operation.
  - There are LDs in the system created with the -cg parameter.
  - There is a cage already using the specified ID.

Examples

The following example demonstrates how to assign cage1 a position description:

```
cli% setcage position "Cabinet 0 Bay 5 " cage1
```

The following example demonstrates how to assign model names to the power supplies in cage1:

```
cli% setcage ps Magnetek cage1
```

The following example demonstrates how to swap the IDs of two cages:

```
cli% showcage -enclosure
Cage ------WWN------- Locate FailInd WarnInd FailReq WarnReq -State-
0 51402EC010A0EC80          No     Off     Off     Off     Off
OK
1 51402EC010A90F80          No     Off     Off     Off     Off
OK
2 51402EC010A97A00    No     Off     Off     Off     Off     OK
```

```
cli% setcage id 1 cage2
The new ID specified is already in use by another cage.
Please specify a new ID which is not already in use.

cli% setcage id 3 cage2

cli% showcage -enclosure
Cage ------WWN------- Locate FailInd WarnInd FailReq WarnReq -State-
0 51402EC010A0EC80          No     Off     Off     Off     Off
OK
1 51402EC010A90F80          No     Off     Off     Off     Off
OK
3 51402EC010A97A00          No     Off     Off     Off     Off     OK

cli% setcage id 2 cage1        cli% showcage -enclosure
Cage ------WWN------- Locate FailInd WarnInd FailReq WarnReq -State-
0 51402EC010A0EC80    No     Off     Off     Off     Off     OK
2 51402EC010A90F80    No     Off     Off     Off     Off     OK
3 51402EC010A97A00    No     Off     Off     Off     Off     OK

cli% setcage id 1 cage3
cli% showcage -enclosure
Cage ------WWN------- Locate FailInd WarnInd FailReq WarnReq -State-
0 51402EC010A0EC80    No     Off     Off     Off     Off     OK
1 51402EC010A97A00    No     Off     Off     Off     Off     OK
2 51402EC010A90F80    No     Off     Off     Off     Off     OK
```

The following example demonstrates how to set the ID of a cage:

```
cli% showcage
        Id Name  LoopA Pos.A LoopB Pos.B Drives Temp  RevA RevB Model FormFactor
         0 cage0 0:1:1     0 1:1:1     0     10 27-31 4082 4082 DCN2  SFF
         1 cage1 0:1:1     1 1:1:1     1     10 32-34 4082 4082 DCS8  SFF
         3 cage3 0:1:2     1 1:1:2     0     10 32-34 4082 4082 DCS8  SFF
         4 cage4 0:1:2     0 1:1:2     1     16 37-42 4082 4082 DCS7  LFF

cli% setcage id 1 cage4
        The new ID specified is already in use by another cage.
        Please specify a new ID which is not already in use.

cli% setcage id 2 cage4
cli% showcage
        Id Name  LoopA Pos.A LoopB Pos.B Drives Temp  RevA RevB Model FormFactor
         0 cage0 0:1:1     0 1:1:1     0     10 27-31 4082 4082 DCN2  SFF
         1 cage1 0:1:1     1 1:1:1     1     10 32-34 4082 4082 DCS8  SFF
         2 cage2 0:1:2     0 1:1:2     1     16 37-42 4082 4082 DCS7  LFF
         3 cage3 0:1:2     1 1:1:2     0     10 32-34 4082 4082 DCS8  SFF
```

setcim

Syntax

`setcim [options]`

Description

Sets properties of the CIM server, including options to enable/disable the HTTP and HTTPS ports for the CIM server. `setcim` allows a user to enable/disable the SLP port. The command also sets the CIM server policy.

Options

`-f`

Forces the operation of the `setcim` command, bypassing the typical confirmation message.

Requires at least one of the following options:

-slp {enable | disable}

Enables or disables the SLP port 427.

-http {enable | disable}

Enables or disables the HTTP port 5988.

-https {enable | disable}

Enables or disables the HTTPS port 5989.

-pol <policy>[,<policy>...]

Sets the cim server policy:

- `replica_entity`—complies with SMI-S standard for usage of Replication Entity objects in associations. This is the default policy setting.
- `no_replica_entity`—does not comply with SMI-S standard for Replication Entity usage. Use only as directed by HPE support personnel or Release Notes.
- `one_hwid_per_view`—calling exposePaths with multiple initiatorPortIDs to create new view will result in the creation of multiple SCSCIProtocolControllers (SPC), one StorageHardwareID per SPC. Multiple hosts will be created each containing one FC WWN or iscsiname. This is the default policy setting.
- `no_one_hwid_per_view`—calling exposePaths with multiple initiatorPortIDs to create new view will result in the creation of only one SCSCIProtocolController (SPC) that contains all the StorageHardwareIDs. One host will be created that contains all the FC WWNs or iscsinames.
- `use_pegasus_interop_namespace`—use the pegasus defined interop namespace root/PG_interop. This is the default policy setting.
- `no_use_pegasus_interop_namespace`—use the SMI-S conformant interop namespace root/interop.
- `tls_strict`—only TLS connections using TLS 1.2 with secure ciphers will be accepted if HTTPS is enabled.
- `no_tls_strict`—TLS connections using TLS 1.0 - 1.2 will be accepted if HTTPS is enabled. This is the default policy setting.

Specifiers

None.

Authority

- Super, Service
- Any role granted the `cim_set` right

Usage

- Access to all domains is required to run this command.
- You cannot disable both of the HTTP and HTTPS ports.
- When the CIM server is active, a warning message will be prompted to inform you of the current status of the CIM server and asks for the confirmation to continue or not. The `-f` option forces the action without a warning message.

Examples

The following example disables the HTTPS ports:

```
cli% setcim -https disable
Warning: The CIM server is active and will restart.
Are you sure you want to continue?
select q=quit y=yes n=no: y
```

The following example enables the HTTPS port:

```
cli% setcim -https enable
Warning: The CIM server is active and will restart.
Are you sure you want to continue?
select q=quit y=yes n=no: y
```

The following example disables the HTTP port and enables the HTTPS port:

```
cli% setcim -http disable -https enable
WARNING: CIM server is active and it will restart.
Are you sure you want to continue?
select q=quit y=yes n=no: y
```

The following example sets the `no_use_pegasus_interop_namespace` policy:

```
cli% setcim -pol no_use_pegasus_interop_namespace
WARNING: CIM server is active and it will restart.
Are you sure you want to continue?
select q=quit y=yes n=no: y
```

The following example sets the `replica_entity` policy:

```
cli% setcim -pol replica_entity
```

setclienv

Syntax

`setclienv <parameter> <value>`

Description

The `setclienv` command sets the CLI environment parameters.

Options

None.

Specifiers

The specifiers include the parameters and values to which the parameters should be set. Valid parameters and their values are as follows:

<parameter>	<values>
currentdomain	Enter one of the following: • The name of the domain that you wish to set as the working domain for the current CLI session. • -unset to set no current domain.
listdom	Enter one of the following: • 0 - (Default) Do not include the domain column in the output. • 1 - Include domain column where relevant.
csvtable	Enter one of the following: • 0 - (Default) Normal table printing format. • 1 - Comma Separated Values (CSV) format.
nohdtot	Enter one of the following: • 0 - (Default) Show header and total lines. • 1 - Does not show the header and total lines.
matchbulkobjs	Enter one of the following: • 0 - (Default) Does not perform operation on bulk objects. • 1 - Perform operations on bulk objects.
hafter	<nlines> - Specifies the number of lines of data to display before an output header is displayed. If <nlines> is 10 or more, print the header after every <nlines> of data. If <nlines> is less than 10, print the header only at the beginning.
histstatnum	<nlines> - For the hist and stat commands, specifies the number of lines of data to display as soon as the data is available. Default is 5000 if histstatnum is not set, or set to 0. The data sorting will be within the <histstatnum> rows.
editor	This parameter is only supported when connected via SSH. Specifies the command line editing mode. Enter one of the following: • emacs - (Default) Use emacs-style line editing. • vi - Use vi-style line editing.

Authority

- Super, Service, Edit, Browse
- Any role granted the clienv_set right

Usage

This command is only available when you are using a CLI shell or SSH.

Examples

In the following example, the CLI environment is set to display domains information:

```
cli% setclienv listdom 1
```

setcorequest

Syntax

```
setcorequest queuesize <size>
setcorequest {approve|deny} [-comment <comment>] <requestid>
setcorequest {execute|cancel} <requestid>
```

Description

The `setcorequest` command is used to approve, deny, execute, or cancel a dual sign-on request or modify the queue size.

Options

`-comment <comment>`

Specifies the comment to be included in the request. The maximum size of `<comment>` is 255 characters.

Specifiers

`<requestid>`

Specifies the request ID.

Subcommands

`approve`

Approves a pending dual sign-on request.

`deny`

Rejects a pending dual sign-on request.

`execute`

Executes an approved dual sign-on request.

`cancel`

Cancels a pending dual sign-on request.

`queuesize`

Modifies the requests queue size.

Authority

- CO
- Any role granted the `corequest_set` right can approve/deny a request or modify the maximum queue size
- Any role granted the `cotask_set` right can execute or cancel a request

Usage

- Only approved requests can be executed.
- Only the requests having approval state as "pending" can be canceled.

Examples

The following example shows how to approve a request with id 12345:

```
cli% setcorequest approve 12345
```

The following example shows how to deny a request with id 12345:

```
cli% setcorequest deny 12345
```

The following example shows how to modify queue size to 25:

```
cli% setcorequest queuesize 25
```

The following example shows how to execute an approved request with id 12345:

```
cli% setcorequest execute 12345
```

The following example shows how to cancel a pending request with id 12345:

```
cli% setcorequest cancel 12345
```

setcpg

Syntax

```
setcpg [options <arg>] <CPG_name>
```

Description

The `setcpg` command modifies existing Common Provisioning Groups (CPG).

Options

`-sa <LD_name>...`

Specifies that existing logical disks are added to the CPG and are used for snapshot admin (SA) space allocation. The <LD_name> argument can be repeated to specify multiple logical disks.

This option is deprecated and will be removed in a subsequent release.

`-sd <LD_name>..`

Specifies that existing logical disks are added to the CPG and are used for snapshot data (SD) space allocation. The <LD_name> argument can be repeated to specify multiple logical disks.

This option is deprecated and will be removed in a subsequent release.

`-f`

Forces the command. The command completes the process without prompting for confirmation.

`-aw <percent>`

Specifies the percentage of used snapshot administration or snapshot data space that results in a warning alert. A percent value of 0 disables the warning alert generation. The default is 0.

This option is deprecated and will be removed in a subsequent release.

`-sdgs <size>[g|G|t|T]`

Specifies the growth increment, the amount of logical disk storage created on each auto-grow operation. The default growth increment may vary according to the number of controller nodes in the system. If <size> is non-zero it must be 8G or bigger. The size can be specified in MB (default) or GB (using g or G) or TB (using t or T). A size of 0 disables the auto-grow feature. The following table displays the default and minimum growth increments per number of nodes:

Table 6: Growth Increment Per Number of Nodes

Number of Nodes	Default	Minimum
1-2	32 GB	8 GB
3-4	64 GB	16 GB
5-6	96 GB	24 GB
7-8	128 GB	32 GB

-sdgl <size>[g|G|t|T]

Specifies that the auto-grow operation is limited to the specified storage amount. The storage amount can be specified in MB (default) or GB (using g or G) or TB (using t or T). A size of 0 (default) means no limit is enforced. To disable auto-grow, set the limit to 1.

-sdgw <size>[g|G|t|T]

Specifies that the threshold of used logical disk space, when exceeded, results in a warning alert. The size can be specified in MB (default) or GB (using g or G) or TB (using t or T). A size of 0 (default) means no warning limit is enforced. To set the warning for any used space, set the limit to 1.

-t <RAID_type>

Specifies the RAID type of the logical disk: r0 for RAID-0, r1 for RAID-1, r5 for RAID-5, or r6 for RAID-6. If no RAID type is specified, then the default is r6.

-ssz <size_number_chunklets>

Specifies the set size in terms of chunklets. The default depends on the RAID type specified: 2 for RAID-1, 4 for RAID-5, and 8 for RAID-6.

-rs <size>

Specifies the number of sets in a row. The <size> is a positive integer.

If not specified, no row limit is imposed.

-ss <size_KB>

Specifies the step size from 32 KB to 512 KB. The step size should be a power of 2 and a multiple of 32. The default value depends on raid type and device type used. If no value is entered and FC or NL drives are used, the step size defaults to 256 KB for RAID-0 and RAID-1, and 128 KB for RAID-5. If SSD drives are used, the step size defaults to 32 KB for RAID-0 and RAID-1, and 64 KB for RAID-5. For RAID-6, the default is a function of the set size.

-ha port|cage|mag

Specifies that the layout must support the failure of one port pair, one cage, or one drive magazine (mag). This option has no meaning for RAID-0. The default is cage availability.

-ch first|last

Specifies the chunklet location characteristics: either first (attempt to use the lowest numbered available chunklets) or last (attempt to use the highest numbered available chunklets). If no argument is specified, the default characteristic is first.

-p <pattern>

Specifies a pattern for candidate disks. Patterns are used to select disks that are used for creating logical disks. If no pattern is specified, the option defaults to Fast Class (FC) disks. If specified multiple times, each instance of the specified pattern adds additional candidate disks that match the pattern. The -devtype pattern cannot be used to mix Nearline (NL), FC, and Solid State Drive (SSD) drives. An item is specified as an integer, a comma-separated list of integers, or a range of integers specified from low to high.

The following arguments can be specified as patterns for this option:

An item is specified as an integer, a comma-separated list of integers, or a range of integers specified from low to high.

-nd <item>

Specifies one or more nodes. Nodes are identified by one or more integers (item). Multiple nodes are separated with a single comma (e.g. 1,2,3). A range of nodes is separated with a hyphen (e.g. 0-7). The primary path of the disks must be on the specified node(s).

-st <item>

Specifies one or more PCI slots. Slots are identified by one or more integers (item). Multiple slots are separated with a single comma (e.g. 1,2,3). A range of slots is separated with a hyphen (e.g. 0-7). The primary path of the disks must be on the specified PCI slot(s).

-pt <item>

Specifies one or more ports. Ports are identified by one or more integers (item). Multiple ports are separated with a single comma (e.g. 1,2,3). A range of ports is separated with a hyphen (e.g. 0-4). The primary path of the disks must be on the specified port(s).

-cg <item>

Specifies one or more drive cages. Drive cages are identified by one or more integers (item). Multiple drive cages are separated with a single comma (e.g. 1,2,3). A range of drive cages is separated with a hyphen (e.g. 0-3). The specified drive cage(s) must contain disks.

-mg <item>

Specifies one or more drive magazines. The "1." or "0." displayed in the CagePos column of showpd output indicating the side of the cage is omitted when using the -mg option. Drive magazines are identified by one or more integers (item). Multiple drive magazines are separated with a single comma (e.g. 1,2,3). A range of drive magazines is separated with a hyphen(e.g. 0-7). The specified drive magazine(s) must contain disks.

-pn <item>

Specifies one or more disk positions within a drive magazine. Disk positions are identified by one or more integers (item). Multiple disk positions are separated with a single comma (e.g. 1,2,3). A range of disk positions is separated with a hyphen(e.g. 0-3). The specified position(s) must contain disks.

-dk <item>

Specifies one or more physical disks. Disks are identified by one or more integers (item). Multiple disks are separated with a single comma (e.g. 1,2,3). A range of disks is separated with a hyphen (e.g. 0-3). Disks must match the specified ID(s).

-tc_gt <number>

Specifies that physical disks with total chunklets greater than the number specified be selected.

-tc_lt <number>

Specifies that physical disks with total chunklets less than the number specified be selected.

-fc_gt <number>

Specifies that physical disks with free chunklets greater than the number specified be selected.

-fc_lt <number>

Specifies that physical disks with free chunklets less than the number specified be selected.

-devid <model>

Specifies that physical disks identified by their models be selected. Models can be specified in a comma-separated list.

Models can be displayed by issuing the "`showpd -i`" command.

-devtype <type>

Specifies that physical disks must have the specified device type (FC for Fast Class, NL for Nearline, SSD for Solid State Drive) to be used. Device types can be displayed by issuing the "`showpd`" command. If it is not specified, the default device type is FC.

-rpm <number>

Disks must be of the specified speed. Device speeds are shown in the RPM column of the `showpd` command. For FC and NL drives, the number corresponds to both a performance measure and actual rotational speed. For SSD and SCM drives, the field is "N/A."

Disks that satisfy all of the specified characteristics are used.

For example -p -fc_gt 60 -fc_lt 230 -nd 2 specifies all the disks that have greater than 60 and less than 230 free chunklets and that are connected to node 2 through their primary path.

-sax <LD_name>[,<LD_name>...]

Specifies that the logical disk, as identified with the <LD_name> argument, used for snapshot administration space allocation be removed.

The <LD_name> argument can be repeated to specify multiple logical disks-sdx <LD_name>[,<LD_name>...]

Specifies that the logical disk, as identified with the <LD_name> argument, used for snapshot data space allocation be removed. The <LD_name> argument can be repeated to specify multiple logical disks.

-name <newname>

Specifies the name of the Common Provisioning Group (CPG) to be modified to. <newname> can be up to 31 characters in length.

Specifiers

<CPG_name>

Specifies the name of the common provisioning group being modified.

Authority

- Super, Service, Edit
- Any role granted the `cpg_set` right.

Usage

- Access to all domains is required to run this command.
- With this command, you can change the device type of a CPG from logical disks of one device type to logical disks of another device type (device types are Fast Class (FC), Nearline (NL), Solid State Drive (SSD).
- This implies that, within a CPG, one can have logical disks of type FC, type NL, and type SSD. However, this is only permitted so that, if a user wants to change the type of a CPG from FC to NL or SSD, they can first change the new logical disk creation characteristics using `setcpg` and then use region moves to change the device types of existing logical disks. For this same reason, users are permitted to add logical disks of a different device type to a CPG.

- This command sets the parameters that will be used for future grows. Unless the only growth parameters (-sdgs, -sdgw or -sdgl) are being changed, this command will OVERWRITE all previously set parameters.
- The options -sdgs, -sdgl and -sdgw control the auto logical disk creation for the Common Provisioning Group's snapshot data regions. Auto logical disk creation occurs when the amount of free logical disk space falls below the specified grow (enlarge) size setting options (-sdgs, -sdgl).
- If auto-grow is enabled, new logical disks will be created automatically in addition to any logical disks specified with -sa or -sd options.
- To remove the pattern, enter "" after the <pattern>.
- To move a CPG from one domain to another, use the **movetodomain** command.

Examples

The following example displays the modification of the auto-growth parameters for common provisioning group cpg1:

```
cli% setcpg -sdgs 16 -sdgl 48 -sdgw 36 cpg1
```

The following example removes the -cg pattern:

```
cli% setcpg -p -cg "" cpg1
```

The following example shows that this command sets the parameter that will be used for future grows and it will OVERWRITE all previously set parameters. (use "showcpg -sdg" to verify):

```
cli% createcpg -sdgl 48 -sdgw 36 cpg1
cli% showcpg -sdg cpg*
        ------(MB)------
        Id Name Warn Limit  Grow Args
         2 cpg1   36    48 32768 -p -devtype FC
cli% setcpg -p -nd 0,1 -devtype FC cpg1
cli% showcpg -sdg cpg*
        ------(MB)------
        Id Name Warn Limit  Grow Args
         2 cpg1   36    48 32768 -p -nd 0,1 -devtype FC
```

setdate

Syntax

The syntax for the setdate command can be one of the following:

```
setdate <MMDDhhmm>[[<CC>]<YY>][<.ss>]
setdate -tzlist [group]
setdate -tz <tzname>
```

Description

The setdate command allows you to set the system time and date on all nodes.

Options

-tzlist [group]

Displays a timezone within a group, if a group is specified. If a group is not specified, displays a list of valid groups.

-tz <tzname>

Sets the timezone on all nodes. The option must have a valid tzname from the list provided with the -tzlist command.

-f

Specifies that the command is forced. If this option is not used, the command requires confirmation that proposed time is correct before proceeding to change the time.

Specifiers

<MMDD>

Specifies the month (MM) and day (DD).

<hhmm>

Specifies the hour (hh) and minute (mm) on a 24-hour clock.

<CC>

Specifies the century (CC) and cannot be used unless a year is specified (YY). This specifier is not required.

<YY>

Specifies a year (YY). This specifier is not required.

<.ss>

Specifies seconds (ss). This specifier is not required.

Authority

- Super, Service
- Any role granted the `date_set` right

Usage

- Access to all domains is required to run this command.
- Specifiers must be provided in one of the combinations listed as follows when issuing the `setdate` command:
  - `MMDDhhmm`
  - `MMDDhhmmYY`
  - `MMDDhhmmCCYY`
  - `MMDDhhmm.ss`
  - `MMDDhhmmYY.ss`
  - `MMDDhhmmCCYY.ss`
- Check node dates by issuing the **`showdate`** command.
- To set the date using Network Time Protocol (NTP), use the "setnet ntp" command to specify an NTP server. The date will not change immediately, as NTP corrects the time slowly to avoid issues.

Examples

The following example displays the current date on the node:

```
cli% showdate
```

The following example displays the timezones with the -tzlist option:

```
cli% setdate -tzlist
```

The following example narrows down the list to the required timezone of Etc:

```
cli% setdate -tzlist Etc
```

The following example shows the timezone being set:

```
cli% setdate -tz Etc/GMT
```

The following example verifies the timezone is set to the required setting:

```
cli% showdate
```

setdomain

Syntax

```
setdomain [options] <domain_name>
```

Description

The `setdomain` command sets the parameters and modifies the properties of a domain.

Options

`-name <name>`

Changes the name of the domain.

`-comment <comment>`

Specifies comments or additional information for the domain. The comment can be up to 511 characters long and must be enclosed in quotation marks. Unprintable characters are not allowed within the <comment> specifier.

`-vvretentiontimemax <time>{h|H|d|D}`

Specifies the maximum value that can be set for the retention time of a volume in this domain. <time> is a positive integer value and in the range of 0 - 43,800 hours (1825 days). Time can be specified in days or hours providing either the 'd' or 'D' for day and 'h' or 'H' for hours following the entered time value.

To remove the maximum volume retention time for the domain, enter '-vvretentiontimemax ""'. As a result, the maximum volume retention time for the system is used instead.

To disable setting the volume retention time in the domain, enter 0 for <time>.

Specifiers

`<domain_name>`

Indicates the name of the domain.

Authority

- Super
- Any role granted the `domain_set` right

Usage

- Access to all domains is required to run this command.
- For moving objects into a domain, use the **`movetodomain`** command.

Examples

In the following example, the name of a domain named Domain1 is changed to DomainX:

```
cli% setdomain -name DomainX Domain1
```

The following example displays the addition of a comment to the domain Engineering:

```
cli% setdomain -comment "This is a comment for engineering." Engineering
```

setdomainset

Syntax

```
setdomainset [options <arg>] <setname>
```

Description

The `setdomainset` command sets the parameters and modifies the properties of a domain set.

Options

`-comment <comment>`

Specifies any comment or additional information for the set. The comment can be up to 255 characters long. Unprintable characters are not allowed.

`-name <newname>`

Specifies a new name for the domain set, using up to 27 characters in length.

Specifiers

`<setname>`

Specifies the name of the domain set to modify.

Authority

- Super, Edit
- Any role granted the `domainset_annotate` right can only annotate the comment of a domain set
- Any role granted the `domainset_set` right can set any domain set property

Usage

Access to all domains is required to run this command.

Examples

The following example renames a set from test1 to test2:

```
cli% setdomainset -name test2 test1
```

The following example changes the comment on set test2:

```
cli% setdomainset -comment "This used to be set test1" test2
```

setflashcache

Syntax

```
setflashcache {enable|disable} {vvset:<name|pattern>|sys:all} ...
setflashcache {clear} {sys:all}
```

Description

The `setflashcache` command allows you to set the policy of the flash cache for virtual volumes. The policy is set by using virtual volume sets (vvset). The sys: all is used to enable the policy on all virtual volumes in the system.

Options

None.

Specifiers

`vvset:<name|pattern>`

<name> refers to the target object name as listed in the `showvvset` command. Pattern is glob-style (shell-style) patterns. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.)

`sys:all`

The policy is applied to all virtual volumes.

Subcommands

`enable`

Turns on the flash cache policy for the target object.

`disable`

Turns off flash cache policy for the target object.

`clear`

Turns off policy and can only be issued against the sys: all target.

Authority

- Super, Edit
- Any role granted the `flashcache_set` right

Usage

- It is possible that a virtual volume can be in multiple virtual volume sets and that flash cache policy for those sets may contradict each other.
- In these cases, the "enable" setting will take precedence. This is also true when the policy is enabled at the system level.
- Issuing disable for the sys: all will disable the policy at the system level.
- Issuing clear for the sys: all will turn off the global policy for virtual volumes but will leave the policy for any vvset that is in place.

Examples

The following example will enable flash cache policy for vvset1:

```
cli% setflashcache enable vvset:vvset1
```

The following example will enable flash cache policy for vvset1 and vvset2:

```
cli% setflashcache enable vvset:vvset1 vvset:vvset2
```

The following example will disable flash cache policy for vvset1 and vvset1:

```
cli% setflashcache disable vvset:vvset1 vvset:vvset2
```

The following example will enable the flash cache policy for all volumes in the system:

```
cli% setflashcache enable sys:all
```

The following example will clear the sys: all will turn off the global policy for virtual volumes:

```
cli% setflashcache clear sys:all
```

The following example will disable the policy for all virtual volume sets:

```
cli% setflashcache disable sys:all
```

setfpg

Syntax

```
setfpg [options] <fpgname>
```

Description

Enable and disable various properties associated with a File Provisioning Group.

Options

-f

Silences system confirmation of the command request.

-comment <comment string>

Specifies any addition textual information.

-rmcomment

Clears the comment string.

-activate

Makes the File Provisioning Group available.

-deactivate

Makes the File Provisioning Group unavailable.

-primarynode <nodeid>

Specifies the primary node to which the FPG will be assigned. Appropriate `<nodeid>` values are defined as those on which File Persona has been enabled.

-failover

Specifies that the FPG must be failed over to its alternate node. If it has previously failed over to the secondary, this will cause it to fail back to the primary node. Will fail if a graceful failover is not possible.

-forced

In the event of failure to failover, this will attempt a forced failover.

-upgrade

Upgrades the On-disk version.

-unmap {enable | disable}

Enables or disables Thin Reclamation for an FPG. Thin Reclamation should not be enabled for FPGs that use fully provisioned volumes.

-defaultcpg <cpgname>

Sets the CPG for future VV creation, and is reflected in `showfpg`. Not available with other options.

-version <version>

Specifies the on-disk upgrade version of FPG. If not specified, FPG is upgraded to the latest supported version.

-rcopygroup <rcopy_group>

Removes the FPG from its current Remote Copy group and add it to the one specified. With no group specified, removes the FPG from its current group.

Specifiers

<fpgname>

The fsname is the name of the file provisioning group to be modified.

-defaultcpg <cpgname>

Sets the CPG for future VV creation. This will be reflected in **showfpg**. It is not available with other options.

Authority

- Super, Edit
- Any role granted `fpg_set` right

Usage

- Access to all domains is required to run this command.
- The `-primarynode` and `-failover` options are mutually exclusive.
- When assigning primary nodes, the secondary node will be implicit as a couplet pair [0,1] [2,3] [4,5] [6,7]. This action will fail if the graceful failover is not possible.
- The `-failover` and `-primarynode` options will result in temporary unavailability of the Virtual File Servers associated with the FPG being migrated, and also the unavailability of any associated shares. An implicit `-deactivate` and `-activate` process is undertaken during a migration to the alternate node.
- The default CPG must be changed after a **tunevv** of any FPG VVs.

Examples

- Showing the assignment of `examplefpg` to node 1:

```
cli% setfpg -primarynode 1 examplefpg
```

- Setting the comment for `examplefpg`:

```
cli% setfpg -comment "Example comment string" examplefpg
```

- Activating `examplefpg`:

```
cli% setfpg -activate examplefpg
```

- Upgrading the current on disk version of `examplefpg`:

```
cli% setfpg -upgrade examplefpg
```

- Upgrading the current on disk version of `examplefpg` to version 12.1:

```
cli% setfpg -upgrade -version 12.1 examplefpg
```

setfs

Syntax

The syntax for the `setfs` command can be one of the following:

```
setfs ad [-passwd <password>] <user> <domain>
setfs ad -leave [-f]
setfs ad [-f] [[-disable] [-passwd <password>] -username <user>] -leave
```

```
setfs auth -clearcache

setfs auth [-f] <provider> ...

setfs auth [-f] -enabledomainprefix {true|false}

setfs auth [-f] -encryptiontype {aes|aes-rc4|rc4-des|all}

setfs bond <bond_mode>

setfs dns {<ipaddress-list> [<suffix-list>]|-delete}

setfs gw {<ipaddress>|-delete}

setfs idmap <nfsv4domain>

setfs idmap -delete [-f]

setfs ldap [-passwd <binddnpwd>] [-schema <schema>] [{-usessl|-usetls} {-
certfile <file>|-certdata <data>} -certcn <certcn>] [-cloneservers
[+|-]<cloneip1|cloneFQDN1>,...] <server> <binddn> <searchbase> <netbios>

setfs ldap -delete [-f]

setfs mtu <mtu_size>

setfs nfs -servergid {true|false}

setfs nodeip -ipaddress <ipaddress> [-subnetoverlap] -subnet <subnet> [-
vlantag <vlanid>] <nodeid>

setfs nodeip -delete <nodeid>

setfs obj [-keepalive {true|false}] [-timeout <secs>] [-maxclients <num>] [-
rblksize <size>] [-wblksize <size>] [-f]

setfs pdc [-f] -enable {true|false}

setfs pdc [-f] -delete

setfs pdc -hostname <hostname> -ipaddress <ipaddress> -domain <domain> [-site
<site>]

setfs rfc2307 [-f] {enable|disable}

setfs smb [-f] [-enableoplocks {true|false}] [-signingenabled {true|false}]
[-signingrequired {true|false}] [-ignorewritethroughrequests {true|false}] [-
supportpersistenthandles {true|false}] [-smb3dialectenable {true|false}] [-
enablesmb1ad {true|false}] [-enablesmb2ad {true|false}] [-enablesmbleases
{true|false}] [-enabledirleases {true|false}] [-enablesmb1 {true|false}] [-
enablesmb2 {true|false}]

setfs usermap [-f] -enable {true|false}

setfs usermap -importconf <file_path_on_client>

setfs usermap -export {users|groups} -provider <provider>
```

Description

Configures the various components through subcommands.

Options

The following options are for the **ad** subcommand:

`-username <user>`

Specifies the authorized user to access Active Directory domain controller. Specify the option with `-leave` option to have the changes to the computer account reflected on the domain controller.

-passwd <password>

Specifies the password of the authorized user. If this is not specified, the user will be prompted for a cleartext password.

-leave

Disconnects File Persona nodes from Active Directory domain. If this option is not specified, the subcommand joins File Persona nodes to Active Directory domain.

-disable

Disables AD Computer account. Specify the option with credentials using `-username` option.

-f

Specifies that the operation is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

The following options are for the **auth** subcommand:

<provider> ...

Specifies the provider order. Values must be: Ldap, Local, or ActiveDirectory. One provider must be Local.

-f

Specifies that the operation is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

-clearcache

Clears name cache of the cluster. This option cannot be used with any other option. This option can be used if there is any modifications to the external authentication provider such as user deletion or removal of user from groups.

The following options are for the **pdc** subcommand:

-enable {true | false}

Enable or disable the PDC configuration. Default is disable. Do not specify the option with any other options except `-f`.

-delete

Deletes all the configured PDCs. Do not specify the option with any other options except `-f`.

-f

Specifies that the operation is forced. If not specified, the command requires confirmation before proceeding with its operation.

-hostname

Specifies the host name of the Windows domain controller.

-ipaddress

Specifies the IP address of the Windows domain controller.

-domain

Specifies the domain name to which the PDC being added belongs to or resides in.

-site

Specifies the site name of the Windows domain controller. If not specified, the default site will be used.

The following options are for the **dns** subcommand:

<ipaddress-list>

Specifies the DNS addresses. The value for ipaddress-list is a comma-separated list.

<suffix-list>

Specifies the DNS suffixes. The value for suffix-list is a comma-separated list.

-delete

Clears the DNS settings (addresses and suffixes).

The following options are for the **gw** subcommand:

<ipaddress>

Specifies the IP Address for the Gateway.

-delete

Removes the gateway to allow reconfiguration of node addresses to another subnet.

The following options are for the **obj** subcommand:

-keepalive {true|false}

Specifies whether persistent connections must be allowed or not for the default Object profile. The default value is "true".

-timeout <secs>

Specifies the timeout value in seconds for the persistent connections for the default Object profile. The valid range is from 1 to 2592000. The default value is 5 seconds.

-maxclients <num>

Specifies the maximum number of simultaneous connections for the default Object profile. The valid range is from 1 to 128. The default value is 50 connections.

-rblksize <size>

Specifies the socket read block size for the default Object profile. The size can be specified with a "K" or "M" to indicate the unit in kilobyte or megabyte (for example, `-rblksize 10K`). The valid range is from 8K to 2048M bytes. The default value is 8K.

-wblksize <size>

Specifies the file write block size for the default Object profile. The size can be specified with a "K" or "M" to indicate the unit in kilobyte or megabyte (for example, `-wblksize 1M`). The valid range is from 8K to 2047M bytes. The default value is 8K.

-f

Specifies that the operation is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

The following options are for the **ldap** subcommand:

-passwd <binddnpwd>

Specifies the password associated with the bind DN `<binddn>`.

-schema {posix|samba}

Specifies the schema used to create user and group accounts on the LDAP server. If this is not specified, the default is "`posix`".

-usetls

Indicates to use TLS connection between File Persona nodes and the LDAP server.

-usessl

Indicates to use SSL connection between File Persona nodes and the LDAP server.

-certcn <certcn>

Specifies the CN or common name used when the certificate is generated, which has to be the fully qualified hostname of the LDAP server. When option `-usetls` or `-usessl` is specified, this option must be specified.

-certfile <file>

Specifies the certificate file name on the client machine. The command imports certificate information from this file. When option `-usetls` or `-usessl` is specified, this option or `-certdata` must be specified.

-certdata <data>

Specifies the certificate data. When option `-usetls` or `-usessl` is specified, this option or `-certfile` must be specified.

-delete

Deletes LDAP configuration for File Persona.

-f

Specifies that the operation is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

-cloneservers [+|-]<cloneip1|cloneFQDN1>,...

Specifies the list of LDAP clone server IP addresses or fully qualified domain names (FQDN). For Simple connection, IP addresses or FQDNs or both can be specified. For SSL/TLS, only FQDNs are allowed. If <cloneip|cloneFQDN> has a prefix: + add IP addresses, FQDNs to the existing cloneservers list. - remove IP addresses, FQDNs from the existing cloneservers list. If no prefix, the list will be used as the new list.

The following options are for the **nodeip** subcommand:

-ipaddress <ipaddress>

Specifies the IP Address to be used for the node.

-subnet <subnet>

Specifies the subnet mask to be used for the node.

-vlantag <vlanid>

Specifies the VLAN tag used for the node. Defaults to 0.

-delete

Deletes a specified node IP.

-subnetoverlap

Allow overlapping of subnets. If not specified, overlapping subnets are not allowed.

NOTE: This is not a typical configuration. If not configured properly, the configuration can cause data unavailability.

-f

Specifies that the operation is forced. If not specified, the command requires confirmation before proceeding with its operation.

The following options are for the **idmap** subcommand:

-delete

Deletes NFSv4 domain name.

-f

Specifies that the operation is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

The following options are for the **rfc2307** subcommand:

{disable|enable}

Specifies if RFC2307 must be disabled or enabled for the "ActiveDirectory" provider. If enabled, the AD provider expects UIDs and GIDs to be found in the Active Directory and not formulated using the SID as is the default behavior. This allows users to customize UIDs and GIDs. If rfc2307 being enabled and the user does not have a configured UID/GID, the user is not given write access, even if granted through an ACL. If this subcommand is not run, the default is "`disable`".

-f

Specifies that the operation is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

The following options are for the **smb** subcommand:

-enableoplocks {true|false}

Enables or disables Oplocks.

-signingenabled {true|false}

SMB Signing can be either true or false.

-signingrequired {true|false}

If signing is required, set this to true else set this to false.

-ignorewritethroughrequests {true|false}

Ignores Write-through Requests if this option is set to true.

-supportpersistenthandles {true|false}

Enables or disables support for Persistent handles.

-smb3dialectenable {true|false}

Enables or disables SMB3 Dialect.

-enablesmb1ad {true|false}

Enables or disables SMB1 connections to Active Directory servers only. Default value for option "enablesmb1ad" is false.

-enablesmb2ad {true|false}

Enables or disables SMB2 connections to Active Directory servers only.

-enablesmbleases {true|false}

SMB leases can be enabled or disabled by setting this option.

-enabledirleases {true|false}

Directory leases are enabled or disabled.

-enablesmb1 {true|false}

Enables or disables SMB1 globally for client connections. Default value for option "enablesmb1" is false.

-enablesmb2 {true|false}

Enables or disables SMB2 globally for client connections.

-f

Suppresses the confirmation from the user.

The following options are for the **nfs** subcommand:

-servergid {true|false}

Enables or disables Server GID.

The following options are for the **usermap** subcommand:

-enable {true|false}

Enables or disables the user mapping. Default is disabled. This option cannot be used with any other option except `-f`.

-importconf <filepath_on_client>

Imports the user-mapping configuration from the file specified on client. This option cannot be used with any other option except `-f`.

-export {users|groups}

Exports the users/groups entries to a file. If this option is specified, `-provider` must also be specified. This option generates a task id. See "`showfs -usermap`" for the command to copy the exports file to the client.

-provider {Local|Ldap|ActiveDirectory}

Specifies the type of `<provider>` to be used to export users/groups entries. This must be one of the configured authentication providers. If this option is specified, `-export` must also be specified.

-f

Suppresses the warning message. If this option is not used, the command may require a confirmation before proceeding with its operation.

Specifiers

The following specifiers are for the **ad** subcommand:

<user>

Specifies the authorized user to access Active Directory domain controller.

<domain>

Specifies the domain name of Active Directory which File Persona nodes are to join.

The following specifiers are for the **ldap** subcommand:

<server>

Specifies the fully qualified hostname or IPv4 address of the LDAP server you want to configure. If nonstandard port is used (that is, not 389 or 636), the port number must be specified with the server in the format of `<server>:<port>`.

<binddn>

Specifies the distinguished name (that is, DN), which will be used to authenticate to the LDAP server to read data. This account must have privileges to read the subtree specified by `<searchbase>`. Write permissions are not required.

<searchbase>

Specifies the LDAP base used to search user/group accounts.

<netbios>

Netbios specifies the name of the LDAP Domain. It can be up to 15 alphanumeric characters with no spaces. The name must be unique on the network. To access an SMB share, you specify "`<netbios>\<ldap-username>`" as the user name.

The following specifier is for the **idmap** subcommand:

<nfsv4domain>

Specifies NFSv4 domain name for ID mapping. For NFSv4 to work correctly, the NFSv4 client and the NFSv4 server must be in the same NFSv4 domain.

The following specifier is for the **mtu** subcommand:

<mtu_size>

Specifies the size of the MTU for all the File Persona nodes.

The following specifier is for the **bond** subcommand:

<bond_mode>

Specifies the bond mode for all the File Persona nodes.

Subcommands

ad

Joins or leaves Active Directory domain.

auth

Sets user authentication settings.

dns

Sets or clears DNS settings.

gw

Sets or deletes Gateway IP address.

idmap

Sets NFSv4 domain name for NFS ID mapping.

ldap

Configures or deletes LDAP configuration.

mtu

Sets the mtu size for all the File Persona nodes.

nodeip

Sets or deletes Node network settings.

obj

Sets Object settings.

pdc

Sets a PDC configuration.

rfc2307

Enables or disables RFC2307 for the Active Directory provider.

smb

Allows setting of tunable parameters for SMB component. At least one parameter must be supplied. Asks for confirmation from user if they want to proceed.

usermap

Configures the user mapping.

Authority

- Super
- Any role granted `fs_set` right.

Usage

- Access to all domains is required to run this command.
- For the `obj` subcommand, Object service daemon will be restarted for the change to take effect.
- For the `auth` subcommand, HPE-SMB services will be restarted to change the stacking order or to enable/disable usage of LOCAL_CLUSTER prefix, which causes client I/O errors.
- "ActiveDirectory" must be in the provider stacking order to use "`setfs ad`". After joining Active Directory domain, the computer names for the nodes will be displayed as "`<serial number>-<node ID>`" by Windows "Active Directory Users and Computers". For example, if the CLI "`showsys`" shows ID "0" and the serial number "1234567", Windows will display the computer name "1234567-0" for the node.
- When running the `setfs ad -leave` command:

  1. Providing username and password without the `-disable` option will delete the Active Directory computer account.
  2. Providing username and password with the `-disable` option will disable the Active Directory computer account.
  3. Not specifying the username and password will keep the Active Directory computer account enabled. This option can be used to clean up File Persona Active Directory join state when File Persona cannot communicate to the Active Directory.

  When usage of LOCAL_CLUSTER prefix is enabled, local domain prefix must be supplied for the local user authentication while accessing SMB shares. To disable usage of LOCAL_CLUSTER prefix, "Local" must be at last in the auth order. When usage of LOCAL_CLUSTER prefix is disabled, an attempt to move "Local" to the "nonlast" position will fail with an error.

  The Active Directory encryption type cannot be changed with the auth subcommand when already joined to Active Directory. Leave, configure the encryption type, and join Active Directory again to apply this setting.
- For the `ldap` subcommand, if neither option `-usessl` nor `-usetls` is specified, the connection between File Persona nodes and the LDAP server will not be encrypted, and the certificate specified by `-certfile` or `-certdata` will be ignored.
- For SSL/TLS connection if user wants to add FQDNs using -cloneservers option, -certfile or -certdata must have existing LDAP servers certificates appended at start followed by new FQDN certificates. Maximum number of IPs/FQDNs including the primary server supported is 5. Option -certfile is not available with SSH access to CLI.
- "Ldap" must be in the provider stacking order to use "`setfs ldap`".
- To use NFSv4, you must run "`setfs idmap`" to set NFS domain name.
- For the `usermap` subcommand, HPE-SMB services will be restarted to enable/disable the user mapping, which causes client I/O errors.
- For the `pdc` subcommand, HPE-SMB services will be restarted to enable/disable PDC configuration, which causes disruption to other services using registry.
- The configuration file for `-importconf` can have maximum of 1024 entries in the format "domainname1\username1<operator>domainname2\username2" and the number of characters in domain name must be less than 14 and that of username must be less than 104.
- The option `-importconf <file_path_on_client>` is not available with SSH access to CLI.
- The supported operators for `-importconf` are as follows:

- **=>**: Unidirectional replace rule for static mapping. Once the "From" user is authenticated, replace the "From" user identity with the "To" user identity. If the "To" user is missing any part of the identity, if the AD provider is in rfc2307 mode and the UID or primary GID is missing the rule will report a failure when it is evaluated. If the "To" user is an LDAP user and LDAP is configured to be in POSIX schema, SID will be synthesized. If the "To" user is an AD user and AD is configured to be in unprovisioned mode, UID/GID will be synthesized. This type of mapping consolidates cross-protocol access to the same AD or LDAP account. It assures common access across protocols without duplicating accounts and group memberships across name services. Bidirectional mapping is not useful in this scenario.
- **==**: Bidirectional join/merge rule for static and dynamic mapping. Joins the native IDs from both the "From" user and the "To" user. There is no positional relationship based on the rule as it is a bidirectional rule. If an AD user logs in, the user identity includes the SID for the AD user and UID/GID of the mapped LDAP user. If an LDAP user logs in, the user identity includes UID/GID of the LDAP user and SID of the mapped AD user. By using a wildcard for both "From" and "To", instead of a specific name this rule can be used to support dynamic mapping. For example, *==* will result in any user from a provider to be mapped to another user with the same name from another provider.
- Coexistence of replace rules along with bidirectional rules in the mapping file is not a supported configuration. Bidirectional rules also require group-mapping rules. If the group names between both providers are the same and if the mapping file has a dynamic rule that would address this scenario as the dynamic mapping rule also applies to groups. However if the group names are not the same across the providers, then there has to be a specific mapping rule for the desired group names. While the mapping rule for supplemental groups may be added as needed, for a bidirectional mapping to work, the primary group has to be mapped either through a specific static rule or through dynamic rule.

- SMBv1 and SMBv2
- At any moment, both values in a key pair "enablesmb1ad/enablesmb2ad" and/or "enablesmb1/enablesmb2" cannot be false. Below invalid configurations will result in loss of capability to interact with AD servers and smb clients respectively, 1. "enablesmb1ad = false" and "enablesmb2ad = false" 2. "enablesmb1 = false" and "enablesmb2 = false"
- SMBv2 and SMBv3
- SMBv2 must be enabled to enable SMBv3. Hence the following invalid configuration will result in unavailability of SMBv3 since SMBv2 is disabled, "smb3dialectenable = true" and "enablesmb2 = false"

Examples

The following examples show usage of `setfs` command with different subcommands and respective options:

```
cli% setfs ad admin ad.example.com
cli% setfs ad -leave
cli% setfs ad -username admin -leave
cli% setfs ad -disable -passwd adminpw -username admin -leave
cli% setfs auth Ldap Local
cli% setfs auth -clearcache
cli% setfs auth -f -enabledomainprefix true
cli% setfs auth -encryptiontype aes
cli% setfs bond 1
cli% setfs dns 192.168.8.80,127.127.5.50 foo.com,bar.com
cli% setfs gw -delete
cli% setfs idmap example.com
cli% setfs obj -keepalive false
cli% setfs ldap 1.1.1.1 cn=Manager,dc=example,dc=com dc=example,dc=com netbios
cli% setfs ldap -passwd pwd343 -schema posix -usessl -certfile /home/testuser1/cert.txt
                    -certcn ldap.example.com -cloneservers +ldap2.example.com,
                    ldap3.example.comldap1.example.com
cli% setfs ldap -cloneservers -1.1.1.0,2.2.2.0,ldap1.example.com 1.1.1.1 cn=Manager,
                    dc=example,dc=com dc=example,dc=com hostname
cli% setfs ldap -cloneservers 1.1.1.0,2.2.2.0,ldap1.example.com,ldap2.example.com
                    1.1.1.1 cn=Manager,dc=example,dc=com dc=example,dc=com hostname
cli% setfs mtu 1500
cli% setfs nodeip -subnetoverlap -ipaddress 10.10.10.11-subnet 255.255.240.0 0
cli% setfs nodeip -delete 0
```

```
cli% setfs obj -keepalive false
cli% setfs pdc -enable true
cli% setfs pdc -hostname Europe.sample.com -ipaddress 10.35.155.237 -domain sample.com
cli% setfs pdc -delete
cli% setfs rfc2307 enable
cli% setfs smb -enablesmb2ad true -signingrequired true
cli% setfs smb -enablesmb1ad true -enablesmb1 true
```

The following example sets the nfs servergid:

```
cli% setfs nfs -servergid true
```

The following example enables the user mapping:

```
cli% setfs usermap -f -enable true
```

The following example imports the user-mapping configuration from the file /home/usermap/mymapping.cfg:

```
cli% setfs usermap -importconf /home/usermap/mymapping.cfg
```

The following example exports the LDAP user entries to a file:

```
cli% setfs usermap -export users -provider Ldap
```

setfsarchive

Syntax

The syntax for the `setfsarchive` command can be one of the following:

```
setfsarchive pol [-mode {enterprise|compliance} -defperiod <defaultperiod> -
minperiod <period> -maxperiod <period> -autocommperiod <period>] [-
retenvalidation {enable|disable}] [-fstore <fstorename>] [-fpg <fpgname>]
<vfs>
```

```
setfsarchive pol_inherit -inheritance {enable|disable} -fstore <fstorename>
[-fpg <fpgname>] <vfs>
```

```
setfsarchive legalhold {-set | -clear} [-basepath <basepath>] {-files
<filepath>[,<filepath>]... | -importfile <source_path> | -inputfile
<pathoffile>} -fstore <fstorename> [-fpg <fpgname>] <vfs>
```

```
setfsarchive retention [-basepath <basepath>] {-files
<filepath>[,<filepath>]... | -importfile <source_path> | -inputfile
<pathoffile>} -expdate <expirydate> -fstore <fstorename> [-fpg <fpgname>]
<vfs>
```

Description

The `setfsarchive` commands are used to set retention policy and retention properties for files in FPG at VFS or FSTORE level.

Options

<period>

When using a period option in the command, a decimal number is optionally followed by one of the following characters:

- s (seconds)
- m (minutes)
- h (hours)

- d (days)
- w (weeks)
- M (months)
- y (years)

If no character is specified, the decimal number is interpreted as seconds.

-fpg <fpgname>

Specifies the FPG that contains the VFS. This is required if VFS exists under multiple FPGs.

-fstore <fstorename>

File store name.

-mode {enterprise}

Retention mode.

-defperiod <period>

Default retention period. The valid range is in between minimum and maximum retention period set.

-minperiod <period>

Minimum retention period. The valid range is 2s - 100y.

-maxperiod <period>

Maximum retention period. The valid range is 2s - 100y.

-autocommperiod <period>

Retention auto commit period. The valid range is 5m - 1y. It can be set to 0 to disable autocommit period.

-inheritance {enable|disable}

Enables or disable the policy inheritance.

-basepath <basepath>

Base path of files for which retention settings are to be modified.

-set

Sets the legal hold for the file(s) specified.

-clear

Releases the legal hold for the file(s) specified.

-files <filepath>[,<filepath>]...

Relative path of files from the above base path for which retention settings are to be set or modified. If basepath is not specified absolute path of the files must be specified.

-inputfile <pathoffile>

Path of the input file containing list of files for which retention setting is to be changed or legal hold is to be set or released. The input file must be placed in file persona namespace (.admin file store). It should be absolute path of the input file, which is visible to the file persona namespace. This input file should contain absolute path of files if basepath is not specified. If basepath is specified, it should be relative path of the files starting from the basepath.

-importfile <source_path>

Path of the input file containing list of files for which retention setting is to be changed or legal hold is to be set or released. The input file should be present on client machine. The max size of import file can be up to 2GiB.

-expdate <expirydate>

Retention expiry date can be any valid future date provided in any of the formats as supported by POSIX touch/date command. See the touch(1) man page documentation for the time/date supported formats. The new expiry date should be less than current date and time plus the maximum retention period specified in the retention policy for that namespace. In case the expiry date value exceeds the maximum limit, the file retention expiry date will be set to the current date and time plus max retention period.

-retenvalidation {enable|disable}

Enables or disable the retention validation scan.

Specifiers

<vfs>

Specifies the name of the VFS.

Subcommands

pol

Configures the retention policy for the FSTORE or VFS in the file provisioning group.

legalhold

Sets or release the legal hold for the file(s) specified.

retention

Changes retention expiration date for the specified retained file(s). Retention expiry date can be valid future date provided in any of the formats as supported by POSIX date command. See the Linux touch(1) man page under -t for supported time/date formats. The new expiry date should be less than current date and time plus the maximum retention period specified in the retention policy for the namespace. In case the expiry date value exceeds the maximum limit, the file retention expiry date will be set to the current date and time plus max retention period.

pol_inherit

It is used to enable or disable the policy inheritance at the FSTORE level. Policy inheritance can be disabled only on FSTORE which does not have files and local retention policy.

Authority

- Super, Edit
- Any role granted `fsarchive_set` right.
- Super for `-importfile <source_path>`.

Usage

- Access to all domains is required to run this command.
- For pol subcommand one of `-mode`, `-defperiod`, `-minperiod`, `-maxperiod`, and `-autocommperiod` or `-retenvalidation` or both must be specified.
- The option `-importfile <source_path>` is not available with SSH access to CLI.

Examples

Configuring the retention policy for the files in the file provisioning group at VFS or File store level:

```
cli% setfsarchive pol -mode enterprise -defperiod 5M -minperiod 3d \
                    -maxperiod 5y -autocommperiod 1h -fstore fstore1 \
                    -fpg fpg1 vfs1
```

Setting inheritance of policy at File store level:

```
cli% setfsarchive pol_inherit -inheritance disable -fstore fstore1 \
                  -fpg fpg1 vfs1
```

Setting legal hold for the specified file:

```
cli% setfsarchive legalhold -set -basepath /fpg1/vfs1
                  -files fstore1/abcd.txt,fstore2/xyz.txt -fstore fstore1
                  -fpg fpg1 vfs1

cli% setfsarchive legalhold -set -basepath /fpg1/vfs1
                  -inputfile /fpg1/vfs1/.admin/abcd.txt -fstore fstore1
                  -fpg fpg1 vfs1

cli% setfsarchive legalhold -set -importfile /home/worm/file.txt \
                  -fstore fstore1 -fpg fpg1 vfs1

cli% setfsarchive legalhold -set -importfile /home/worm/file.txt \
                  -fpg fpg1 vfs1
```

Releasing legal hold for the specified file:

```
cli% setfsarchive legalhold -clear -basepath /fpg1/vfs1
                  -files fstore1/abcd.txt,fstore2/xyz.txt -fstore fstore1
                  -fpg fpg1 vfs1

cli% setfsarchive legalhold -clear -basepath /fpg1/vfs1
                  -inputfile /fpg1/vfs1/.admin/abcd.txt -fstore fstore1
                  -fpg fpg1 vfs1
```

Changing the retention expiration date for the specified retained file(s):

```
cli% setfsarchive retention -basepath /fpg1/vfs1
                  -files fstore1/abcd.txt,fstore2/xyz.txt -expdate 3years
                  -fstore fstore1 -fpg fpg1 vfs1

cli% setfsarchive retention -basepath /fpg1/vfs1 \
                  -inputfile /fpg1/vfs1/.admin/abcd.txt -expdate 3years
                  -fstore fstore1 -fpg fpg1 vfs1

cli% setfsarchive retention -importfile /home/worm/file.txt \
                  -expdate 3years -fstore fstore1 -fpg fpg1 vfs1

cli% setfsarchive retention -basepath /fpg1/vfs1 -files 3.txt
                  -expdate "2016-08-22 22:59:00" -fstore fstore1
                  -fpg fpg1 vfs1
```

setfsaudit

Syntax

The syntax for the `setfsaudit` command can be one of the following:

```
setfsaudit client -type {internal|external} [-fpg <fpgname>] <vfs>

setfsaudit logpol [options] <vfs>

setfsaudit replica -enable {true|false} [-fpg <fpgname>] <vfs>

setfsaudit smb {[-audit {operation1:value[,operation2:value]...}] [-
auditadminshare {operation1:value[,operation2:value]...]}} [-fpg <fpgname>]
<vfs>

setfsaudit user [-delete] [-fpg <fpgname>] <username> <vfs>
```

```
setfsaudit log -delete [-fpg <fpgname>] <logfile> <vfs>
setfsaudit tz [-clear]
```

Description

Modifies File Access Audit settings.

Options

The following options are for **all** subcommands:

-fpg <fpgname>

Specifies the FPG that <vfs> belongs. If this is not specified, the command will determine the FPG based on the specified <vfs>. However, if <vfs> exists under multiple FPGs, `-fpg` must be specified.

The following options are specific to the **client** subcommand:

-type {<internal|external>}

Specifies type of client authorized to consume logs. Internal client is an application built to consume stored logs periodically.

External client is an ISV that will consume File Access Audit logs. Default is internal.

The following options are specific to the **logpol** subcommand:

-size <number>

Maximum size of log file to be generated in MBs. Once the log file reaches a size greater than or equal to the size specified, it is compressed and exported to `.admin/fileauditing`. The supported size range is between 50 MB and 500 MB.

-time <number>

Maximum time in minutes for which log file to be generated. When the log file is created, time counter is started. Once it reaches the number specified in -time argument, log file is compressed and exported to `.admin/fileauditing`. The supported time range is between 1 and 300.

-retention <number>[m|d|h]

Retention is either a number of logs to keep, or duration in months/days/hours to keep log files. When the duration is exceeded, the compressed log file is deleted. Auto deletion of log files is disabled by default.

If only a number is specified, the log file count is "number" (maximum 50). After the specified number is reached, the oldest log file will be removed.

Maximum values for duration in months/days/hours are 1m/30d/720h respectively. 1 month is equivalent to 30 days.

-format {proto|xml|json}

Supports log file format. Supported formats are: proto, xml, and json. Default format is json.

The following option is specific to the **replica** subcommand, which is no longer used:

-enable {true|false}

Enables or disables the replica of message queue. Replica is disabled by default.

The following options are specific to the **smb** subcommand:

-audit {operation1:value[,operation2:value]...}

Specifies the events to be audited on a global level for SMB protocol on a specified VFS. Supported operations are logon, logoff and global. Supported values are success, failure, none (no logging), and all (logging for both success and failure). The operation/values are case insensitive. By default, all operations are set to "none". Operations are:

logon

Logon (authenticate) to the SMB server.

If a logon fails (invalid user/password), the user `sid` and `name` are `null` and `status = 0`.

`logoff`

Log off from the SMB server.

Changing this setting while a session has already been established (logon) has no effect. This setting is only evaluated at the beginning of a session.

`global`

Setting flag to failure allows auditing of TreeConnects to shares that are not defined within the VFS.

Only the failure flag is supported.

The following options are specific to the **user** subcommand:

`-delete`

Deletes authorization to access File Access Audit log messages for the specified user.

The following options are specific to the **log** subcommand:

`-delete`

Deletes the specified compressed audit log file.

Specifiers

The following specifiers are for **all** subcommands:

`<vfs>`

The VFS for which file access audit settings to be modified.

`-auditadminshare {operation1:value[,operation2:value]...}`

Specifies the events to be audited on the administrative share (C$). Supported operations are: TreeConnect, TreeDisconnect, Open, Close, Delete, Rename, Read, Write, and ChangeSecurity. Supported values are success, failure, none (No logging) and all (logging for both success and failure). The operation values are case insensitive. By default, all operations are set to "none".

The operations will generate events for the listed conditions:

- `TreeConnect`: Connected to a share. To create an audit record for a user trying to access a share that is not configured on the VFS, use the command:

  `setfsaudit smb -audit Global:failure`.

- `TreeDisconnect`: Disconnected from a share.

  NOTE: Certain clients that use SMBv1 may not trigger a `TreeDisconnect` when the client logs off.

- `Open`: Open/Create a file or directory.
- `Close`: Close a file or directory.
- `Rename`: Rename a file or directory.
- `Delete`: Delete a file or directory.
- `Read`: Read from a file. Only the first read after the file was opened is reported.
- `Write`: Write to a file. Only the first write after the file was opened is reported.
- `ChangeSecurity`: Change security attributes.

The following specifiers are specific to the **user** subcommand:

<username>

Name of the user to be added/removed for authorization. <username> must either be in NT4 (<domain>\\<user>) or UPN (<user>@<domain>) format.

If the user is configured on Active Directory, use "setfs ad" to join Active Directory domain with <domain> if it has not been done, and then specify the user.

If the user is configured on the LDAP server, use "setfs ldap" to create LDAP configuration with <ldap-netbios> if it has not been done, and then specify the user.

For users configured locally, specify "LOCAL_CLUSTER\<user>".

The following specifiers are specific to the **log** subcommand:

<logfile>

Name of a compressed audit log file.

The following specifiers are specific to the **tz** subcommand:

-clear

Clears the local timezone and sets back the UTC timezone for audit log files.

Subcommands

client

Specifies type of client authorized to consume logs.

logpol

Configures the log policy for defined VFS.

replica

Enables or disables replica of a message queue to another node.

smb

Set global level file access audit configuration for SMB protocol. Set global level file access audit configuration and the administrative share (C$) level file access audit configuration for SMB protocol.

user

Adds or removes authorization for a user to access file access audit log messages. The user name is only applicable when the client type is external.

log

Deletes the specified compressed audit log file.

tz

Sets the local timezone for audit log files. Log file names are created with local timezone and each audit log contains UTC offset to the Local timezone as a first line. If not executed, audit log file names and logs will have date corresponding to UTC.

Authority

- Super, Edit
- Any role granted the fsaudit_set right.

Usage

- Log policy is created implicitly during VFS creation.
- When client is set as external, execution of setfsaudit logpol command will result in error.

- When `replica` is enabled, there will be performance degradation on both active and passive FPG nodes.
- Whenever log format (for example, json, xml, or proto) changes, rotation of log file occurs and new log file is created.
- Options `-size` and `-time` are mutually exclusive.
- To display File access audit settings (such as client type, logpol, replica, and authorized user) use "`showfsaudit pol`".
- When both `-size` and `-time` are specified, and log size reaches its max limit before the specified time limit, log rotation happens based on size. Otherwise, log rotation occurs when the max time limit is reached.

Examples

The following example enables internal client as a consumer for the VFS vfs1:

```
cli% setfsaudit client -type internal vfs1
```

The following example enables or overwrites existing file access audit log policy settings for internal client on specified VFS for a max log file size of 50MB, retention period of seven days and log format as xml:

```
cli% setfsaudit logpol -size 50 -retention 7d -format xml vfs1
```

The following example enables replica for vfs1 related audit logs:

```
cli% setfsaudit replica -enable true -fpg fpg1 vfs1
```

The following example disables replica for vfs1 related audit logs:

```
cli% setfsaudit replica -enable false -fpg fpg1 vfs1
```

The following example enables auditing event on vfs1 for successful logon operation through SMB protocol:

```
cli% setfsaudit smb -audit logon:success vfs1
```

The following example authorizes local user user1 to access vfs1 related audit logs:

```
cli% setfsaudit user user1 vfs1
```

The following example authorizes AD user user2 to access vfs1 related audit logs:

```
cli% setfsaudit user "2008ad.lab\\user2" vfs1
```

The following example authorizes LOCAL_CLUSTER user Administrator to access vfs1 related audit logs:

```
cli% setfsaudit user "LOCAL_CLUSTER\Administrator" vfs1
```

The following example unauthorizes local user user1 from accessing vfs1 related audit logs:

```
cli% setfsaudit user -delete user1 vfs1
```

The following example sets the local timezone for audit log files:

```
cli% setfsaudit tz
```

The following example clears the local timezone for audit log files:

```
cli% setfsaudit tz -clear
```

The following example enables auditing of successful TreeConnect event for the administrative share (C$) through the SMB protocol:

```
cli% setfsaudit smb -auditadminshare treeconnect:success vfs1
```

setfsav

Syntax

The syntax for the `setfsav` command can be one of the following:

```
setfsav vse [[+|-]<vselist>]
```

```
setfsav pol [-scan {enable|disable|inherit}] [-vendor <vendor_name>] [-fileop
{open|openclose|inherit}] [-unavail {allow|deny|inherit}] [-excludesize
{<size>|inherit}] [-excludeext {<ext>[,<ext>...]|inherit}] [-inheritall] [-
fpg <fpgname>] [-fstore <fstore>] <vfs>
```

```
setfsav quar {exportlist|move|reset|delete|clearcount} [-fpg <fpgname>] [-
fstore <fstore>] [-quar_file <filepath>] <vfs>
```

Description

The `setfsav` command configures antivirus properties for File Persona.

Options

`[+|-]<vselist>`

Specifies the IPv4 address and port number of an external VSE. "port" is allowed only while adding `<vselist>` to the existing VSE list.

If `<vselist>` has a prefix:

- `+` add `<vselist>` to the existing VSE list. The "vseip:port" tuple in `<vselist>` must not be present in the existing VSE list. It must be specified in the format of: "`<vseip:port>[,<vseip:port>]...`".
- `-` remove `<vselist>` from the existing VSE list. The "vseip" in `<vselist>` must be present in the existing VSE list. It must be specified in the format of: "`<vseip>[,<vseip>]...`".

  If `<vselist>` has no prefix, `<vselist>` will be used as the new list.

  Adding "vseip:port" to an empty VSE list will automatically start the antivirus service.

  When the last "vseip" is removed from the list, the antivirus service will be stopped automatically.

`-fpg <fpgname>`

The name of the file provisioning group (FPG) in which the VFS was created.

`-fstore <fstore>`

Specifies the file store name.

`-inheritall`

Inherits all the settings from the VFS for the specified file store, overriding any previous settings in the file store.

Only valid when `-fstore` option is specified.

`-scan {enable|disable|inherit}`

Enables or disables the antivirus scan, or inherits the scan setting from VFS.

If not specified, the default value is disable for VFS, inherit for file store.

`-vendor <vendor_name>`

Specifies the antivirus vendor name.

Valid values are MCAFEE, SYMANTEC, TRENDMICRO, SOPHOS, or KASPERSKY.

Only valid in VFS context.

-quar_file <filepath>

Specifies the path of file on VFS containing list of quarantined files. The file contains quarantined files, one per line. This option can only be used with move, reset and delete operations.

-fileop {open|openclose|inherit}

Specifies the policy that determines which file operations trigger antivirus scans. The policies are:

- `open`—Scan on file open.
- `openclose`—Scan on file open and file close.
- `inherit`—Inherits the fileop setting from VFS.

If it is not specified, the default is "open" when applied to a VFS, "inherit" when applied to a file store.

-unavail {allow|deny|inherit}

Specifies the scan policy to determine how targeted file operations are handled when an external VSE is not available.

The policies are:

- `allow`—All operations triggering scans are allowed to run to completion.
- `deny`—All operations triggering scans are blocked and returned with an error.
- `inherit`—Inherits the unavail setting from VFS.

If it is not specified, the default is "allow" when applied to a VFS, "inherit" when applied to a file store.

-excludesize {<size>|inherit}

Excludes all files larger than the specified size (MB). The value of size is an integer from 0 to 2147483647.

If this option is not specified or size is 0, all files will be included in antivirus scan.

If "inherit" is specified, the "excludesize" setting will be inherited from VFS.

-excludeext {<ext>[,<ext>,...]|inherit}

Excludes all files having the specified extension.

If this option is not specified or ext is '', all files will be included in antivirus scan.

If "inherit" is specified, the "excludeext" setting will be inherited from VFS.

Specifiers

<vfs>

The VFS name, using up to 31 characters.

Subcommands

vse

Specifies the virus scan engine (VSE).

pol

Specifies the antivirus policies.

quar

Specifies the management of quarantined files. It has following subcommands

exportlist

Exports the list of quarantined files to default (.admin/AV/Quarantine in specified VFS) location.

move

Moves quarantined files to default location (.admin/AV/Quarantine folder in specified VFS) with timestamp.

reset

Resets quarantined files under specified virtual file server (VFS)/file store.

delete

Deletes quarantined files under specified VFS/file store.

clearcount

Deletes AV statistics for a specified VFS.

Authority

- Super, Edit
- Any role granted the `fsav_set` right

Usage

- Access to all domains is required to run this command.
- For all options allowing the specification of "inherit": This is only valid when applied to a file store, i.e., in conjunction with the "`-fstore`" option.
- The following steps need to be followed to manage the quarantined files:
  1. Administrator archives the list of quarantined files in `<file>` by using "`setfsav quar exportlist`" option.
  2. Administrator evaluates the infected files in `<file>` and decides what operation to be done (`reset`, `move`, or `delete`) for each of the quarantined files.
- When only "`setfsav vse`" is specified, it will clear the VSE list.

Examples

The following example shows how to configure VSE IP and port. Antivirus service starts automatically:

```
cli% setfsav vse 10.2.2.156:1344
```

The following example shows how to add a new VSE IP and port:

```
cli% setfsav vse +10.2.2.157:1344
```

The following example shows how to remove a specific VSE IP:

```
cli% setfsav vse -10.2.2.156
```

The following example shows how to configure various antivirus properties for VFS:

```
cli% setfsav pol -scan enable -vendor SYMANTEC -fileop open  -unavail allow -excludesize 10 -excludeext htm,jpg  testvfs
```

The following example shows how to override VFS antivirus properties "fileop" and "excludesize" into file store:

```
cli% setfsav pol -fstore engineering -fileop openclose  -excludesize 100 testvfs
```

The following example shows how to inherit VFS antivirus properties "fileop" and "excludesize" into file store:

```
cli% setfsav pol -fstore engineering -fileop inherit  -excludesize inherit testvfs
```

The following example shows how to inherit all of the VFS antivirus properties into file store:

```
cli% setfsav pol -inheritall -fstore engineering testvfs
```

The following example shows how to set no exclusion for size in file store:

```
cli% setfsav pol -fstore engineering -excludesize 0 testvfs
```

The following example shows how to export list of quarantined filed under unityvfs:

```
cli% setfsav quar exportlist unityvfs
```

The following example shows how to move quarantined files under VFS unityvfs:

```
cli% setfsav quar move unityvfs
```

The following example shows how to reset quarantined files under file store engineering under VFS unityvfs:

```
cli% setfsav quar reset -fstore engineering unityvfs
```

The following example shows how to delete quarantined files under file store engineering under VFS unityvfs:

```
cli% setfsav quar delete -fstore engineering unityvfs
```

The following example shows how to delete AV statistics for a VFS vfs1:

```
cli% setfsav quar clearcount vfs1
```

The following example shows how to move the quarantined files listed in file_move:

```
cli% setfsav quar move -quar_file .admin/AV/Quarantine/file_move \
                    unityvfs
```

The following example shows how to reset the quarantined files listed in file_reset:

```
cli% setfsav quar reset -quar_file .admin/AV/Quarantine/file_reset \
                    unityvfs
```

The following example shows how to delete the quarantined files listed in file_delete:

```
cli% setfsav quar delete -quar_file .admin/AV/Quarantine/file_delete \
                    unityvfs
```

setfsgroup

Syntax

```
setfsgroup [options] <groupname>
```

Description

The `setfsgroup` command modifies a local group account.

Options

`-memberlist <list>`

Specifies user members of the group. It is a set of comma separated strings.

If <list> has a prefix (for example, +user1):

- + — add <list> to the existing user list. Users in <list> must not be in the existing list.
- – — remove <list> from the existing list. Users in <list> must be already in the existing list.

If specified, the prefix will be applied to the entire list.

If <list> has no prefix, <list> will be used as the new user list. If "" is specified user list will be emptied.

-f

Specifies that the operation is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

Specifiers

<groupname>

Specifies the local group name using up to 31 characters. Valid characters are alphanumeric characters, periods, dashes (except first character), and underscores.

Authority

- Super
- Any role granted the `fsgroup_set` right.

Usage

Access to all domains is required to run this command.

Examples

The following example shows how to add a user jdoe to local group named accounting:

```
cli% setfsgroup -memberlist +jdoe accounting
```

setfshare

Syntax

```
setfshare {smb|nfs|obj|ftp} [options <arg>] <vfs> <sharename>
```

Description

Modifies File Share properties for supported protocols.

Options

The following options are for **all** subcommands:

-acl [+|-]<permlist>

Specifies the ACL permissions that are allowed on a share directory.

This option cannot be used when `-mode` is used. Setting ACL permissions will break cross protocol functionality and may cause loss of modebits information. A warning prompt is displayed to the users asking if they want to proceed.

The <permlist> contains the list of ACEs. Use commas to separate ACEs.

Each ACE contains four values named type, flag, principal and permissions.

These four values must be separated by ":".

Example: `A:fd:OWNER@:rwax,A:fdg:GROUP@:rwax`

If <permlist> has a prefix (for example: `+A:fd:OWNER@:rwa,A:g:GROUP@:rwxa,`.):

- + — add <permlist> to the existing permlist. The ACEs in <permlist> may be present in the existing list. The new <permlist> is inserted at the head of the existing ACL, in the same order specified.
- – — remove <permlist> from the existing list. The ACEs in <permlist> must be already in the existing list. The first matching entry will be removed.

If specified, the prefix will be applied to the entire <permlist>.

If <permlist> has no prefix, it will be used as the new permlist.

Duplicate ACEs are allowed. The order specified for the ACEs will be maintained.

The values for <permlist> fields type:flag:principal:permissions will be as follows.

The "type" field can take only one of the following values:

- A—allow
- D—deny
- U—audit
- L—alarm

The "flags" is optional and can take one or more of the following values:

- f—file-inherit
- d—directory-inherit
- p—no-propagate-inherit
- i—inherit-only
- S—successful-access
- F—failed-access
- g—group (denotes that <principal> is a group)

The "principal" field can be any named user or group or one of the following values:

- OWNER@
- GROUP@
- EVERYONE@

The "permissions" field can take one or more of the following values:

- r—read-data | list-directory
- w—write-data | create-file
- a—append-data | create-subdirectory
- x—execute

- d—delete
- D—delete-child (directories only)
- t—read-attrs
- T—write-attrs
- n—read-named-attrs
- N—write-named-attrs
- c—read-ACL
- C—write-ACL
- o—write-owner
- y—synchronize

-mode <modebits>

Specifies the modebits permissions that are allowed on a share directory.

This option cannot be used when `-acl` is used. Setting mode bits will break cross protocol functionality and may cause loss of ACL information.

A warning prompt is displayed to the users asking if they want to proceed.

-owner <name>

Specifies the name of the owner to which the share directory belongs.

If the owner is configured on Active Directory, use "`setfs ad`" to join Active Directory domain with <domain> if it has not been done, and use "<domain>\\<uname>" or "<ad-netbios>\\<uname>" to specify the owner (for example, `-owner example.com\aduser`).

The "<ad-netbios>" is Active Directory NetBIOS name, which can be found by running "`showfs -ad`".

If the owner is configured on the LDAP server, use "`setfs ldap`" to create LDAP configuration with <ldap-netbios> if it has not been done, and use "<ldap-netbios>\\<owner>" to specify the owner (for example, `-owner ldaphost\\ldapuser`). The "<ldap-netbios>" is the LDAP server NetBIOS name, which can be found by running "`showfs -ldap`".

-group <name>

Specifies the name of the group to which the share directory belongs.

If the group is configured on Active Directory, use "`setfs ad`" to join Active Directory domain with <domain> if it has not been done, and use "<domain>\\<gname>" or "<ad-netbios>\\<uname>" to specify the user (for example, -group example.com\adgroup).

The <ad-netbios> is Active Directory NetBIOS name, which can be found by running "`showfs -ad`".

If the group is configured on the LDAP server, use "`setfs ldap`" to create LDAP configuration with <ldap-netbios> if it has not been done, and use "<ldap-netbios>\\<gname>" to specify the user (for example, `-group ldaphost\\ldapgroup`).

-f

This option specifies that the command is forced. When setting ACL permissions or modebits of a share directory, if this option is not used, the command requires confirmation before proceeding with its operation.

-fpg <fpgname>

Specifies the file provisioning group (FPG) that <vfs> belongs. If this is not specified, the command will find out the FPG based on the specified <vfs>.

However, if <vfs> exists under multiple FPGs, `-fpg` must be specified.

-fstore <fstore>

Specifies the file store that the share to be modified belongs.

If this is not specified, the <sharename> will be used as the file store name to identify the share.

-comment <comment>

Specifies any comments or additional information for the share. The comment can be up to 255 characters long. Unprintable characters are not allowed.

The following options are specific to the **smb** subcommand:

-abe {true|false}

Access-Based Enumeration. Specifies if users can see only the files and directories to which they have been allowed access on the shares.

-allowip [+|-]<iplist>

Specifies client IP addresses that are allowed access to the share. Use commas to separate the IP addresses. If <iplist> has a prefix (for example: +1.1.1.0,2.2.2.0):

- + — add <iplist> to the existing allowed list. The IP addresses in <iplist> must not be in the existing allowed list.
- – — remove <iplist> from the existing allowed list. The IP addresses in <iplist> must be already in the existing allowed list.

If specified, the prefix will be applied to the entire <iplist>. If <iplist> has no prefix, <iplist> will be used as the new allowed list.

-denyip [+|-]<iplist>

Specifies client IP addresses that are denied access to the share. Use commas to separate the IP addresses. If <iplist> has a prefix (for example: +1.1.1.0,2.2.2.0):

- + — add <iplist> to the existing denied list. The IP addresses in <iplist> must not be in the existing denied list.
- – — remove <iplist> from the existing denied list. The IP addresses in <iplist> must already be in the existing denied list.

If specified, the prefix will be applied to the entire <iplist>. If <iplist> has no prefix, <iplist> will be used as the new denied list.

-allowperm [+|-|=]<permlist>

Specifies the permissions that users/groups are allowed to access the share. <permlist> must be specified in the format of: "<user1>:<perm1>,<user2>:<perm2>,...". The <user> can be a user or group name specified using the same format as described in `createfshare`. <perm> must be "fullcontrol", "read", or "change".

If <permlist> has a prefix (for example: +Everyone:read):

- + — add <permlist> to the existing allowed list. Users/groups in <permlist> must not be in the existing allowed list.
- − — remove <permlist> from the existing allowed list. Users/groups in <permlist> must be already in the existing allowed list.
- = — modify the existing allowed list with <permlist>. Users/groups in <permlist> must be already in the existing allowed list.

If specified, the prefix will be applied to the entire <permlist>. If <permlist> has no prefix, <permlist> will be used as the new allowed list.

-denyperm [+|-|=]<permlist>

Specifies the permissions that users/groups are denied to access the share. <permlist> must be specified in the format of: "<user1>:<perm1>,<user2>:<perm2>,...". The <user> can be a user or group name specified using the same format as described in createfshare. <perm> must be "fullcontrol", "read", or "change".

If <permlist> has a prefix (for example, +Everyone:read):

- + — add <permlist> to the existing denied list. Users/groups in <permlist> must not be in the existing denied list.
- − — remove <permlist> from the existing denied list. Users/groups in <permlist> must be already in the existing denied list.
- = — modify the existing denied list with <permlist>. Users/groups set in <permlist> must be already in the existing denied list.

If specified, the prefix will be applied to the entire <permlist>. If <permlist> has no prefix, <permlist> will be used as the new denied list.

-audit {operation1:value1[,operation2:value2]...}

Specifies the operations to be audited at share level for SMB protocol on specified VFS. The specified operations will be modified and all the operations that are not specified will be set to default values.

Supported operations are mentioned below. The event can be audited based on operations: success, failure, none or all outcome. The operation/values are case insensitive.

The operations listed below will generate events when:

- TreeConnect—Connected to a share. To create an audit record for a user trying to access a share that is not configured on the VFS, use the command: "`setfsaudit smb -audit Global:failure`".
- TreeDisconnect—Disconnected from a share.

  NOTE: Certain clients that use SMBv1 may not trigger a TreeDisconnect when the client logs off.
- Open—Open/Create a file or directory.
- Close—Close a file or directory.
- Rename—Rename a file or directory.
- Delete—Delete a file or directory.
- Read—Read from a file. Only the first read after the file was opened is reported.
- Write—Write to a file. Only the first write after the file was opened is reported.
- ChangeSecurity—Change security attributes.

For more details, see the *File Persona User Guide*.

-cache {off|manual|optimized|auto}

Specifies client-side caching for offline files. Valid values are:

- "off"—The client must not cache any files from this share. The share is configured to disallow caching.
- "manual"—The client must allow only manual caching for the files open from this share.
- "optimized": The client may cache every file that it opens from this share. Also, the client may satisfy the file requests from its local cache. The share is configured to allow automatic caching of programs and documents.
- "auto"—The client may cache every file that it opens from this share. The share is configured to allow automatic caching of documents.

-ca {true|false}

Specifies if SMB3 continuous availability features must be enabled for this share.

The following options are specific to the **nfs** subcommand:

-options <options>

Specifies the new options to use for the share. This completely overwrites the options you set previously. Standard NFS export options except "no_subtree_check" are supported. Do not enter option "fsid", which is provided. If not specified, the following options will be automatically set: sync, auth_nlm, wdelay, sec=sys, no_all_squash, crossmnt, secure, subtree_check, hide, root_squash, ro.

See linux exports(5) manpage for detailed information on valid options.

-audit {operation1:value1[,operation2:value2]...}

Following operations can be audited on a share level for NFS protocol on specified VFS.

Acceptable value is either `on` or `off`. Defaults to `off`. Operations are case insensitive.

The following operations generate events when:

- audit_open—Open a file. This is NFSv4 only.
- audit_close—Close a file. This is NFSv4 only.
- audit_meta—Rename a file or directory, creating a hardlink for a file, creating a file or directory and removing a file or a directory.
- audit_attr—Changing file/directory attributes.
- audit_read—Read of a file/directory. Subsequent file reads are suppressed for a specific time (currently 10m). Note, reads from a directory are never suppressed.
- audit_write—Write to a file. Subsequent writes are suppressed for a specific time (currently 10m).
- audit_acl— Allow audit for ACL (AccessControlList) changes on a file or directory.

For more details, see, *File Persona User Guide*.

-clientip [+|-]<iplist>

Specifies the clients that can access the share. The NFS client can be specified by the name (for example, sys1.example.com), the name with a wildcard (for example, *.example.com), or by its IP address. Use comma to separate the IP addresses.

If <iplist> has a prefix (for example, +1.1.1.0,2.2.2.0):

- + — add <iplist> to the existing list. IP addresses in <iplist> must not be in the existing list.
- – — remove <iplist> from the existing list. IP addresses in <iplist> must be already in the existing list.

If specified, the prefix will be applied to the entire <iplist>. If <iplist> has no prefix, <iplist> will be used as the new list.

The following options are specific to the **obj** subcommand:

-ssl {true|false}

Specifies to enable or disable SSL.

The following options are specific to the **ftp** subcommand:

-options [+]<option1=value1>[,<option2=value2>]...

Specifies the configuration options to be modified for the FTP share. With no prefix specified, only the specified options are modified and all other options are set to default values. With the + prefix specified, only the specified values are modified. Setting options to "" sets all values to the default.

See `-options` under `createfshare ftp` for supported values.

-ssl {true|false}

Specifies to enable or disable SSL.

-shareip [+|-]<iplist>

This option is used to add, remove, or overwrite IP addresses of the FTP share. To add/overwrite, the IP addresses mentioned in the <iplist> must be assigned to the specified VFS. Use commas to separate the IP addresses.

If <iplist> has a prefix (for example, +1.1.1.0,2.2.2.0):

- + — add <iplist> to the existing list. IP addresses in <iplist> must not be in the existing list.
- – — remove <iplist> from the existing list. IP addresses in <iplist> must be already in the existing list.

If no prefix is specified, it overwrites the existing IP list with the specified <iplist>. If specified, the prefix will be applied to the entire <iplist>.

Specifiers

<vfs>

Specifies the virtual file system (VFS) that the share to be modified belongs.

<sharename>

Specifies the name of the share to be modified.

Subcommands

smb

Sets file share options for SMB.

nfs

Sets file share options for NFS.

obj

Sets file share options for Object.

NOTE: WARNING: The Object Access API capability has been deprecated and will be removed in a future release.

ftp

Sets file share options for FTP.

Authority

- Super, Edit
- Any role granted the `fshare_set` right.

Usage

- Access to all domains is required to run this command.
- For setting SMB permissions, the same user cannot be specified with the same permission in both "allowperm" and "denyperm".

Examples

The following example allows the host with IP address "100.1.1.1" to access SMB file share "myshare":

```
cli% setfshare smb -allowip 100.1.1.1 -fstore myfstore myvfs myshare
```

The following example modifies options of NFS file share with export path /<myfpg>/myvfs/myfstore to become "rw", where <myfpg> is the file provisioning group that myvfs belongs:

```
cli% setfshare nfs -options rw -fstore myfstore myvfs myshare
```

The following example changes the Object share "myshare" to enable SSL:

```
cli% setfshare obj -ssl true -fstore myfstore myvfs myshare
```

The following example sets ACL permissions and owner name on NFS file share "myshare":

```
cli% setfshare nfs -acl +A:fd:OWNER@:rwax -owner user1 -fstore \
    myfstore myvfs myshare
```

The following example sets NFS share myshare with auditing option for open, close operation through NFS protocol:

```
cli% setfshare nfs -audit audit_open:on,audit_close:on -fstore \ myfstore myvfs myshare
```

The following example sets modebits and group name on SMB file share "myshare":

```
cli% setfshare smb -mode 755 -group group1 -fstore myfstore myvfs \
    myshare
```

The following example sets owner and group on OBJ file share "myshare":

```
cli% setfshare obj -owner user1 -group group1 -fstore myfstore myvfs myshare
```

The following example adds IP 1.1.1.1 to FTP file share "myshare":

```
cli% setfshare ftp -shareip +1.1.1.1 -fstore myfstore myvfs myshare
```

The following example disables ssl on FTP file share "myshare":

```
cli% setfshare ftp -ssl false -fstore myfstore myvfs myshare
```

The following example modifies the option "anon_other_write_enable" to "yes" and resets all other options to default, on FTP file share "myshare":

```
cli% setfshare ftp -options anon_other_write_enable=yes -fstore \
    myfstore myvfs myshare
```

The following example modifies the option "anon_other_write_enable" to "yes" and leaves all other options as they are on FTP file share "myshare":

```
cli% setfshare ftp -options +anon_other_write_enable=yes -fstore \
    myfstore myvfs myshare
```

The following example modifies SMB share share1 with auditing enabled for successful TreeConnect operation through SMB protocol:

```
cli% setfshare smb -audit TreeConnect:success vfs1 share1
```

setfsip

Syntax

```
setfsip [options] <vfs> <id|ip>
```

Description

The `setfsip` command modifies the network config of a virtual file system (VFS).

Options

`-vlantag <tag>`

Specifies the VLAN Tag to be used.

`-ip <ipaddr>`

Specifies the new IP address.

`-subnet <subnet>`

Specifies the new subnet mask.

`-fpg <fpgname>`

Specifies the file provisioning group (FPG) in which the VFS was created.

`-subnetoverlap`

Allow overlapping of subnets.

NOTE: This is not a typical configuration. If the configuration is not proper, the configuration can cause data unavailability. If the option is not specified, overlapping subnets will not be allowed.

`-f`

Specifies that the operation is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

Specifiers

`<id|ip>`

Specifies the ID/IP for the network config.

`<vfs>`

Specifies the VFS which is to have its network config modified.

Authority

- Super, Edit
- Any role granted the `fsip_set` right.

Usage

Access to all domains is required to run this command.

Examples

The following example modifies the subnet mask to 255.255.255.254 of VFS vfs_1:

```
cli% setfsip -subnet 255.255.255.254 vfs_1 12345678
cli% setfsip -subnet 255.255.255.254 vfs_1 10.10.10.1
```

The following example allows overlapping subnets when modifying subnet mask to 255.255.255.251 for VFS, vfs_1:

```
cli% setfsip -subnetoverlap -subnet 255.255.255.251 vfs_1 12345678
```

setfsndmp

Syntax

The syntax for the `setfsndmp` command can be one of the following:

`setfsndmp conf`

- `[-dma [+|-]<dma_ip_list>] [-username <username>]`
- `[-passwd <password>] [-enable_sessions {true|false}]`
- `[-loglevel <level>] [-tcpwinsize <bytes>]`
- `[-maxsessions <numsessions>]`

`setfsndmp vtl [-force] [+|-]<vtliplist>`

`setfsndmp sessions -cancel <session_list>`

Description

The `setfsndmp` command configures NDMP properties for file persona like DMA, VTL and session cancellation. When no option is specified after `setfsndmp` command, it is an error scenario.

Options

The following options are specific to the conf subcommand:

`[-dma [+|-]<dma_ip_list>] [-username <username>]`

`[-passwd <password>] [-enable_sessions {true|false}]`

`[-loglevel <level>] [-tcpwinsize <bytes>]`

`[-maxsessions <numsessions>]`

Specifies DMA . At least one of the parameters should be given after subcommand "conf" to modify the DMA configuration.

`-dma`

Valid IP address of the DMA machine.

`+<dma_ip_list>`

Adds dma IP to the existing configuration.

`-<dma_ip_list>`

Removes dma IP from the existing configuration.

`-username`

Username to be verified while connecting from the DMA.

-passwd

Password to be verified while connecting from the DMA.

-enable_sessionstrue | false

(case insensitive)

- false—No further NDMP sessions will be allowed
- true—Allow further NDMP sessions.

-loglevel

Sets the NDMP trace log level. Range from 0 to 10

- 10—represents the most logging
- 0—represents only critical logs and error messages

-tcpwinsize

Sets the TCP window size for NDMP data transfer.

Range from 65534 - 163840

-maxsessions

Sets the maximum number of concurrent sessions per NDMP server.

Range from 1 - 128

vtl [-force] [+|-]<vtliplist>

Discovers and add or remove target VTL device with the specified IP to or from the Unity Cluster. VTL devices are added to the cluster list even if target VTL device is not discovered.

+<vtliplist>

discovers and add VTL IP to the unity cluster list.

-<vtliplist>

removes the VTL IP from the list.

-force -<vtliplist>

forcefully remove the VTL IP from the list.

sessions -cancel <session_list>

Cancels NDMP sessions specified with session IDs.

Specifiers

None.

Subcommands

conf

Configures NDMP parameters on the cluster.

vtl

Discovers and add a target VTL device with the specified IP to the Unity Cluster. It can also be used to delete an already added target VTL device.

sessions -cancel

Cancels NDMP sessions specified with session IDs.

Authority

- Super, Edit
- Any role granted the `fsndmp_set` right

Usage

- Access to all domains is required to run this command.
- If no option is specified after `setfsndmp`, help will be displayed as it is an error case.
- `setfsndmp -conf <parameters>` configures NDMP parameters on the cluster.
- VTL IPs can be modified by using vtl subcommand and NDMP sessions can be cancelled by using `sessions -cancel` subcommand.

Examples

The following example shows how to configure NDMP parameters on the cluster:

```
cli% setfsndmp conf -dma +1.2.3.4,4.5.6.7 -username ndmp
     -passwd ndmp -enable_sessions true -loglevel 5
     -tcpwinsize 163840 -maxsessions 128
```

The following example shows how to discover and add/remove target vtl device with the specified IP to the Unity Cluster:

```
cli% setfsndmp vtl +1.2.3.4,4.5.6.7
 cli% setfsndmp vtl -4.5.6.7
 cli% setfsndmp vtl -force -1.2.3.4
```

The following example cancels NDMP sessions specified with session IDs:

```
cli% setfsndmp sessions -cancel 34545,6766
```

setfsquota

Syntax

```
setfsquota [options] <vfsname>
```

Description

The `setfsquota` command modifies the quotas for a given virtual file server (VFS).

Options

`-fpg <fpgname>`

Specifies the name of the file provisioning group (FPG) hosting the VFS.

`-username <uname>`

The username of the quotas to be modified.

If the user is configured on Active Directory, use "setfs ad" to join Active Directory domain with <domain> if it has not been done, and use "<domain>\\<uname>" or "<ad-netbios>\\<uname>" to specify the user (for example, -username example.com\aduser).

The "<ad-netbios>" is Active Directory NetBIOS name, which can be found by running "showfs -ad".

If the user is configured on the LDAP server, use "setfs ldap" to create LDAP configuration with <ldap-netbios> if it has not been done, and use "<ldap-netbios>\\<username>" to specify the user (for example, -username ldaphost\\ldapuser). The "<ldap-netbios>" is the LDAP server NetBIOS name, which can be found by running "showfs -ldap".

-groupname <gname>

The groupname of the quotas to be modified.

If the group is configured on Active Directory, use "setfs ad" to join Active Directory domain with <domain> if it has not been done, and use "<domain>\\<gname>" or "<ad-netbios>\\<uname>" to specify the user (for example, -groupname example.com\adgroup).

The <ad-netbios> is Active Directory NetBIOS name, which can be found by running "showfs -ad".

If the group is configured on the LDAP server, use "setfs ldap" to create LDAP configuration with <ldap-netbios> if it has not been done, and use "<ldap-netbios>\\<gname>" to specify the user (for example, -groupname ldaphost\ \ldapgroup).

-uid <uid>

The user id of the quotas to be deleted. This is supported only with -clear. Options -uid, -gid, -username and -groupname are mutually exclusive.

-gid <gid>

The group id of the quotas to be deleted. This is supported only with -clear. Options -uid, -gid, -username and-groupname are mutually exclusive.

-fstore <fstore>

The file store to which you wish to apply quotas.

-scapacity <soft capacity limit>

An integer value in MiB for the soft capacity storage quota. The maximum value is the maximum size supported by FPG in MiB. 0 - No quota

-hcapacity <hard capacity limit>

An integer value in MiB for the hard capacity storage quota. The maximum value is the maximum size supported by FPG in MiB. 0 - No quota

-sfile <soft file limit>

An integer limit of the number of files for the soft file quota. The maximum value is 250000000. 0 - No quota

-hfile <hard file limit>

An integer limit of the number of files for the hard file quota. The maximum value is 250000000. 0 - No quota

-clear

Clears the quotas of the specified object.

-archive

Stores the quota information associated with the VFS in a file.

-restore <file>

Applies the quota information stored in the file to the VFS.

Specifiers

<vfsname>

Specifies the name of the VFS associated with the quotas.

Authority

- Super, Edit.
- Any role granted `fsquota_set` right.

Usage

- Access to all domains is required to run this command.
- Only one set of quotas may be specified per call.

Examples

The following example sets the quotas for group gExample to 1024MB soft and 2048MB hard:

```
cli% setfsquota -groupname gExample -scapacity 1024 -hcapacity 2048 \
-sfile 1024 -hfile 2048 examplevfs
```

The following example deletes the quotas for user with uid 11001:

```
cli% setfsquota -uid 11001 -clear examplevfs
```

The following example restores a set of quotas for the examplevfs VFS:

```
cli% setfsquota -restore /examplefpg/examplevfs/.admin/Quotas/quotas_examplefpg_examplevfs.exports examplevfs
```

setfsroute

Syntax

```
setfsroute modifygw [-f] {<targetaddr>,{<subnetmask>|<prefixlen>},<vlantag> |
<routeidentifier>} <gateway>
```

Description

The `setfsroute` command modifies a route for a target identified with either route id or a unique combination of target address, vlantag and subnet mask. Currently only gateway can be modified. A warning is displayed to user before modifying the route.

Options

`-f`

Suppresses confirmation from user before modifying the route.

Specifiers

`<targetaddr>`

The target IPv4/IPv6 address of the route to be modified.

`<subnetmask>|<prefixlen>`

The subnet mask or prefix length for the target IP Address.

`<vlantag>`

The VLAN tag associated with the route to be modified.

`<routeidentifier>`

Instead of providing a combination of targetaddr,subnetmask|prefixlen and vlantag, a route identifier can be provided. Obtain the route identifier from the "`showfsroute -d`" command.

`<gateway>`

New gateway to be assigned to the target IP address.

Subcommands

`modifygw`

Modifies the gateway of the route specified.

Authority

- Edit
- Any role granted `fsroute_set` right.

Usage

Access to all domains is required to run this command.

Examples

The following example modifies a route with following details:

```
cli% setfsroute modifygw 10.16.23.101,16,8 10.16.23.1
```

The following example modifies a route without confirmation:

```
cli% setfsroute modifygw -f 10.16.23.101,255.255.192.0,12 10.16.23.1
```

The following example modifies a route with specified route ID:

```
cli% setfsroute modifygw staticRoute4e426403-98a7-4025-9b11-ca59afd9971c 10.16.23.1
```

setfstore

Syntax

```
setfstore [options] <vfs> <fstore>
```

Description

Allows modification of the specified file store.

Options

-f

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

-comment <comment>

Specifies any addition textual information.

-fpg <fpgname>

The name of the parent file provisioning group (FPG).

-secop_errsuppress {true|false}

Enables or disables the security operations error suppression for file stores in "ntfs" security mode. In "ntfs" security mode, permission changing operations will report error messages for non-root users when -secop_errsuppress is "false". This may not be the ideal behavior for some applications trying to implicitly perform a permission changing operation such as "chmod" over an NFSv3 client. This tunable can be set to "true" to allow those applications to work with NFSv3 clients. This option cannot be used for file stores in "legacy" security mode.

-secmode ntfs

Specifies the security mode of the file store. Security mode of a file store can be changed to NTFS from LEGACY. Valid values are:

ntfs

Windows and non-Windows clients can create/read/write/delete/ rename file objects (files and directories). Permissions on new file objects will be based on NTFS inheritance rules. Windows clients can view and set permissions (Windows ACLs). Non-windows clients can view permissions, but cannot modify them (change of ownership, group, mode bits or POSIX ACLs are not permitted).

Specifiers

<vfs>

The name of the containing virtual file server (VFS).

<fstore>

The name of the file store to be modified.

Authority

- Super, Edit
- Any role granted `fstore_set` right.

Usage

Access to all domains is required to run this command.

Examples

The following example sets the comment string for the specified file store:

```
cli$ setfstore -comment "Example comment string." -fpg fpg1 vfs1 examplefstore
```

The following example modifies the security operation error suppression tunable to "true":

```
cli$ setfstore -secop_errsuppress "true" vfs1 examplefstore
```

setfsuser

Syntax

```
setfsuser [options] <username>
```

Description

The `setfsuser` command modifies a local user account associated with File Persona.

Options

-primarygroup <groupname>

Specifies the user's primary group.

-passwd <password>

Specifies the user's password. The -passwd option and -passprompt option are mutually exclusive options.

-passprompt

Indicates a password prompt request to enter password. The -passwd option and -passprompt option are mutually exclusive options.

-enable {true|false}

Specifies if the user is enabled or not.

-grplist <list>

Specifies a list of additional groups which the user is to be a member. It is a set of comma separated strings.

If <list> has a prefix (for example, +group1): d

+

add <list> to the existing group list. Groups in <list> must not be in the existing list.

-

remove <list> from the existing list. Groups in <list> must be already in the existing list.

If specified, the prefix will be applied to the entire list.

If <list> has no prefix, <list> will be used as the new group list. If "" is specified the group list will be emptied.

-f

Specifies that the operation is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

Specifiers

<username>

Specifies the local user name using up to 31 characters. Valid characters are alphanumeric characters, periods, dashes (except first character), and underscores.

Authority

- Super
- Any role granted the `fsuser_set` right

Usage

- Access to all domains is required to run this command.
- User's primary group cannot be Built-in groups:
- "Administrators", "Users", "Guests" and "Backup Operators".

Examples

The following example modifies a local user named jdoe:

```
cli% setfsuser -primarygroup finance jdoe
```

sethost

Syntax

```
sethost [options <arg>] <hostname>
sethost initchap [-f] [options <arg>] <secret> {<hostname> | <pattern> ...}
sethost targetchap [-f] [options <arg>] <secret> {<hostname> | <pattern> ...}
sethost removechap [-target] [-f] {<hostname> | <pattern> ...}
sethost clearagent <WWN | iscsi_name>
sethost rst <hostname>
sethost rst -wwn <WWN> <node:slot:port>
sethost rst -iscsi_name <iscsi_name> <node:slot:port>
```

Description

The `sethost` command sets properties on existing system hosts, including options to annotate a host with descriptor information such as physical location, IP Address, Operating System, Model, etc. The command also provides the ability to configure or remove iSCSI CHAP authentication information and to reset a host, aborting all its pending I/O.

Options

-loc <location>

Specifies the location of the host.

-ip <IP address>

Specifies the IP address of the host.

-os <OS>

Specifies the operating system running on the host.

-model <model>

Specifies the model of the host.

-contact <contact>

Specifies the contact information for the host.

-comment <comment>

Specifies any additional information for the host.

-name <hostname>

Specifies the new name of the host up to 31 characters in length.

-persona <hostpersonaval>

Sets the host persona that specifies the personality for all ports which are part of the host set. This selects certain variations in scsi command behavior which certain operating systems expect.

<hostpersonaval> is the host persona id number with the desired capabilities. These can be seen with `showhost -listpersona`.

The following options are for use with the initchap and targetchap subcommands:

-f

Do not ask for confirmation before performing the operation.

-chapname <chapname>

Used to specify the initiator or target CHAP name. If this option is not specified, then the initiator CHAP name defaults to the host name and the target CHAP name defaults to the HPE 3PAR system name.

-hex

The CHAP secret is treated as a hex number.

The following options are for use with the removechap subcommand:

-target

Removes only the target CHAP authentication.

Specifiers

<hostname>

Name of the host up to 31 characters in length.

<pattern>

Specifies that properties are set for all hosts matching the specified pattern.

<secret>

The CHAP secret for the host or the target. If -hex is specified, it is treated as a hex number. Otherwise it should be a printable ASCII string 12 to 16 characters in length with no spaces, or 16 bytes in HEX.

<WWN | iscsi_name>

The World Wide Name (WWN) or iSCSI name of the host which should be reset or for which Host Explorer data should be cleared.

<node:slot:port>

node

Specifies the node using a number from 0 through 7.

slot

Specifies the PCI slot in the specified node. Valid ranges are:

0 - 9 for the HPE 3PAR 10000 storage system platform.

0 - 3 for the HPE 3PAR 7200 and HPE 3PAR 7400 storage system platforms.

0 - 3 for the HPE 3PAR 8200 and HPE 3PAR 8400 storage system platforms.

port

Specifies the port using a number from 1 through 4.

Subcommands

initchap

Sets the initiator CHAP authentication information on one or more hosts.

targetchap

Sets the target CHAP authentication information on one or more hosts

removechap

Removes CHAP authentication on one or more hosts. By default, this removes all CHAP information for all specified hosts. Using this subcommand with the -target option removes only target CHAP information.

clearagent

Clears any Host Explorer data associated with the host.

rst

Resets a host or host initiator, aborting all pending commands from the host initiator(s) and releasing any SCSI-2 reservations held by the initiator(s).

Authority

- Super, Edit
- Any role granted the `host_set` right

Usage

- Access to all domains is required for the clearagent subcommand.
- The CHAP configuration operations are applied to all hosts whose names match one or more of the specified <hostname> or <pattern>. The patterns are treated as glob-style (shell-style) patterns. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.)
- The options that allow for adding descriptor information are for annotation purposes only; the information provided here is not actively used by the storage server.
- Remove a descriptor by passing an empty string to the command.
- Verify modification of host properties by issuing the **showhost** command.

Examples

The following examples change the settings of a host:

```
cli% sethost -contact "Joe Smith" -model "Sun Ultra 60" queasy10
cli% sethost initchap "MyChapSecret" queasy10
cli% sethost targetchap -hex "3031323334353637383939303132333435" queasy10
cli% sethost removechap -target queasy10
cli% sethost clearagent 210100E08B32A58A
cli% sethost -persona 1 queasy10
cli% sethost rst -wwn 210100E08B32A58A 0:2:3
```

sethostset

Syntax

```
sethostset [options <arg>] <setname>
```

Description

The `sethostset` command sets the parameters and modifies the properties of a host set.

Options

`-comment <comment>`

Specifies any comment or additional information for the set. The comment can be up to 255 characters long. Unprintable characters are not allowed.

`-name <newname>`

Specifies a new name for the host set, using up to 27 characters in length.

Specifiers

`<setname>`

Specifies the name of the host set to modify.

Authority

- Super, Edit
- Any role granted the `hostset_annotate` right can only annotate the comment of a host set
- Any role granted the `hostset_set` right can set any host set property

Usage

None.

Examples

The following example shows how to rename a set from foo to bar:

```
cli% sethostset -name bar foo
```

The following example shows how to change the comment on a set:

```
cli% sethostset -comment "This used to be set foo" bar
```

setlicense

Syntax

```
setlicense [options <arg>]
```

Description

The `setlicense` command sets the license key information.

Options

`-f <filename>`

Specifies the file from which the license key is read.

`-noconfirm`

Specifies that the system does not prompt for confirmation for the new license key.

Specifiers

None.

Authority

- Super, Service
- Any role granted the `license_set` right

Usage

- Access to all domains is required to run this command.
- HPE's license terms and conditions must be accepted before proceeding with this command.
- This command prompts for a new license key. To finish entering the license key, press ENTER on a blank line.
- When the license key is being interpreted, all characters other than letters (without case sensitivity) and numbers are ignored, and the letters are not case-sensitive.
- After the new license key has been entered, the changes between the existing license key and the new license key are displayed. There is a prompt to confirm the changes unless the -noconfirm option is given, in which case the information is not displayed, and the new license key will be entered immediately.
- The entered license key is accepted only if it is recognized as a valid key.
- A valid license key includes an appropriate serial number and is associated with the number of nodes in the system for which the license key is being entered.

Examples

The following example displays the setting of a license key:

```
cli% setlicense

If this software is being provided to you for a limited evaluation period,
then your license shall be governed by the current HPE Software License
Terms (or as otherwise agreed between us) with the exception that the
term of the license shall expire upon the earlier of the evaluation period
notified to you or 120 days. Upon expiration of the license, you must cease
using the software and HPE reserves the right to disable the software
without notice. By using or activating the software you are agreeing to
these terms.

Do you agree to these terms and conditions? y=yes n=no: yes

Please enter the new license key below. When finished, press enter on an
empty line. If the key is entered by hand, note that characters other than
letters and numbers are ignored, and the key is not case-sensitive.

60R3-0C1G...
```

setmaint

Syntax

```
setmaint [options <maintype>]
```

Description

Allows modification of the maintenance window record with the specified options for the maintenance type.

Options

-comment <comment>

Specifies any comment or additional information for the maintenance window record. The comment can be up to 255 characters long. Unprintable characters are not allowed.

-duration <time>{m|h}

Extends the duration of the maintenance window record by the specified time. May be specified in minutes (e.g. 20m) or hours (e.g. 6h). If unspecified, the window duration is unchanged. This option cannot be specified with the `-end` option.

-end

Ends the window record for the specified maintenance type. If the maintenance window record has been created more than once with `createmaint`, this option reduces its reference count by one without ending the window record. This option cannot be specified with the `-duration` option.

Specifiers

<mainttype>

The maintenance type for the maintenance window record to be modified.

Authority

Super, Service

Any role granted `maint_set` rights.

Usage

If no maintenance window record exists with the specified maintenance type, the command fails.

Examples

To end the maintenance window record Other:

```
cli% setmaint -end Other
```

To modify the comment for the maintenance window record Restart:

```
cli% setmaint -comment "Node 1 restart" Restart
```

setnet

Syntax

The syntax for the `setnet` command can be one of the following:

```
setnet startaddr <old_IP> <new_IP> <new_netmask>|<prefix_len>
setnet startgateway <new_gateway>
setnet finish [-f]
setnet abort
```

```
setnet cleargateway [-f] [-v4] [-v6]
setnet speed <IP_addr> auto|<mbps> <duplex>
setnet failoverping <IP_addr> {<ping_addr>|none}
setnet ntp {none | [-add|-set|-remove] [-force] <server_addr>}
setnet changenode [<node_ID>]
setnet addaddr <old_IP> <new_IP> {<new_netmask>|<prefix_len>}
setnet removeaddr <old_IP>
setnet dns {none | [-add|-remove] <server_addr>}
setnet disableports {yes|no}
```

Description

The `setnet` command sets the administration network interface configuration.

Options

`-add`

Specifies that the server should be added to an existing set.

`-remove`

Specifies that the server should be removed from an existing set.

`-set`

Specifies that the server should be the only server in a set.

`-f`

Specifies that the operation is forced even if verification has not occurred. This option can only be used with the finish and cleargateway subcommands.

Specifiers

`<old_IP>`

Specifies an existing IP address that is to be changed. This specifier is used in conjunction with the <new_IP> and <new_netmask> specifiers and can only be used with the startaddr, addaddr and removeaddr subcommands.

`<new_IP>`

Specifies a new IP address to which the system is configured. This specifier is used in conjunction with the <old_IP> and <new_netmask> specifiers and can only be used with the startaddr and addaddr subcommands.

`<new_netmask>`

Specifies a new netmask to which the system is configured. This specifier is used in conjunction with the <old_IP> and <new_IP> specifiers and can only be used with the startaddr and addaddr subcommands.

`<prefix_len>`

Specifies a new prefix_len to which the system is configured. This specifier is used in conjunction with the <old_IP> and <new_IP> specifiers. This specifier can only be used when <new_IP> is of IPv6 type and can only be used with the startaddr and addaddr subcommands.

`<new_gateway>`

Specifies the IP address of the new gateway for the system. This specifier can only be used with the startgateway subcommand.

auto|<mbps> <duplex>

Specifies that the speed of the network interface is either auto negotiated (auto), or specified manually using the <duplex> and <mbps> specifiers.

<mbps>

Specifies the speed of the network interface. Valid values are either 10, 100, or 1000. This specifier can only be used with the <duplex> specifier and with the speed subcommand.

<duplex>

Specifies the duplex of the network interface. Valid values are either half or full. This specifier can only be used with the <mbps> specifier and with the speed subcommand.

<IP_addr>

Specifies the IP address of the node. This specifier can only be used with the failoverping subcommand. <IP_addr> must be an IPv4 address.

<ping_addr>|none

Specifies that during an IP failover, a ping either be sent to the specified IP address (<ping_addr>) or not sent at all (none). This specifier can only be used with the failoverping subcommand.

<ping_addr> must be an IPv4 address.

[<node_ID>]

Specifies the node, by ID, that has an active Ethernet interface. This specifier can only be issued with the changenode subcommand. This specifier is not required.

Subcommands

startaddr

Specifies that the system start switching the old IP address (<old_IP> specifier) to the new IP address (<new_IP> specifier) with the specified netmask (<new_netmask> specifier).

startgateway

Specifies that the gateway is immediately set to the specified IP address if no gateway is currently defined. If a gateway is currently defined, the system starts switching the old gateway to the new gateway.

Both IPv4 and IPv6 can be changed with this command.

finish

Specifies that outstanding changes from the startaddr and startgateway subcommands be completed.

abort

Specifies that any attempt to configure a new IP address or gateway fails. The system returns to its previous state.

cleargateway

Specifies that all the existing system gateways are removed if neither the -v4 or -v6 options are given. The -v4 option specifies that only the IPv4 gateway be removed. The -v6 option specifies that only the IPv6 gateway be removed. The -f option overrides the prompt confirming to clear the gateway(s).

speed

Specifies that the network interface is set to the specified speed and duplex as indicated with the <mbps> and <duplex> specifiers.

failoverping

Specifies that on IP failover, a ping is sent to the specified IP address as indicated with the <ping_addr> specifier.

ntp {none | [-add|-set|-remove] [-force] <server_addr>}

Specifies the NTP server the system should use to synchronize its clocks. The server may be an IP address or hostname. To add or set a NTP server it must be resolvable. The -force option can be used to override this check.

changenode

Forces the system to change which node has an active Ethernet interface.

If a node ID is specified, it switches to that node. Otherwise, any node with a connected Ethernet interface is chosen.

addaddr

Specifies that the new IP (<new_IP> specifier)is added to the interface that is currently assigned the old_IP (<old_IP> specifier). The old IP and new IP should be of different IP versions (IPv4 or IPv6).

removeaddr

Specifies that the old IP (<old_IP> specifier) be de-configured.

The old IP can only be an IPv6 address.

dns {none | [-add|-remove] <serveraddr>}

Specifies the DNS server the system should use to resolve host names. The server must be specified as an IP address.

disableports {yes|no}

Specifies the network configuration functionality where if the option is specified as "yes" will disable the non-encrypted ports, if option "no" is specified, it will enable the non-encrypted ports. Disabling non-encrypted ports will also prevent the service processor from monitoring events, which will prevent the generation of email notifications about system issues.

Authority

- Super, Service
- Any role granted the `net_set` right

Usage

- Access to all domains is required to run this command.
- If successful, this command causes any clients currently connected to lose their connection.
- To make it possible to change the network configuration without running the risk of losing contact with the system because of misconfiguration, the `setnet` command uses a two step process. When a new IP address is specified with the `setnet startaddr` command, the system is configured to listen to both the old and new IP addresses. When a new gateway is specified with the `setnet gateway` command, the system switches between the old and new gateways when it sees packets addressed to it being routed through those gateways.
- After a connection has been made with the new configuration, the `setnet finish` command can be used to remove the old configuration. While in the middle of this process, additional work must be done by the system. It is preferable to run the `setnet finish` command after the new configuration has been verified.
- When changing gateways, starting a connection takes longer than usual, as the first reply packet is typically routed through the previously used gateway address.
- NTP servers may be an IP address or hostname. When adding or setting a NTP server the 3PAR storage system checks to see if the NTP server is resolvable.

Examples

The following example displays the switching of the old IP address with a new IP address and netmask:

```
cli% setnet startaddr 10.0.23.42 192.168.5.218 255.255.252.0
Change of IP address successfully started.
```

The following example displays the gateway being cleared:

```
cli% setnet cleargateway
If the machine that the HPE 3PAR CLI is running on is not on the same
subnet as the HPE 3PAR Storage system it is connected to, clearing the
gateway will render the storage system unreachable from that machine,
and any future connections will need to be made from a system which
is on the same subnet as the storage system.

Are you sure you want to clear the gateway (y/n)?
y
Gateway modified successfully.
```

The following example modifies the disableports option:

```
cli% setnet disableports yes
Disabling non-encrypted ports will disable SP event handling,
Recovery Manager for VMWare and SRA. Disabling SP event handling will
prevent support personnel from being notified of system problems.
Disabling non-encrypted ports should only be done if there is a strict
requirement for all connections to be encrypted.

Are you sure you want to disable non-encrypted ports?
select q=quit y=yes n=no: y
```

setnode

Syntax

```
setnode ps <PS_ID> [options] <node_ID>
```

Description

The `setnode` command sets properties of the node components such as serial number of the power supply.

Options

`-s <serial_number>`

Specifies the serial number. It can be up to 8 characters in length.

Specifiers

`<PS_ID>`

Specifies the power supply ID.

`<node_ID>`

Specifies the node ID.

Subcommands

`ps`

Sets the power supply properties.

Authority

- Super, Service
- Any role granted the `node_set` right

Usage

- Access to all domains is required to run this command.
- At least one option must be specified.

Examples

The following example displays setting the serial number of the node power supply:

```
cli% shownode ps
Node  PS  -Serial-  -PSState--  FanState ACState DCState -BatState- ChrgLvl(%)
   0   0  --        NotPresent  --       --      --      NotPresent          0
   0   1  FFFFFFFF  OK          OK       OK      OK      OK                  0
   1   0  FFFFFFFF  OK          OK       OK      OK      OK                  0
   1   1  --        NotPresent  --       --      --      NotPresent          0
cli%
cli% setnode ps 1 -s 12345678 0
cli% shownode -ps
Node  PS  -Serial-  -PSState--  FanState ACState DCState -BatState- ChrgLvl(%)
   0   0  --        NotPresent  --       --      --      NotPresent          0
   0   1  12345678  OK          OK       OK      OK      OK                  0
   1   0  FFFFFFFF  OK          OK       OK      OK      OK                  0
   1   1  --        NotPresent  --       --      --      NotPresent          0
cli%
cli% setnode ps 0 -s aabbccdd 1
cli% shownode -ps
Node  PS  -Serial-  -PSState--  FanState ACState DCState -BatState- ChrgLvl(%)
   0   0  --        NotPresent  --       --      --      NotPresent          0
   0   1  12345678  OK          OK       OK      OK      OK                100
   1   0  AABBCCDD  OK          OK       OK      OK      OK                100
   1   1  --        NotPresent  --       --      --      NotPresent          0
```

setpassword

Syntax

```
setpassword [options <arg>]
```

Description

The `setpassword` command allows a user with Super role to change the password for any user and create a password file on a client. Edit-, browse-, or service-level users can use the `setpassword` command to change their own passwords or save their password files on a client.

Options

`-u <username>`

Specifies the login name of the user whose password is being changed. If a login name is not specified, the command defaults to the current user.

`-minlen <N>`

Sets the minimum password length for all users. Where N is an integer from 6 to 32. This option can only be used by itself. The current minimum password length is displayed by `showsys -d`.

Specifiers

None.

Authority

- Super, Service, Edit, Browse
- Any role granted the `password_setany` right can set any user password.

- Any role granted the `password_setown` right can only set their own password.
- Only users with Super role can set the minimum password length.

Usage

- Only a user with Super role can set another user's password and set the minimum password length. The default minimum password len is 15.
- Passwords can be no longer than 32 printable characters.
- The format of the entry in the file is <username> <encrypted_password>. This file may be referenced by the TPDPWFILE environment variable or -pwf command line option for subsequent commands.
- Without any options, the command will prompt to change the invoking user's password on the storage system.
- Changing a user's password has no effect on SSH access if the user has set a valid key with the **`setsshkey`** command. Until the SSH key is removed the user will not have to provide the new password.

Examples

The following example displays the prompts encountered when changing a user's (user1) password:

```
cli% setpassword -u user1
password:
Old password:
```

The following example displays the how to set the minimum password length:

```
cli% setpassword -minlen 12
```

setpd

Syntax

```
setpd ldalloc on|off <PD_ID>...
```

Description

The `setpd` command marks a PD as allocatable or non allocatable for logical disks (LD).

Options

None.

Specifiers

`<PD_ID>...`

Specifies the PD identification using an integer.

Subcommands

`ldalloc on|off`

Specifies that the PD, as indicated with the PD_ID specifier, is either allocatable (on) or non-allocatable (off) for LDs.

Authority

- Super, Service
- Any role granted the `pd_set` right

Usage

- Access to all domains is required to run this command.
- This command can be used when the system has disks that are not to be used until a later time.
- Verify the status of PDs by issuing the `showpd -state` command (see the **showpd** command).

Examples

The following example displays PD 0 marked as non allocatable for LDs:

```
cli% setpd ldalloc off 0
```

setqos

Syntax

```
setqos [options] [{{vvset|domain}:{<name>|<pattern>}|sys:all_others}]...
```

Description

The `setqos` command creates and updates QoS rules in a system.

Options

-pri {high|normal|low}

Sets the QoS scheduling priority of the QoS rule (of this target object).

The default priority is "normal".

-io {none|[<MinGoal>[k|K|m|M|g|G]-]<MaxLimit>[k|K|m|M|g|G]]}

Sets the I/O issue count Min goal and Max limit for QoS throttling.

Sets the I/O issue count Min goal and Max limit for QoS throttling. If only <MaxLimit> is given, sets both I/O issue count rate Min goal and Max limit to the given value. If "none" is specified, there is no limit on I/O issue count. Note even when there is no limit for I/O issue count, I/O-bandwidth-count based throttling (-bw) can still dynamically put a limit on it. The integer can optionally be followed with k or K to indicate a multiple of 1000, m or M to indicate a multiple of 1,000,000, or g or G to indicate a multiple of 1,000,000,000.

-bw {none|[<MinGoal>[k|K|m|M|g|G]-]<MaxLimit>[k|K|m|M|g|G]]}

Sets the I/O issue bandwidth rate Min goal and Max limit for QoS throttling. If only <MaxLimit> is given, sets both I/O issue bandwidth rate Min goal and Max limit to the given value. If "none" is specified, there is no limit on I/O issue bandwidth rate. Note even when there is no limit for I/O issue count, I/O-bandwidth-rate based throttling (-io) can still dynamically put a limit on it. The default unit is byte. The integer can optionally be followed with k or K to indicate a multiple of 1000, m or M to indicate a multiple of 1,000,000, or g or G to indicate a multiple of 1,000,000,000.

-lt {<target>[s|ms|us]|default}

Sets the QoS I/O committing target latency to be <target>. <target> is in units of milliseconds (ms) if no suffix is specified. If "default" is specified, the latency goal will be cleared.

{-on|-off}

Controls QoS scheduling for the target object. By default QoS scheduling is "on".

-clear

Turns "off" QoS scheduling of the rule and clear its setting.

-vv {<VV_name>|<pattern>}[,{<VV_name>|<pattern>}]...

Applies only to QoS rules whose targets include virtual volumes with names matching any of the names or patterns specified.

Specifiers

{{vvset|domain}:{<name>|<pattern>}|sys:all_others}

The target objects of QoS setting. <name> and <pattern> refer to the target object name, as listed in the show{vvset|domain} commands. If sys:all_others is specified, the QoS rule will be applied to the group of all virtual volumes, which do not have any specific QoS "on".

Authority

- Super, Edit
- Any role granted the `qos_set` right can set QoS configurations.

Usage

- Maximum limit - this is the maximum amount of IOPS or bandwidth, or both, which a given VVset or Domain is allowed to achieve. Best practice is to use the System Reporter data to quantify the volumes' performances and set QoS rules accordingly.
- Minimum goal - this is the minimum amount of IOPS or bandwidth, or both, below which the system will not throttle a given VVset or Domain in order to meet the latency goal of a higher priority workload. These VVsets or Domains may consume more IOPS/bandwidth than what the goal is (up to the Max), but will be throttled to the given goal as the system gets busier.
- Latency Goal - this is the Svctime goal the system will target to fulfill for a given QoS rule. In order for the goal to work, rules with a minimum goal must exist so the system can throttle those workloads.
- Three Priority Levels exist: high, normal and low. As the system gets busier it will start targeting lower priority workloads and throttling their performances to meet higher priority workloads' latency goals. High priority level should be used against critical applications, lower priority on less critical applications.
- For all the <pattern> fields, the patterns are glob-style (shell-style) patterns. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.) When -vv option is present, the setting is only applied to target objects with existing QoS configuration. To limit the new settings to target objects with existing QoS configuration, it is useful to specify -vv * in the command line.
- QoS rules only manage host I/O.
- For this command KB = 1000 bytes.
- Using this command to create new QoS rules requires a Priority Optimization license. Contact your local HPE representative for information.

Examples

The following example sets the IOPs Min goal and Max limit of vvset vsa to 300:

```
cli% setqos -io 300 vvset:vsa
```

The following example turns "off" the QoS scheduling for vvset vsa:

```
cli% setqos -off vvset:vsa
```

The following example removes all QoS settings of all vvsets:

```
cli% setqos -clear vvset:*
```

setrcopygroup

Syntax

The syntax for the `setrcopygroup` command can be one of the following:

- `setrcopygroup pol [option] [<pattern>] <policy> [<group_name>]`
- `setrcopygroup period [option] [<pattern>] <period_value> <target_name> [<group_name>]`
- `setrcopygroup mode [option] [<pattern>]<mode_value> <target_name> [<group_name>]`
- `setrcopygroup snap_freq [option] [<pattern>]<freq_value> <target_name>[<group_name>]`
- `setrcopygroup <dr_operation> [options] [<pattern>] [<target_name|group_name>...]`
- `setrcopygroup cpg -usr_cpg <cpg_name> <target_name>:<cpg_name>... -snp_cpg <cpg_name> <target_name>:<cpg_name>... <group_name>`
- `setrcopygroup cpg -usr_cpg_unset -snp_cpg_unset <group_name>`
- `setrcopygroup vvol -removetest <group_name>`

Description

The `setrcopygroup` command performs the following actions:

- Sets the policy of the Remote Copy volume group for dealing with I/O failure and error handling.
- Switches the direction of transfer between volume groups.
- Sets a resynchronization period for volume groups in asynchronous periodic mode.
- Sets the minimum stop period and defined stop order for groups in asynchronous streaming mode.
- Sets the group's mode.
- Sets a frequency which Remote Copy takes coordinated snapshots in asynchronous streaming mode.

Options

`-t <tname>`

When used with `<dr_operation>` subcommands, specifies the target to which the `<dr_operation>` command applies. This is optional for single target groups, but is required for multi-target groups. If no groups are specified, it applies to all relevant groups. When used with the `pol` subcommand, specified for a group with multiple targets, then the command only applies to that target; otherwise it will be applied to all targets.

NOTE: The `-t` option, without the groups listed in the command, will only work in a unidirectional configuration. For bidirectional configurations, the `-t` option must be used along with the groups listed in the command.

`-f`

Does not ask for confirmation for disaster recovery commands.

`-nostart`

Specifies that groups are not started after role reversal is completed. This option can be used for `failover`, `recover`, and `restore` subcommands.

`-nosync`

Specifies that groups are not synced after role reversal is completed through the `recover`, `restore`, and `failover` specifiers.

-discard

Specifies not to check a group's other targets to see if newer data should be pushed from them if the group has multiple targets. The use of this option can result in the loss of the most recent changes to the group's volumes and should be used carefully. This option is only valid for the `failover` specifier.

-nopromote

This option is only valid for the `failover` and `reverse` specifiers. When used with the `reverse` specifier, specifies that the synchronized snapshots of groups that are switched from primary to secondary should not be promoted to the base volume. When used with the `failover` specifier, it indicates that snapshots of groups that are switched from secondary to primary should not be promoted to the base volume in the case where all volumes of the group were not synchronized to the same time point.

The incorrect use of this option can lead to the primary secondary volumes not being consistent.

-nosnap

Specifies that snapshots are not taken of groups that are switched from secondary to primary. Additionally, existing snapshots are deleted if groups are switched from primary to secondary. The use of this option may result in a full synchronization of the secondary volumes. This option can be used for `failover`, `restore`, and `reverse` subcommands.

-stopgroups

Specifies that groups are stopped before running the `reverse` subcommand.

-local

The `-local` option only applies to the `reverse` operation, and then only when the `-natural` or `-current` options to the `reverse` operation are specified. Specifying `-local` with the `reverse` operation and an associated `-natural` or `-current` option will only affect the array where the command is issued, and will not be mirrored to any other arrays in the Remote Copy configuration.

-natural

Specifies the `-natural` option with the `reverse` operation changes the role of the groups but not the direction of data flow between the groups on the arrays. For example, if the role of the groups are "primary" and "secondary", issuing the `-natural` option with the `reverse` operation will result in the role of the groups becoming "primary-rev" and "secondary-rev" respectively. The direction of data flow between the groups is not affected, only the roles. Since the `-natural` option does not change the direction of data flow between groups, it does not require the groups be stopped.

-current

Specifies the `-current` option with the `reverse` operation changes both the role and the direction of data flow between the groups. For example, if the roles of the groups are "primary" and "secondary", issuing the `-current` option to the `reverse` operation will result in the roles of the group becoming "secondary-rev" and "primary-rev" respectively, and the direction data flows between the groups is reversed. Since the `-current` option actually reverses the direction of data replication, it requires the group be stopped.

Both the `-natural` and `-current` options must be used with care to ensure the Remote Copy groups do not end up in a non-deterministic state (like "secondary", "secondary-rev" for example) and to ensure data loss does not occur by inadvertently changing the direction of data flow and resyncing old data on top of newer data.

-waittask

Waits for all tasks created by this command to complete before returning. This option applies to the `failover`, `recover`, `restore`, and `reverse` subcommands.

-pat

Specifies that specified patterns are treated as glob-style patterns and all Remote Copy groups matching the specified pattern will be set. The `-pat` option can specify a list of patterns. This option must be used if `<pattern>` specifier is used.

-usr_cpg <cpg name> <target_name>:<cpg_name>

Specifies the local user CPG and target user CPG that will be used for volumes that are auto-created. The local CPG will only be used after `failover` and `recover`.

-snp_cpg <cpg name> <target_name>:<cpg_name>

Specifies the local snap CPG and target snap CPG that will be used for volumes that are auto-created. The local CPG will only be used after `failover` and `recover`.

-usr_cpg_unset

Unsets all user CPGs that are associated with this group.

-snp_cpg_unset

Unsets all snap CPGs that are associated with this group.

-forceppfailover

Specifies that the Peer Persistence failover operation is forced overriding data inconsistency warnings. All I/O to the existing primary volumes should be quiesced when using this option.

Specifiers

<pattern>

Specifies a glob-style pattern. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.) If this specifier is not used, the `<group_name>` or `<target_name>` specifier must be used.

<policy>

Specifies the policy to assign to the group. This specifier can only be used with the `pol` subcommand.

Valid policies are:

auto_failover

Configures automatic failover on a Remote Copy group. This feature will perform an automatic failover on a Remote Copy group when used in conjunction with the Quorum Witness functionality.

Secondary volumes will be subject to a Remote Copy failover operation. Any of the secondary volumes that are exported to attached hosts will become accessible.

If a network fails and prompts a failover, any primary volumes will be placed into a failsafe state to prevent data corruption and inconsistency between primary and secondary volumes.

no_auto_failover

Remote-copy groups will not be subject to automatic failover (default).

auto_failover_ext

Configure extended automatic failover on a remote copy group. This feature will perform an automatic failover any time the primary system becomes isolated from the quorum witness and communication between arrays subsequently goes down (for example, during a rolling network failure that is outside the fail_timeout window where the quorum witness link fails before the remote copy links fail). The remote or secondary system initiates a failover when the remote copy links are down and information from the quorum witness indicates that the primary is not communicating with the quorum witness.

The `auto_failover` policy is a prerequisite and is automatically set when the `auto_failover_ext` policy is enabled. Clearing the `auto_failover` policy also clears this policy setting.

no_auto_failover_ext

Remote-copy groups will not be subject to extended automatic failover (default).

auto_recover

Specifies that if the Remote Copy group is stopped as a result of the Remote Copy links going down, the group is restarted automatically after the links come back up. If this policy is enabled for a group while the group is stopped after link failures, it will only be started when the links come up for the failed target. If the links are already up at the time the policy is set, then the group will not be restarted at that time.

no_auto_recover

Specifies that if the Remote Copy is stopped as a result of the Remote Copy links going down, the group must be restarted manually after the links come back up (default).

auto_synchronize

Remote Copy system automatically recovers and synchronizes all volumes in the group after a system failover, for either automatic or manual failover scenarios. Synchronization occurs after system recovery completes and the Remote Copy links recover. This policy also allows the failover command to be used when synchronous groups are started and are online. It is no longer necessary to stop the synchronous groups before initiating a failover command to the secondary system.

no_auto_synchronize

Remote Copy groups will not be subject to automatic synchronization after failover (default).

over_per_alert

If a synchronization of a periodic Remote Copy group takes longer to complete than its synchronization period, then an alert will be generated. This is the default behavior.

no_over_per_alert

If a synchronization of a periodic Remote Copy group takes longer to complete than its synchronization period, then an alert will **not** be generated.

path_management

Volumes in the specified group will be enabled to support ALUA. The Target Port Group states of the volumes will be presented as ACTIVE on the primary and as STANDBY on the secondary.

no_path_management

ALUA behavior will be disabled for volumes in the group. The Target Port Group state of the volumes will be presented as ACTIVE (default).

mt_pp

Specifies that the group is participating in a multi-target Peer Persistence configuration. The group must have two targets, one of which must be synchronous. The path_management and auto_failover policies will also be set for the synchronous group target.

no_mt_pp

Disables the multi-target Peer Persistence configuration. This will also remove the path_management and auto_failover policies.

active_active

Volumes in the specified group will be enabled for host reads and writes from both the primary and the secondary. This policy will also set the `auto_synchronize` and `auto_recover` policies to the group.

no_active_active

Disable the active_active policy, allowing host writes only to the primary. Host reads from secondary depends on whether the path_management is still set. Disabling this policy, does not disable auto_synchronize and auto_recover policies. (default).

<group_name>

Specifies the name of the volume group whose policy is set, or whose target direction is switched.

<target_name>

Specifies the target name for the target definition created with the createcopytarget command.

<mode_value>

Specifies the mode, sync, periodic, or async to which the group is set. This specifier can only be used with the mode subcommand.

<period_value>s|m|h|d

Specifies the time period in units of seconds (s), minutes (m), hours (h), or days (d), for automatic resynchronization (e.g., 14h for 14 hours). The time can be a minimum of 30 seconds (for asynchronous periodic groups) and up to a maximum of one year, or set to zero. This specifier can only be used with the period subcommand.

For asynchronous periodic mode, this value is used to specify the time period for automatic resynchronization. A <period_value> of 0 will prevent automatic resynchronization. The minimum period value for asynchronous streaming groups is thirty seconds.

For asynchronous streaming mode, the <period_value> is used to define the order in which these groups will be stopped and restarted automatically, as well as the minimum elapsed time after which each specific group will be restarted once the replication resources are restored. Groups with the largest <period_value> will be stopped first and restarted last. The <period_value> can be a minimum of 30 seconds, the default value is 5 minutes. Values between 1 and 29 seconds will be rejected. A <period_value> of 0 is allowed and will prevent automatic restarting of the group; these groups will need to be started manually. If groups have the same <period_value>, the group using the most replication resources will be stopped first and restarted last.

When changing mode from asynchronous streaming to asynchronous periodic, <period_value> will be reconfigured if it is not valid. For example if an asynchronous streaming mode group has a <period_value> of 30 seconds when changed to asynchronous periodic mode, <period_value> will be changed to 5 minutes.

<freq_value>s|m|h|d

Specifies the time period in units of seconds(s), minutes (m), hours (h), or days (d), for automatic coordinated snapshots taken on both local and remote systems. The time must be longer than or equal to 30 seconds and not more than one year in duration. If the value is set to zero, automatic coordinated snapshots will be turned off. Only applicable in asynchronous streaming mode.

Subcommands

pol

Sets the policy of the Remote Copy volume groups for dealing with I/O failure and error handling.

period

Specifies that groups that are in asynchronous periodic mode should be periodically synchronized in accordance with the specified <period_value>.

Groups that are in asynchronous streaming mode will transition to "Logging" state when replication resources drop below sustainable limits. The <period_value> is used to define the order that these groups will be stopped and restarted automatically, as well as the minimum elapsed time each specific group will be restarted once the replication resources are restored.

Groups with the largest <period_value> will be stopped first and restarted last. A <period_value> of 0 will prevent automatic restarting of the group; these groups will need to be started manually.

snap_freq

Specifies the interval at which Remote Copy takes coordinated snapshots. Only applicable in asynchronous streaming mode. The interval can be a minimum of five minutes and up to a maximum of one year.

mode

Specifies the mode to which the volume group is set.

cpg

Specifies the local and target cpg.

<dr_operation>

Specifies the operation of the group(s). Valid operations are:

reverse

Changes the natural and current direction of all specified groups. The operation is mirrored resulting in a direction change on both systems. This option is very flexible depending on the options that are provided; however it should not be used as a part of the normal disaster recovery process. It should be used only to reverse direction on consistent copies.

failover

Changes secondary volume groups to primary volume groups on the active system in the event of a server failure. If the group has multiple targets, it will also attempt to pull more recent data from other targets and start Remote Copy to those targets when complete.

switchover

Migrates the Remote Copy group from primary to secondary without impacting the host IO. This command requires that associated hosts are connected to both the primary and secondary arrays. The WWNs of primary and secondary volumes are consistent and the host persona of the host must support RTPG.

recover

Used for groups on which the `failover` operation has already been run. Changes matching primary volume groups on the backup system to secondary volume groups and then starts and synchronizes all groups.

restore

Used for groups on which the `failover` operation has already been run. The `restore` command will implicitly execute the `recover` operation on any group after a failover, which has not yet been recovered, prior to performing the restore. It identifies and differentiates between those groups requiring a simple restore and those requiring both a recover and then a restore, and orchestrates the execution of the required tasks to return all groups to their natural direction and start them. In addition to being context-specific, the execution of the `restore` operation now takes a best effort approach. If one or more groups in the list of groups passed to the operation is not in the correct state for a restore or a combined recover and restore, it will return to the user with an error code. However, it will process those groups that are deemed eligible for the simple or combined operation before returning. If a primary AA/PP group is in Failsafe mode, `restore` can be called on the secondary instead of calling `restore`, then `reverse -natural`, and then `switchover`.

override

This command overrides the failsafe state which is applied to Remote Copy group, allowing the associated volumes to be exported to attached hosts.

vvol

This command specifies a VMware vVol specific operation on the Remote Copy group. Valid operations are:

-removetest

Specifies that the INTEST state of this VMware vVol group be cleaned up.

Authority

- Super, Edit
- Any role granted the `rcopygroup_set` right.

Usage

- This command requires the HPE 3PAR Remote Copy Software license. Contact your local service provider for more information.
- When issuing the `setrcopygroup <dr_operation>` command, either the `<group_name>` specifier, the `<target_name>` specifier, or the `<pattern>` specifier must be specified.
- Do not use the `reverse` value of the `<dr_operation>` as part of the normal disaster recovery process.
- The period can be set only for groups whose mode is asynchronous periodic or asynchronous streaming (see **`creatercopygroup`**).
- For groups whose mode is asynchronous streaming, using a period value of 0 means that groups which are stopped due to low levels of replication resources will not be automatically started. You must specify a period using `setrcopygroup period <value>` if you want specific groups to be considered first for stopping and also considered for automatically starting again when availability of replication resources improve.
- The minimum period for asynchronous streaming groups is 30 seconds.
- The default period for asynchronous streaming groups is five minutes.
- Reversing the direction of primary volumes will result in the loss of any data changed after the group was stopped.
- There is no default resynchronization period. For groups whose mode is asynchronous periodic, you must specify a resynchronization period using `setrcopygroup period <value>`, or resynchronizations will not automatically take place.
- The minimum interval for periodic resynchronizations is five minutes.
- If the mirror_config policy is set for this group's target and the group is a primary group, then the `setrcopygroup` command is mirrored to the target when the `period` and `pol` subcommands are used.
- Use the `-nosnap` option when the primary server has failed or where the disks are ruined or in an unknown state. For example, an uncontrolled shutdown can result in loss of data. If you suspect that the primary volumes are not in a known good state, you should use this option to force a FULL RESYNC when the primary system is restored.
- The `-nosnap` option can be used when making a secondary group take over as the primary after a disaster takes down the primary (`setrcopygroup failover -nosnap...`). This option indicates that no incremental resynchronization of the primary group is possible while the primary system is coming back online. Without this option, a snapshot is taken when the secondary server takes over as the primary. That snapshot is used to do an INCREMENTAL synchronization of the primary after it is restored. This assumes that there was no loss of data in the primary volumes when the primary server went down.
- The `switchover` command is only supported for synchronous Remote Copy groups that are started and synced. The command must be issued on the primary array, and upon completion of the command, the Remote Copy group will be restarted in the reverse direction.

Examples

The following example sets the group policy for Group1:

```
setrcopygroup pol auto_recover Group1
```

The following example sets the group policy for all the groups that start with the name testgroup:

```
setrcopygroup pol  -pat testgroup* auto_recover
```

The following example reverses the current direction of secondary group (Group1) so that I/O might be applied to the group after disaster recovery:

```
setrcopygroup failover Group1.r121
```

The following example reverses the current direction of all secondary groups that start with the name testgroup so that I/O can be applied to the groups during disaster recovery:

```
setrcopygroup failover  -pat testgroup*
```

The following example sets Group1 to be automatically synchronized every 30 minutes to System2:

```
setrcopygroup period 30m System2 Group1
```

The following example sets volume groups that start with the name testgroup to synchronize to its asynchronous periodic mode target InServ2 every 30 minutes:

```
setrcopygroup period  -pat testgroup* 30m InServ2
```

The following example sets the local user and snap CPG, and the target usr and snp CPG for the group testgroup:

```
setrcopygroup cpg -usr_cpg u_cpg RC603:remote_u_cpg  -snp_cpg s_cpg RC603:remote_s_cpg testgroup
```

The following example will unset the local user and snap CPG for the group testgroup:

```
setrcopygroup cpg -usr_cpg_unset -snp_cpg_unset testgroup
```

The following example will configure Remote Copy to take coordinated snapshots of all volumes in an asynchronous streaming group every two hours:

```
setrcopygroup snap_freq 2h target1 group1
```

setrcopytarget

Syntax

The syntax for the `setrcopytarget` command can be one of the following:

```
setrcopytarget pol <policy> <target_name>
setrcopytarget name <new_name> <target_name>
setrcopytarget tput <tput_value> <target_name>
setrcopytarget tunelinks <bandwidth> <latency> <target_name>
setrcopytarget autotunelinks <target_name>
setrcopytarget {enable|disable} <target_name>
setrcopytarget witness create [-ssl [-port <port_number>]] [-remote]
<witness_ip> <target>
setrcopytarget witness {start|stop|remove} [-remote] <target>
setrcopytarget witness check [-ssl [-port <port_number>]] [-remote] [-node
node_id] <witness_ip> [target]
```

Description

The `setrcopytarget` command adjusts the characteristics of a target definition.

Options

`-remote`

Uses to forward a witness subcommand to be executed on the remote HPE 3PAR Storage system. When used in conjunction with the "witness check" subcommand, the target must be specified - when executing on the local storage system target specification is not required to check connectivity with the Quorum Witness. This option cannot be used with the `-ssl` option.

`-node`

Used in conjunction with the "witness check" subcommand to test the connectivity to the Quorum Witness through the Quorum Announcer process running on the specified node. Otherwise, the command simply verifies that there is at least one operational route to the witness.

`-ssl`

Used in conjunction with the "`witness create`" and "`witness check`" subcommands to specify SSL connectivity to the Quorum Witness through the Quorum Announcer process. This option cannot be used with the `-remote` option. The `-ssl` option is valid on QW server 4.0 and later, and is required if creating a secure quorum connection. The proper certificates must be installed before creating a secure witness.

`-port`

An optional port number to be used to communicate with SSL to the Quorum Witness application. Valid only when the `-ssl` option is used. The default port is 8443.

Specifiers

`<target_name>`

Specifies the target name for the target definition previously created with the `creatercopytarget` command.

`<policy>`

This specifier can only be used with the pol subcommand. The policy can be one of the following:

`mirror_config|no_mirror_config`

Specifies that all configuration commands (creatercopygroup, removercopygroup, admitrcopyvv, dismissrcopyvv, setrcopygroup pol/period, startrcopygroup, and stoprcopygroup) involving the specified target are duplicated (mirror_config) or not duplicated (no_mirror_config). If not specified, all configuration commands are duplicated.

NOTE: The no_mirror_config specifier must be used only to allow recovery from an unusual error condition and only used after consulting your HPE representative.

`<new_name>`

The new name for the indicated target. This specifier can only be used with the name subcommand.

`<tput_value>`

Specifies the maximum throughput for this target's links, and is used to limit the total throughput of the link. You can optionally specify g or G (gigabytes), m or M (megabytes), or k or K (kilobytes) following the throughput value to indicate size (with no space between the specified value and size type). The default is kilobytes. This specifier can only be used with the tput subcommand. A value of 0 will remove the throughput.

`<bandwidth>`

The measured bandwidth of the connection to the target, specified in kilobytes (KB) per second. This specifier can only be used with the tunelinks subcommand.

<latency>

The measured round-trip latency of the connection to the target, specified in milliseconds (ms). This specifier can only be used with the tunelinks subcommand.

<witness_ip>

The IP address of the Quorum Witness (QW) application, to which the HPE 3PAR Storage system will connect to update its status periodically.

Subcommands

pol

Sets the policy for the specified target using the <policy> specifier.

name

Changes the name of the indicated target using the <new_name> specifier.

tput

Sets the maximum throughput value for each of the target's links.

Applicable to RCIP links only.

tunelinks

Adjusts performance values for the target's links using the <bandwidth> and <latency> specifiers. Applicable to RCIP links only.

autotunelinks

Automatically adjust performance values for the target's links. Applicable to RCIP links only. For consistent results, this must be issued to the target links on both systems.

enable | disable

Enables or disables the target.

witness create

Creates an association between a synchronous target and a Quorum Witness (QW) as part of a Peer Persistence configuration.

witness start|stop|remove

Activates, deactivate, and remove the ATF configuration.

witness check

Checks connectivity to Quorum Witness.

Authority

- Super, Edit
- Any role granted the `rcopytarget_set` right.

Usage

- Access to all domains is required to run this command.
- This command requires the HPE 3PAR Remote Copy Software license. Contact your local service provider for more information.
- If the mirror_config policy is set and the `setrcopytarget` command is issued with the pol subcommand, the duplicated configuration commands cannot be issued on the secondary. Doing so results in an error.

- There must be an active connection between the systems in the Remote Copy pair to issue commands on the primary to be mirrored to the secondary. If there is no connection, the commands return an error.
- The `setrcopytarget` command requires the groups associated with it be stopped prior to using the following options:
  - `setrcopytarget name <new_name> <target_name>`
  - `setrcopytarget tput <throughput> <target_name>`
- The `setrcopytarget` command with the following arguments can be run without bringing down its Remote Copy groups:
  - `setrcopytarget pol <policy> <target_name>`
  - `setrcopytarget tunelinks <bandwidth> <latency> <target_name>`
  - `setrcopytarget autotunelinks <target_name>`
- Under normal operating conditions the mirror_config policy must never be changed to no_mirror_config. This policy option is included only as a method to correct several unusual error conditions that might occur in the course of operation which results in a mismatch in configuration between the two sides of a Remote Copy pair. For instance, it is possible for a group to be created, or a volume to be added to a group, only on one side of the pair if the operation is interrupted by a network failure. In such cases, it might be necessary to temporarily change the policy to no_mirror_config to bring the configurations into alignment. After being corrected the mirror_config policy must be immediately restored. Such operations must only be undertaken on the advice of a HPE representative.
- The `setrcopytarget` witness subcommands are provided to locally create and associate a target with a Peer Persistence quorum and to manage the operational state of the quorum.
- The results of tuning RCIP links are reflected in the output of:
- `showrcopy -d targets`
- The "Socks" field shows the number sockets allocated for the specified tune parameters, and the "Buffsz" field shows the size (in bytes) of the buffers.
- Targets whose links are subject to "autotunelinks" will show a "Buffsz" of zero.
- The tunelinks subcommand will disable the autotunelinks behavior.
- The default performance values for an RCIP target before taking any tuning action are equivalent to those produced by running:
- The proper certificates must be installed before creating a secure witness.

Examples

The following example will set the throughput of each link of target InServ-B to 5 megabytes per second:

```
cli% setrcopytarget tput 5M InServ-B
```

The following example shows how to check the connectivity to the Quorum Witness application, IP address 10.0.0.1 through the Quorum Announcer process on node 0 of the mirror HPE 3PAR storage system over target tar1:

```
cli% setrcopytarget witness check -remote -node 0 10.0.0.1 tar1
```

The following example shows how to create and associate a Peer Persistence quorum with target tar1. The Quorum Announcer process on the local nodes will post health updates to the Quorum Witness application at IP address 10.0.0.1 and will monitor the health of the mirror storage system over target tar1 through its health updates to the witness at 10.0.0.1:

```
cli% setrcopytarget witness create 10.0.0.1 tar1
```

The following example shows how to tune RCIP links for a bandwidth of 2048 KB/s and a round-trip latency of 20 ms on target tar1:

```
setrcopytarget tunelinks 2048 20 tar1
```

The following example shows how to have the system automatically tune RCIP links on target tar1:

```
setrcopytarget autotunelinks tar1
```

The following example will check the secure connection to Quorum Witness at IP address 2620:0:a04:1137:a2d3:c1ff:fef9:22d9:

```
setrcopytarget witness check -ssl 2620:0:a04:1137:a2d3:c1ff:fef9:22d9
```

The following example will create a secure quorum with Quorum Witness at IP address 2620:0:a04:1137:a2d3:c1ff:fef9:22d9 port 8443 on target s2308IP_sec:

```
setrcopytarget witness create -ssl -port 8443 2620:0:a04:1137:a2d3:c1ff:fef9:22d9 s2308IP_sec
```

setsched

Syntax

```
setsched [options] <schedname>
setsched [-suspend | -resume] <schedname>
setsched [-suspend_all | -resume_all]
```

Description

The `setsched` command allows users to suspend, pause, change the schedule, and change the name of currently scheduled tasks.

Options

`-suspend <schedname>`

Suspends execution of a specified task.

This will not halt execution of an already running task.

Cannot be used with any other options.

`-suspend_all`

Suspends execution of all scheduled tasks.

This will not halt execution of an already running task.

Cannot be used with any other options.

`-resume <schedname>`

Resumes scheduling of a specified task.

Cannot be used with any other options.

`-resume_all`

Resumes scheduling of all scheduled tasks.

Cannot be used with any other options.

`-s <newschedule>`

Sets a new schedule for a given task. The newschedule is a cron-style schedule for the task to be run on. Each field can be at most 127 characters in length.

-f

Does not detect when a modified scheduled task exceeds the recommended number of scheduled task starts per minute. Do not ask for confirmation to modify the scheduled task. This option only applies to the `setsched` -s option for setting a new schedule.

-name <newname>

Sets a new name for a given task. The newname is the new name to be given to the specified task. This can be at most 31 characters in length.

-no_alert

Failures of tasks will no longer generate an alert.

-alert

Failures of tasks will generate an alert. This is set by default.

Specifiers

<schedname>

The name of the scheduled task to modify.

Authority

- Super, Service, Edit
- Any role granted the `sched_setany` right can set any user scheduled task properties
- Any role granted the `sched_setown` right can only set their own scheduled task properties

Usage

- To avoid spikes in resource utilization, it is recommended to spread out scheduled task start times as much as possible. `Setsched -s` can predict out to 30 days when the modified scheduled task exceeds the recommended number of scheduled task starts per minute. When this is detected, user confirmation is requested to continue modifying the schedule.
- To skip this detection and confirmation, use the -f option.

Examples

None.

setsnmpmgr

Syntax

```
setsnmpmgr [options <arg>] <manager_IP>
```

Description

The `setsnmpmgr` command changes an SNMP manager's properties.

Options

-alertclear {all|nodup|standard}

all

Enable both kinds of alert clear trap:

1. The original alert which state has changed, but with a severity of "clear" if the new state is acknowledged, fixed, or autofixed.
2. The legacy alert state change `alertNotify` trap, with a `detailedMessage` of "Alert <id> changed from state <prev state> to <new state>"

If an alert is removed through `removealert`, only type 2 trap above is generated.

`nodup`

Enable only the type 1 trap above.

If an alert is removed through `removealert`, the type 2 trap above is generated instead.

`standard`

Enable only the type 2 trap above. This is the default.

`-notify {all|nodup|standard}`

`all`

Enables all trap notification types defined by the HPE 3PAR MIB.

`nodup`

Enables all trap notification types defined by the HPE 3PAR MIB. Only send an alertNotify trap if no other trap type will be sent for an event.

`standard`

Only send alertNotify type traps. This is the default.

`-p <port_number>`

Specifies the port number where the SNMP manager receives traps.

The port number and IP address are used together to uniquely identify the SNMP manager. Use this option if the port number differs from the default of "162".

`-pw <password>`

Specifies the SNMP manager's access community string (password), using up to 32 alphanumeric characters. If not specified, the default is "public".

`-snmpuser <username>`

Specifies the SNMPv3 user name using up to 31 characters. The user name must be enabled for SNMPv3 with the **`createsnmpuser`** command.

`-version <version>`

Integer value that specifies the SNMP version supported by the manager.

Use "2" for SNMPv2 and "3" for SNMPv3. The default is "2".

Specifiers

`<manager_IP>`

Specifies the IP address of the host where the manager runs. It must be valid IPv4 or IPv6 address. IPv6 address is in hexadecimal, is case insensitive, and is separated by colons. An example would be:

`5def:2008:0:0:abcd:0:0:161a`

In addition, a double colon (::) can be used once in an address to replace multiple fields of zeros. For example: `5def:2008:0:0:abcd::161a`

Authority

- Super
- Any role granted the `snmpmgr_set` right.

Usage

- Access to all domains is required to run this command.
- The storage system does not support any form of name resolution. Specify these IP addresses directly.
- Issue the **addsnmpmgr** command to add an SNMP manager.
- Issue the **showsnmpmgr** command to display the list of registered SNMP managers.
- Issue the **setsnmppw** command to change the SNMP passwords.
- Issue the **setsnmpuserpw** command to change the SNMPv3 user passwords.
- Issue the **removesnmppw** command to remove SNMP passwords.
- Issue the **removesnmpmgr** command to remove SNMP managers.

Examples

The following example modifies the SNMP manager IPv4 address 123.45.67.89 with the assigned password of alpha1 and standard trap notification:

```
cli% setsnmpmgr -pw alpha1 123.45.67.89 -notify standard
```

The following example modifies the user and version properties of the SNMP manager identified by the IPv4 address 123.45.67.89. The manager will receive all trap notification types for an event:

```
cli% setsnmpmgr -snmpuser john -version 3 -notify all 123.45.67.89
```

setsnmppw

Syntax

```
setsnmppw [options] <password>
```

Description

The `setsnmppw` command allows a user to update SNMP community string passwords. The SNMP password is required for the system manager to send requests to the SNMP agent.

Options

-rw|-r|-w

Specifies that the read-write (-rw), read-only (-r), or write-only (-w) community password is changed. If not specified, the read-write password is changed.

Specifiers

<password>

Specifies the new user-defined password using up to 50 alphanumeric characters.

Authority

- Super
- Any role granted the `snmppw_set` right

Usage

- Access to all domains is required to run this command.
- The default initial read-write password is public.
- If the read-only or write-only passwords do not exist, they are created.
- Verify SNMP passwords by issuing the **showsnmppw** command.

Examples

The following example changes the read-write SNMP community string password to newpassword1:

```
cli% setsnmppw newpassword1
```

The following example changes the read-only SNMP password to newpassword2 by specifying the -r option on the command line:

```
cli% setsnmppw -r newpassword2
```

The following example changes the write-only SNMP password to newpassword3 by specifying the -w option on the command line:

```
cli% setsnmppw -w newpassword3
```

setsnmpuserpw

Syntax

```
setsnmpuserpw [options] <username>
```

Description

The `setsnmpuserpw` command allows the user to update SNMP v3 user authentication and privacy passwords.

Options

`-authPassword`

Specifies the new authentication password for the SNMP user.

`-privPassword`

Specifies the new privacy password for the SNMP user.

Specifiers

`<username>`

Specifies the SNMP user name.

Authority

- Super
- Any role granted the `snmppw_set` right.

Usage

Access to all domains is required to run this command.

Examples

The following example changes authentication password for SNMP v3 user, user1:

```
cli% setsnmpuserpw -authPassword password1 user
```

The following example changes privacy password for SNMP v3 user, user1:

```
cli% setsnmpuserpw -privPassword password1 user1
```

The following example changes the authentication and privacy password for SNMP v3 user, user1:

```
 cli% setsnmpuserpw -authPassword password1 -privPassword password2 user1
```

setsralertcrit

Syntax

```
setsralertcrit {options} [-pat] {<name>...|<pattern>...|-all}
```

Description

Change the state and attributes of System Reporter threshold alert criteria.

Options

Options available to criteria of **any** types:

-daily

This criterion will be evaluated on a daily basis at midnight.

-hourly

This criterion will be evaluated on an hourly basis.

-hires

This criterion will be evaluated on a high resolution (5 minute) basis. This is the default.

-count <number>

-count_pct <percent>

The number or percent of matching objects that must meet the criteria in order for the alert to be generated. Note that only one alert is generated in this case and not one alert per affected object.

-recur <recurrences>/<samples>

The alert will only be generated if the other conditions of the criterion recur repeatedly. `<recurrences>` is an integer value from 2 to 10, and `<samples>` is an integer from 2 to 10 representing the number of previous System Reporter samples in which the recurrences will be examined. `<samples>` must be at least the requested quantity of recurrences. Note that these samples refer to the selected resolution of the criterion: hires, hourly, or daily.

-btsecs <relative_time>

A negative number indicating the number of seconds before the data sample time used to evaluate conditions which compare against an average. Instead of a number representing seconds, btsecs can be specified with a suffix of m, h, or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h), or days (e.g. -7d). The relative time cannot be more than 10 samples ago: 50 minutes for hires, 10 hours for hourly, or 10 days for daily. If this option is not present the average is only computed for the most recent data sample. The `–btsecs` option may not be combined with the `–recur` option.

-critical

This alert has the highest severity.

-major

This alert should require urgent action.

-minor

This alert should not require immediate action.

-info

This alert is informational only. This is the default.

-enable

Enables the specified criterion.

-disable

Disables the specified criterion.

-comment <comment>

Specifies comments or additional information for the criterion. The comment can be up to 511 characters long.

-defer <duration>

Once an alert is generated, the deferral period prevents the same alert from being repeated for a period of time. The deferral duration can be specified in seconds or with a suffix of m, h, or d to represent time in minutes (e.g. 30m), hours (e.g. 1.5h), or days (e.g. 7d). Note that a single alert criteria can generate multiple alerts if multiple objects exceed the defined threshold. A deferral period applies to each unique alert. Acknowledging an alert with "`setalert ack <id>`" will end its deferral period early.

-name <newname>

Specifies that the name of the SR alert be changed to `<newname>`, with a maximum of 31 characters.

-condition <field><comparison><value>[,...]

Specifies the condition(s) that causes the event to trigger, and is determined by the applicable fields listed below by event type.

Name based filtering options:

-pat

Specifies that certain patterns are treated as glob-style patterns and that all criteria matching the specified pattern will be modified. This option must be used if the pattern specifier is used. This option cannot be combined with -name, -condition, or any of the type-specific filtering options.

-all

Specifies that all criteria will have the designated operation applied to them, changing the state or attributes of all criteria. This option cannot be combined with -name, -condition, or any of the type-specific filtering options.

Type based filtering options:

Options specific to the **PORT** type:

- `-port_type <type>[,<type>...]`
- `-port <npat>:<spat>:<ppat>[,<npat>:<spat>:<ppat>...]`
- `-both | -ctl | -data`
- `-groupby <groupby>[,<groupby>...]`

  Consider the summed value of each combination of `<groupby>` items against the criterion's condition(s). The supported items are:

  PORT_N PORT_S PORT_P PORT_TYPE GBITPS TRANS_TYPE ALL

  If not specified, the default is to group by `PORT_N` and `PORT_S` and `PORT_P`.

Fields applicable to the **PORT** type:

`read_iops, write_iops, total_iops, read_kbps, write_kbps, total_kbps, read_svctms, write_svctms, total_svctms, read_ioszkb, write_ioszkb,`

```
total_ioszkb, qlen, sum_read_iops, sum_write_iops, sum_total_iops,
sum_read_kbps, sum_write_kbps, sum_total_kbps, avg_busy, read_bw_pct,
write_bw_pct, total_bw_pct, read_time0..read_time31,
write_time0..write_time31, total_time0..total_time31,
read_size0..read_size15, write_size0..write_size15, total_size0..total_size15
```

Options specific to the **VLUN** type:

- `-host <host_name|host_set|pattern>[,<host_name|host_set|pattern>...]`
- `-vv <VV_name|VV_set|pattern>[,<VV_name|VV_set|pattern>...]`
- `-l <LUN|pattern>[,<LUN|pattern>...]`
- `-port <npat>:<spat>:<ppat>[,<npat>:<spat>:<ppat>...]`
- `-groupby <groupby>[,<groupby>...]`

  Consider the summed value of each combination of <groupby> items against the criterion's condition(s). The supported items are:

  DOM_NAME VV_NAME HOST_NAME LUN HOST_WWN PORT_N PORT_S PORT_P VVSET_NAME HOSTSET_NAME VM_NAME VM_ID VM_HOST VVOLSC ALL

  If not specified, the default is to group by `HOST_NAME` and `VV_NAME` and `LUN`.

Fields applicable to the **VLUN** type:

```
read_iops, write_iops, total_iops, read_kbps, write_kbps, total_kbps,
read_svctms, write_svctms, total_svctms, read_ioszkb, write_ioszkb,
total_ioszkb, qlen, sum_read_iops, sum_write_iops, sum_total_iops,
sum_read_kbps, sum_write_kbps, sum_total_kbps, avg_busy,
read_time0..read_time31, write_time0..write_time31,
total_time0..total_time31, read_size0..read_size15,
write_size0..write_size15, total_size0..total_size15
```

Options specific to the **LD** type:

- `-node <node>[,<node>...]`
- `-cpg <CPG_name|pattern>[,<CPG_name|pattern>...]`
- `-ld <LD_name>|<pattern>[,<LD_name>|<pattern>]`
- `-groupby <groupby>[,<groupby>...]`

  Consider the summed value of each combination of `<groupby>` items against the criterion's condition(s). The supported items are:

  DOM_NAME LDID LD_NAME CPG_NAME NODE ALL

  If not specified, the default is to group by `LDID`.

Fields applicable to the **LD** type:

```
read_iops, write_iops, total_iops, read_kbps, write_kbps, total_kbps,
read_svctms, write_svctms, total_svctms, read_ioszkb, write_ioszkb,
total_ioszkb, qlen, sum_read_iops, sum_write_iops, sum_total_iops,
sum_read_kbps, sum_write_kbps, sum_total_kbps, avg_busy,
read_time0..read_time31, write_time0..write_time31,
total_time0..total_time31, read_size0..read_size15,
write_size0..write_size15, total_size0..total_size15
```

Options specific to the **PD** type:

- `-disk_type <type>[,<type>...]`
- `-rpm <speed>[,<speed>...]`
- `-pd <PDID>|<pattern>[,<PDID>|<pattern>]`
- `-groupby <groupby>[,<groupby>...]`

  Consider the summed value of each combination of `<groupby>` items against the criterion's condition(s). The supported items are:

  PDID PORT_N PORT_S PORT_P DISK_TYPE SPEED ALL

  If not specified, the default is to group by `PDID`.

Fields applicable to the **PD** type:

`read_iops, write_iops, total_iops, read_kbps, write_kbps, total_kbps, read_svctms, write_svctms, total_svctms, read_ioszkb, write_ioszkb, total_ioszkb, qlen, sum_read_iops, sum_write_iops, sum_total_iops, sum_read_kbps, sum_write_kbps, sum_total_kbps, avg_busy, read_time0..read_time31, write_time0..write_time31, total_time0..total_time31, read_size0..read_size15, write_size0..write_size15, total_size0..total_size15`

Options specific to the **CMP** type:

`-node <node>[,<node>...]`

Fields applicable to the **CMP** type:

`read_hits, write_hits, reads, writes, totals, lock_blocks, read_hitpct, write_hitpct, delack_nl, delack_fc, delack_ssd, free_page, clean_page, write1_page, writen_page, writesched_page, writing_page, recov_page, dirty_nl, dirty_fc, dirty_ssd, maxdirty_nl, maxdirty_fc, maxdirty_ssd,`

The following fields have been deprecated:

`delack_fc10, delack_fc15, delack_ssd150, delack_ssd100, dirty_fc10, dirty_fc15, dirty_ssd150, dirty_ssd100, maxdirty_fc10, maxdirty_fc15, maxdirty_ssd150, maxdirty_ssd100`

Each acts as an alias for the equivalent field without distinction between disk speeds. For example `delack_ssd100` or `delack_ssd150` report equal to `delack_ssd`.

Options specific to the **CPU** type:

- `-node <node>[,<node>...]`
- `-groupby <groupby>[,<groupby>...]`

  Consider the summed value of each combination of `<groupby>` items against the criterion's condition(s). The supported items are:

  NODE CPU ALL

  If not specified, the default is to group by `NODE` and `CPU`.

Fields applicable to the **CPU** type:

`user_pct, sys_pct, idle_pct, interrupts, context_switches`

Options specific to the **LINK** type:

- `-node <node>[,<node>...]`
- `-groupby <groupby>[,<groupby>...]`

  Consider the summed value of each combination of `<groupby>` items against the criterion's condition(s). The supported items are:

  `NODE QUEUE NODE_TO ASIC_FROM ASIC_TO ALL`

  If not specified, the default is to group by `NODE`.

Fields applicable to the **LINK** type:

`xcb_sent, kbps, xcb_size`

Options specific to the **QOS** type:

- `-target {{vvset|domain}:{<name>|<pattern>}|sys:all_others}[,...]`
- `-groupby <groupby>[,<groupby>...]`

  Consider the summed value of each combination of `<groupby>` items against the criterion's condition(s). The supported items are:

  `DOM_NAME QOS_ID TARGET_TYPE TARGET_NAME ALL`

  If not specified, the default is to group by `QOS_ID` and `TARGET_NAME`.

Fields applicable to the **QOS** type:

`read_iops, write_iops, total_iops, read_kbps, write_kbps, total_kbps, read_svctms, write_svctms, total_svctms, read_ioszkb, write_ioszkb, total_ioszkb, qlen, sum_read_iops, sum_write_iops, sum_total_iops, sum_read_kbps, sum_write_kbps, sum_total_kbps, limit_iops, limit_kbps, read_wait, write_wait, total_wait, rej, wait_qlen`

Options specific to the **RCOPY** type:

- `-target <TARGET_NAME|pattern>[,<TARGET_NAME|pattern>...]`
- `-groupby <groupby>[,<groupby>...]`

  Consider the summed value of each combination of `<groupby>` items against the criterion's condition(s). The supported items are:

  `TARGET_NAME LINK_ID LINK_ADDR PORT_N PORT_S PORT_P ALL`

  If not specified, the default is to group by `TARGET_NAME`.

Fields applicable to the **RCOPY** type:

`rc_kbs, rc_kbps, rc_hbrttms`

Options specific to the **RCVV** type:

- `-vv <VV_name|VV_set|pattern>[,<VV_name|VV_set|pattern>...]`
- `-groupby <groupby>[,<groupby>...]`

  Consider the summed value of each combination of `<groupby>` items against the criterion's condition(s). The supported items are:

VV_NAME DOM_NAME TARGET_NAME TARGET_MODE GROUP_NAME GROUP_ROLE PORT_TYPE PORT_N PORT_S PORT_P VVSET_NAME ALL

If not specified, the default is to group by VV_NAME.

Fields applicable to the **RCVV** type:

lcl_read_iops, lcl_write_iops, lcl_total_iops, rmt_read_iops, rmt_write_iops, rmt_total_iops, lcl_read_kbps, lcl_write_kbps, lcl_total_kbps, rmt_read_kbps, rmt_write_kbps, rmt_total_kbps, lcl_read_svctms, lcl_write_svctms, lcl_total_svctms, rmt_read_svctms, rmt_write_svctms, rmt_total_svctms, lcl_read_ioszkb, lcl_write_ioszkb, lcl_total_ioszkb, rmt_read_ioszkb, rmt_write_ioszkb, rmt_total_ioszkb, lcl_total_qlen, lcl_busy_pct, rmt_total_qlen, rmt_busy_pct

Options specific to the **CACHE** type:

-node <node>[,<node>...]

Fields applicable to the **CACHE** type:

reads_per_sec, writes_per_sec, cmp_read_hit_pct, cmp_write_hit_pct, fmp_read_hit_pct, fmp_write_hit_pct, read_back_iops, read_back_mbps, destage_write_iops, destage_write_mbps, fmp_used_pct, fmp_q_dormant, fmp_q_cold, fmp_q_norm, fmp_q_warm, fmp_q_hot, fmp_q_destage, fmp_q_read, fmp_q_flush, fmp_q_wrtback, cmp_q_free, cmp_q_clean, cmp_q_write1, cmp_q_writen, cmp_q_wrtsched, cmp_q_writing, cmp_q_dcowpend, cmp_q_dcowproc

Options specific to the **LDSPACE** type:

- -disk_type <type>[,<type>...]
- -owner <node>[,<node>...]
- -raid_type <type>[,<type>...]
- -cpg <CPG_name|pattern>[,<CPG_name|pattern>...]
- -ld <LD_name>|<pattern>[,<LD_name>|<pattern>]
- -groupby <groupby>[,<groupby>...]

  Consider the summed value of each combination of <groupby> items against the criterion's condition(s). The supported items are:

  DOM_NAME CPG_NAME LDID LD_NAME DISK_TYPE RAID_TYPE SET_SIZE STEP_SIZE ROW_SIZE OWNER ALL

  If not specified, the default is to group by LDID.

Fields applicable to the **LDSPACE** type:

raw_mb, used_mb, free_mb, total_mb

Options specific to the **PDSPACE** type:

- -disk_type <type>[,<type>...]
- -rpm <speed>[,<speed>...]

- `-pd <PDID>|<pattern>[,<PDID>|<pattern>]`
- `-groupby <groupby>[,<groupby>...]`

  Consider the summed value of each combination of `<groupby>` items against the criterion's condition(s). The supported items are:

  PDID CAGEID CAGESIDE MAG DISK DISK_TYPE SPEED ALL

  If not specified, the default is to group by `PDID`.

Fields applicable to the **PDSPACE** type:

`nrm_used_ok, nrm_used_fail, nrm_avail_clean, nrm_avail_dirty, nrm_avail_fail, spr_used_ok, spr_used_fail, spr_avail_clean, spr_avail_dirty, spr_avail_fail, lifeleft_pct, t_degc`

Options specific to the **CPGSPACE** type:

- `-disk_type <type>[,<type>...]`
- `-raid_type <type>[,<type>...]`
- `-cpg <CPG_name|pattern>[,<CPG_name|pattern>...]`
- `-groupby <groupby>[,<groupby>...]`

  Consider the summed value of each combination of `<groupby>` items against the criterion's condition(s). The supported items are:

  DOM_NAME CPGID CPG_NAME DISK_TYPE RAID_TYPE ALL

  If not specified, the default is to group by `CPG_NAME`.

Fields applicable to the **CPGSPACE** type:

`adm_used_mb, snp_used_mb, usr_used_mb, total_used_mb, adm_free_mb, snp_free_mb, usr_free_mb, total_free_mb, adm_mb, snp_mb, usr_mb, total_mb, growth_avail_mb`

Options specific to the **VVSPACE** type:

- `-usr_cpg <CPG_name|pattern>[,<CPG_name|pattern>...]`
- `-snp_cpg <CPG_name|pattern[,<CPG_name|pattern>...]`
- `-prov <prov_type>[,<prov_type>...]`
- `-vv <VV_name|VV_set|pattern>[,<VV_name|VV_set|pattern>...]`
- `-compr <Yes|No|Off|NA>`
- `-groupby <groupby>[,<groupby>...]`

  Consider the summed value of each combination of `<groupby>` items against the criterion's condition(s). The supported items are:

  DOM_NAME VVID VV_NAME BSID WWN SNP_CPG_NAME USR_CPG_NAME PROV_TYPE VV_TYPE VVSET_NAME VM_NAME VM_ID VM_HOST VVOLSC VVOL_STATE COMPR ALL

  If not specified, the default is to group by `VV_NAME`.

Fields applicable to the **VVSPACE** type:

userrawrsvd_mb, snaprawrsvd_mb, adminrawrsvd_mb, totalrawrsvd_mb, userused_mb, userfree_mb, userrsvd_mb, snapused_mb, snapfree_mb, snaprsvd_mb, adminused_mb, adminfree_mb, adminrsvd_mb, totalused_mb, totalrsvd_mb, virtualsize_mb

Options specific to the **SYSSPACE** type:

- -disk_type <type>[,<type>...]
- -groupby <groupby>[,<groupby>...]

  Consider the summed value of each combination of <groupby> items against the criterion's condition(s). The supported items are:

  DOM_NAME DISK_TYPE ALL

  If not specified, the default is to group by ALL, considering the system as a whole.

Fields applicable to the **SYSSPACE** type:

total_mb, alloc_mb, used_mb, system_mb, free_mb, failed_mb, compact_ratio, dedup_ratio, overprov_ratio

Specifiers

<name>

Specifies the name of the criterion to modify. Multiple names can be specified separated by spaces. May not be combined with -all.

<pattern>

Specifies alert criteria by matching names to the glob-style pattern. (For more information, enter clihelp sub,globpat.) Requires the -pat option. Multiple patterns can be specified separated by spaces. May not be combined with the -all option.

Authority

- Super, Service
- Any role granted the sralertcrit_set right.

Usage

- Disabling or changing a criterion also automatically fixes all alerts associated with that criterion.
- All options can be specified in any order.
- Note that the criterion's TYPE cannot be changed.
- Each option replaces the criterion's existing setting. For example -hourly option can change the frequency of a -hires criteria and the -condition option will replace the existing condition(s).
- Options that require a value can be cleared with a value of "". The -name and -condition options cannot be cleared.
- <condition>
- The condition must be of the format <field><comparison><value> where field is one of the fields corresponding to the type (see above), comparison is of the format <,<=,>,>=,=,!= and value is a numeric value, or is a numeric value followed by %_average to indicate that the field is to be compared against the average across multiple objects as specified by filtering options and/or across multiple data sample times as specified by the -btsecs option. See examples. Note that some characters, such as < and >, are significant in most shells and must be escaped or quoted when running this command from another shell. Multiple conditions may be separated by comma (",") to

indicate a logical AND requirement (conjunction). Conditions may be separated by the character "~" to indicate a logical OR requirement (disjunction). AND logic takes precedence over OR logic, and parentheses are not supported to override the natural precedence of the condition terms and logical operators.

Examples

The following example disables the criterion called port_writes:

```
cli% setsralertcrit -disable port_writes
```

The following example removes the RPM filter from a PDSPACE criterion:

```
cli% setsralertcrit -rpm "" crit_name
```

The following example changes the threshold conditions of an existing criterion:

```
cli% setsralertcrit -condition read_iops>500~write_iops>800 pd_crit
```

The following example shows how to change the name of an existing criterion:

```
cli% setsralertcrit -name new_name crit_name
```

The following example changes the severity of a criterion to critical:

```
cli% setsralertcrit -critical crit_name
```

setsshkey

Syntax

```
setsshkey [options]
```

Description

The `setsshkey` command sets the SSH public key for a user. The user will be prompted to provide an SSH public key. To finish entering the public key, press enter on a blank line. The key must have been generated using the ssh-keygen utility. The public key is contained in the user-defined file named with .pub extension. The user can open this file with an ASCII editor to copy the key and paste it. After setting an SSH public key on the storage system, the user can use the corresponding private key to log on without a password. If the -add option is not given, this new key replaces the existing keys if any.

Options

`-add`

Specifies that the given key should be added to the list of authorized keys instead of replacing existing ones.

Specifiers

None.

Authority

- Super, Service, Edit, Browse
- Any role granted the `sshkey_set` right

Usage

- Issuing the `setsshkey` command results in the system prompting for your SSH public key. Copy and paste the key using the aforementioned ASCII editor on the command line and then press ENTER.
- After setting an SSH public key on the system, use the corresponding private key to log on without a password. A single ssh key may be used for multiple hosts with a single use of `setsshkey`, or different keys may be used for

different hosts. To restrict the usage of a key to a particular host, a "from" option can be added to the beginning of a key. For example:

- `from="192.168.1.1" ssh-rsa AAAAB3Nza...`
- The maximum number of characters used to represent the SSH key (including the "from" option, key type, and additional comments) is 4095.
- LDAP users are only allowed to set an SSH key if the **setauthparam** command has been used to set the allow-ssh-key parameter to 1. When an LDAP user runs the `setsshkey` command, the user's role is recorded and is assigned when the user logs in using the key. Changes in the group-to-role mappings set with the **setauthparam** command or changes in the user's data in the LDAP server have no effect as long as the user has an SSH key.
- Removing the user's SSH key forces a new role to be determined at the user's next login.
- Only one key may be entered at a time; to enter multiple keys, run `setsshkey` again with the -add option.

Examples

Setting a user's SSH public key just type "setsshkey" then follow the instructions at the prompt:

```
cli% setsshkey
Please enter the SSH public key below. When finished, press enter twice.
The key is usually long. It's better to copy it from inside an editor
and paste it here. (Please make sure there are no extra blanks.)
The maximum number of characters used to represent the SSH key
(including the "from" option, key type, and additional comments) is 4095.

ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAQEAp+4Z3zT8Sq6t3s08q/MMd2ZnMpRc/3tyHMk63dH
R8b/VgV6ewXNfYhhfTGWRdZ1dtyLFXnuEPuf+z8EtinFStinTzA3FiJ0agK7rLoNtn/F0jBaGWm
SWukqzAQA2VJvq/keaLVMT3+J3nvXEUcjS4fApeeLwFgKczOX511oaHDtn2ys2C5l+mLw1VDUJL
wIKJljOWqr68ToeRgTDmoppOI3cG14ryF0re4xKANHFQLnSCt5ANjqD2jpnyEABuOvBX7G6vI5g
zQFfcAda/a2bq563/AHr8Ehhi5EVy+GSCqjj8cr0/zHSZyX6llQUfe0YxmayPEKlods6bDi5oxG
COw== user1@server1
```

setstatch

Syntax

```
setstatch start|stop <LD_name> <chunklet_num>
```

Description

The `setstatch` command starts and stops the statistics collection mode for chunklets.

Options

None.

Specifiers

<LD_name>

Specifies the name of the logical disk in which the chunklet to be configured resides.

<chunklet_num>

Specifies the chunklet that is configured using the `setstatch` command.

Subcommands

start|stop

Specifies that the collection of statistics is either started or stopped for the specified logical disk and chunklet.

Authority

- Super, Edit
- Any role granted the `statch_set` right

Usage

After the statistic collection mode for the chunklet is set, you can then use either the **histch** command or the **statch** command to view the chunklet's statistics.

Examples

The following example displays the start of statistics collection on chunklet 0 of logical disk test:

```
cli% setstatch start test 0
```

setstatpdch

Syntax

```
setstatpdch start|stop <PD_ID>
```

Description

The `setstatpdch` command sets the statistics collection mode for all in-use chunklets on a physical disk.

Options

None.

Specifiers

<PD_ID>

Specifies the physical disk ID.

Subcommands

start|stop

Specifies that the collection of statistics is either started or stopped for chunklets on the specified physical disk used by logical disks.

Authority

- Super, Edit
- Any role granted the `statpdch_set` right

Usage

After the statistic collection mode for the chunklet is set, you can then use either the **histch** command or the **statch** command to view the chunklet's statistics.

Examples

The following example displays the start of statistics collection on all physical disk chunklets of physical disk 0:

```
cli% setstatpdch start 0
```

setsys

Syntax

```
setsys [options]
```

```
setsys [-f] <parameter>
```

Description

Sets properties of the system. Properties include options to annotate a system with descriptor information such as physical location, owner, contact information, and so on. The command also allows you to set system-wide parameters such as the raw space alert.

Options

-name *<systemname>*

Changes the system name by specifying a new name, up to 31 characters in length.

The following options allow annotations with descriptor information:

-loc *<location>*

Specifies the location of the system.

-owner *<owner>*

Specifies the owner of the system.

-contact *<contact>*

Specifies the contact information for the system.

-comment *<comment>*

Specifies any additional information for the system.

-f

Forces the command. The command completes the process without prompting for confirmation.

Specifiers

Configures the following specifiers on the system and for the *<parameter>* specifier:

RawSpaceAlertFC *<value>*

Sets the user configurable space alert threshold (10 to 100000 GiB) for Fast Class type drives (defaults to 0). When the total space on the available chunklets (both clean and unclean) for the specified drive type falls below the specified value, the system can post one of the four system standard alert levels:

- 50%
- 75%
- 85%
- 95%

The system converts the supplied space alert threshold from GiB free to percentage used. If any of the system standard alert levels are below this computed value, they are suppressed.

A value of 0 cancels the user-specified alert threshold and reverts back to the system standard alert levels. To minimize the alerts, use the specified value of 10 (GiB).

RawSpaceAlertNL *<value>*

Performs the same function as `RawSpaceAlertFC`, but for NearLine type drives.

RawSpaceAlertSSD *<value>*

Performs the same function as `RawSpaceAlertFC`, but for Solid State Drive type drives.

RemoteSyslog *<value>*

Enables or disables sending events as syslog messages to both remote syslog servers. A value of 0 disables the logging, and a value of 1 enables logging. Sends the syslog messages with a facility user and with event severities mapped to syslog priority levels as follows:

Event Severity	syslog Level
fatal	alert
critical	alert
major	crit
minor	err
degraded	warning
info	info

RemoteSyslogHost {{<hostname>|<IPv4>}[:<port>]|<IPv6>|[<IPv6>]:<port>}, {{<hostname>|<IPv4>}[:<port>]|<IPv6>|[<IPv6>]:<port>}, {{<hostname>|<IPv4>}[:<port>]|<IPv6>|[<IPv6>]:<port>}

Sets the hostname or IP address of the remote syslog server to which general events are sent as syslog messages. Optionally sets port. When configuring the destination with both the IPv6 address and port, enclose the IPv6 address in square brackets. You can specify up to three servers by separating the information with commas.

With no port configured, the system uses one of the following default: 514 for UDP, 601 for TCP, 6514 for TLS. To view the status of the syslog servers, enter the `showsys -d` command.

RemoteSyslogSecurityHost {{<hostname>|<IPv4>}[:<port>]|<IPv6>|[<IPv6>]:<port>}, {{<hostname>|<IPv4>}[:<port>]|<IPv6>|[<IPv6>]:<port>}, {{<hostname>|<IPv4>}[:<port>]|<IPv6>|[<IPv6>]:<port>}

Sets the hostname or IP address of the remote syslog server to which security events are sent as syslog messages. Optionally sets port. When configuring the destination with both the IPv6 address and port, enclose the IPv6 address in square brackets. You can specify up to three servers by separating the information with commas.

The system sends security messages over TLS only. With no port configured, the system uses port 6514. To view the status of the syslog servers, enter the `showsys -d` command.

SparingAlgorithm <value>

Sets the sparing algorithm used by the **admithw** command. Valid values are Default (roughly 2.5% with minimums), Minimal (roughly 2.5% without minimums), Maximal (one disk's worth in every cage), and Custom (not managed automatically by the system).

VVRetentionTimeMax <time>{h|H|d|D}

Specifies the maximum value that can be set for the retention time of a volume. `<time>` is a positive integer value and in the range of 0 - 43,800 hours (1825 days). Time can be specified in days or hours providing either the "d" or "D" for day and "h" or "H" for hours following the entered time value. The default value for `<time>` is 14 days.

To disable setting the volume retention time in the system, enter 0 for <time>.

UpgradeNote <value>

Sets a note that will be displayed when `checkupgrade` is run. To include spaces, enclose the value inside double quotes. To erase an existing note, specify "". Maximum length is 127 characters.

PortFailoverEnabled {yes|no}

Enables or disables the automatic failover of target ports to their designated partner ports. If enabled, automatic failover is triggered when a node goes down. The default is "yes".

AutoExportAfterReboot {yes|no}

Enables or disables automatically exporting vluns after a reboot. If disabled, vluns and host ports will not become active after a reboot until "`setsysmgr export_vluns`" is issued. The default is "yes".

AllowWrtbackSingleNode <value>

Allows the system to continue caching when in a single node state for up to the specified number of days. Once past that, all LDs will go into write-through. This does not include single node states that occur as a result of an upgrade. <value> can be 0 to 7 days. The default is 7 which puts LDs into caching when in single node state for up to 7 days.

AllowWrtbackUpgrade <value>

Allows the system to continue caching when in a single node state during an upgrade for up to the specified number of days. This does not include single node states that occur during a non-upgrade situation. <value> can be 0 to 7 days. The default is 7 which puts all LDs into caching when in single node state during an upgrade for up to 7 days.

EventLogNum <value>

Sets the number of retained event logs in addition to the active file. The range is between 1 and 30, with the default value being 10. When the number of event log files reaches this count, the oldest file is removed. Setting this value to a lower count will remove any event log files that no longer meet this constraint. Increasing this value will increase space consumption to accommodate the additional files.

EventLogSize <value>

Sets the size of the event log. The range is between 0.5 MiB and 10 MiB with the default value being 4 MiB. This sets the size of the individual files; by default, 11 total files are retained (the active event log plus 10 archived versions). When the current most recent file hits this size, the files are rolled, with the oldest file being removed if necessary to maintain no more than the specified maximum retention count of files. Increasing this value will increase space consumption to accommodate the larger file size.

SessionTimeout <time>[h|m]

Specifies the value that can be set for the idle session timeout for a CLI session. <time> is a positive integer value and in the range of 180 seconds (3 minutes) - 86399 seconds (23:59:59 hours). Time can be specified in hours, minutes or seconds providing "h" or "m" and just the value for the number of seconds. Note that, when entering the value in seconds, "s" should not be used.

The default value for <time> is 60 minutes (1 hour).

HostDIF {yes|no}

Enables or disables HPE 3PAR implementation of host based Data Integrity Field (DIF) support for all ports (FC only). This will only affect VVs which have their host DIF policy set to "3par_host_dif". The default value is "yes" on systems capable of DIF.

HostDIFTemplate <value>

Sets the default host DIF settings to be used by the `createvv` command. <value> can be "no_host_dif" (no support for host DIF), "3par_host_dif" (HPE 3PAR implementation of DIF) or "std_host_dif" (standard SCSI implementation of DIF). The default value is "3par_host_dif" on systems capable of DIF. Note that the host DIF support is available on FC ports only.

ComplianceOfficerApproval {yes|no}

Specifies whether to enable or disable the compliance officer approval mode.

OverprovRatioWarning <value>

Specifies an overprovisioning ratio, which when exceeded by the system, a device type, or a CPG, results in a warning alert. A ratio of 0 (default) means no warning alert is generated. A ratio of 3 means that there is 3 times the size virtually available than what is physically available. Note that this will be different across each CPG and for the system as a whole.

DisableCompr <value>

New write requests to the Compressed VVs serviced by the system will not be compressed if this parameter is set to "yes". Setting this parameter to "no" will enable writes to the Compressed VVs to be compressed. <value> can be "yes" or "no". The default is "no".

OverprovRatioLimit <value>

Specifies that the system, device types, and all CPGs are limited to the specified overprovisioning ratio. A ratio of 0 (default) means no limit is enforced. A ratio of 3.0 means that there is 3 times the size virtually available than what is physically available. Note that this will be different across each CPG and for the system as a whole. Setting this limit <1.0 will prevent some of the available space from being used. Overprovisioning is an estimate and setting the value to 1.0 isn't a guarantee that a system will never run out of physical space.

RemoteCopyHostThrottling {yes|no}

Enables or disables Remote Copy throttling policy for host IO replicated in asynchronous streaming mode. As Remote Copy resources approach the maximum limits on systems where RemoteCopyHostThrottling is set to "yes" then host IO that is replicated to asynchronous streaming groups are subjected to active host IO management to reduce resource consumption and prevent some or all groups from becoming suspended.

When the policy is set to "no" then host IO replicated to asynchronous streaming groups will not be subjected to any active host IO management which may result in some or all Remote Copy groups becoming suspended sooner.

When the Remote Copy resources reach maximum limits some or all Remote Copy groups may be suspended until system resources become available once more.

Defaults to `no`, for no throttling.

AutoAdmitTune {yes|no}

Enables or disables automatic rebalancing when admithw detects new disks. This setting only applies to 2-node systems where the default is "yes".

SingleLunHost <value>

Enables or disables support to limit VV exports such that each VV can only be exported to a given host one time. `<value>` can be "yes" or "no". Defaults to `no`.

R6LayoutVersion <value|Default>

Specifies which RAID6 implementation to use. This does not affect existing LDs. A <value> of "0" chooses the original RAID6 implementation. A <value> of "1" enables the use of optimal layouts for set size 6, 10, and 16. A <value> of "2" enables the optimal layouts from version "1" and optimal layouts for set size 8 and 12. The default value is "Default", which is currently "2".

AllowDomainUsersAffectNoDomain {yes|no|hostonly}

By default, users with permissions only for specific domains and not for all domains are not allowed to use commands that involve the use of resources that do not belong to specific domains. This includes the createhost and createcpg commands, among others. Setting this value to "no" keeps this behavior. Setting this value to "hostonly" allows these users to create and remove hosts in domains in which they have edit level privilege using WWNs and iSCSI names that have not yet been allocated to a host. Setting this value to "yes" allows these and most other commands that normally require edit level permissions over all domains, such as creating new CPGs. They will not be able to directly view or

modify objects in domains they do not have permissions in, but will have greater ability than normal to indirectly affect other domains. The default value is "no".

AllowSSZ (yes|no)

Enables or disables support for using the `-ssz` option during CPG creation. If disabled, you cannot use the `-ssz` option during CPG creation.

AllowR5OnNLDrives {yes|no}

Enables or disables support for RAID-5 on NL drives. If disabled, new RAID-5 CPGs cannot be created on NL drives. HPE strongly recommends RAID-6 for high capacity NL drives. Defaults to `no`.

AllowR5OnFCDrives {yes|no}

Enables or disables support for RAID-5 on FC drives. If disabled, new RAID-5 CPGs cannot be created on FC drives. HPE strongly recommends RAID-6 for FC drives. Defaults to `no`.

AllowR0 {yes|no}

Enables or disables support for RAID-0. If disabled, new RAID-0 CPGs cannot be created on this system. RAID-0 virtual volumes do not offer data protection. Defaults to `no`.

DisableDedup {yes|no}

New write requests to TDVVs serviced by the system will not be deduplicated if this parameter is set to "yes". Setting this parameter to "no" will enable writes to TDVVs to be deduplicated. Defaults to `no`.

DisableCompr <value>

New write requests to the Compressed VVs serviced by the system will not be compressed if this parameter is set to "yes". Setting this parameter to "no" will enable writes to the Compressed VVs to be compressed. `<value>` can be "yes" or "no". Defaults to `no`.

Authority

- Super, Service
- Any role granted the `sys_set` right

Usage

- Access to all domains is required to run this command.
- Use the `showsys -param` command to see the current settings of the above values (see the **showsys** command).
- Cluster-wide compliance setting (`setsys ComplianceOfficerApproval`) cannot be disabled if there is any compliance enabled VFS or File store in the system. To disable, remove all compliance enabled VFSs and File Stores.

Examples

The following example sets a raw space alert of 800 gigabytes:

```
cli% setsys RawSpaceAlertFC 800
cli% showsys -param
System parameters from configured settings
----Parameter-----            ---Value----
RawSpaceAlertFC            :            800
RawSpaceAlertNL            :              0
RemoteSyslog               :              1
RemoteSyslogHost           :   192.168.6.15
SparingAlgorithm           :        Minimal
CopySpaceReclaim           :              0
```

```
EventLogSize                 :              4M
VVRetentionTimeMax           :      336 Hours
```

The following example sets EventLogSize to 2MB:

```
cli% setsys EventLogSize 2
```

The following examples set the timeout in hours, minutes and seconds:

The value in seconds requires no `s` suffix.

```
cli% setsys SessionTimeout 4h
cli% setsys SessionTimeout 10m
cli% setsys SessionTimeout 240
```

The following examples set RemoteSyslogHost with a port:

```
cli% setsys RemoteSyslogHost https://example.com:8080
cli% setsys RemoteSyslogHost 127.0.0.1:8080
cli% setsys RemoteSyslogHost [2620:0:a04:1120:15:252:204:235]:8080
cli% setsys RemoteSyslogHost 127.0.0.1:8080,127.0.0.2,https://example.com:8080
```

setsysmgr

Syntax

The syntax of the `setsysmgr` command can be one of the following:

```
setsysmgr [option] wipe <system_name>

setsysmgr [option] tocgen [<toc_gen_number>]

setsysmgr [option] force_iderecovery

setsysmgr [option] force_idewipe

setsysmgr [option] export_vluns
```

Description

The `setsysmgr` command sets the system manager startup state.

Options

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

> ⚠ **WARNING:** This option will result in the loss of data and configuration info.

Specifiers

`<system_name>`

Specifies the name of the system to be started, using up to 31 characters.

`<toc_gen_number>`

Specifies the table of contents generation number for the system to start with.

Subcommands

`wipe`

Requests that the specified system be started in the new system state.

CAUTION: If this option is used, critical warning messages may be missed!

tocgen

Specifies that the system is to be started with the specified table of contents generation number.

force_iderecovery

Specifies that the system starts the recovery process from the IDE disk even if all virtual volumes have not been started.

force_idewipe

Specifies that the system wipes the IDE power fail partition. The system is shutdown and restarted, during which time all logical disks and virtual volumes are checked.

export_vluns

If the AutoExportAfterReboot option has been set to no, after a power failure or uncontrolled shutdown vluns will not be automatically exported, and host ports will be in a suspended state. This command will re-export the luns and enable the host ports after this happens.

Authority

- Super
- Any role granted the `sysmgr_set` right

Usage

- Access to all domains is required to run this command.
- Issuing the `setsysmgr` command can potentially remove the entire state of the system causing data loss.
- Use this command only when the system cannot start up normally.
- If the wipe subcommand is specified, all system data and configuration information including customer data and virtual volume layout are destroyed.
- If the force_iderecovery subcommand is specified, the system can delete data for some of the unstarted virtual volumes. The system can run the **checkvv** and **checkld** commands to make the virtual volumes and logical disks consistent, thereby resulting in a possible data loss.

Examples

The following example starts system "mysystem" in the new system state:

```
cli% setsysmgr wipe mysystem
```

The following example starts a system's powerfail recovery despite being unable to start one or more volumes:

```
cli% setsysmgr force_iderecovery
```

settask

Syntax

```
settask [-f] -pri <high|med|low|auto> <task_ID>
```

Description

The `settask` command sets the priority on specified task.

Options

-f

Forces the command. The command completes the process without prompting for confirmation.

-pri <high|med|low|auto>

Specifies the priority of the task.

Specifiers

<task_ID>

Sets the priority on a task specified by the task ID.

The <task_ID> must be an unsigned integer from 1 to 29999.

Authority

- Super, Edit
- Any role granted the task_set right.

Usage

Task priorities can only be set one at a time. If the specified task is not active or valid, attempting to set its priority will result in an error.

Examples

The following example shows how to set a priority on a task using the task ID:

```
cli% settask -pri high 497
Are you sure you want to set priority on task 497?
select q=quit y=yes n=no: y
high priority is set on task id: 497
```

settemplate

Syntax

```
settemplate <option_value>... [option <arg>] <template_name>
```

Description

The settemplate command modifies the properties of existing templates.

Options

-remove <option>...

Indicates that the option(s) that follow -remove are removed from the existing template. When specifying an option for removal, do not specify the option's value. For valid options, see the **createtemplate** command.

Specifiers

<option_value>...

Indicates the specified options and their values (if any) are added to an existing template. The specified option replaces the existing option in the template. For valid options, see the **createtemplate** command.

<template_name>

Specifies the name of the template to be modified, using up to 31 characters.

Authority

- Super
- Any role granted the `template_set` right

Usage

- Access to all domains is required to run this command.
- The -desc option cannot be removed from a template (it can be changed to an empty string).
- The -nrw and -nro options cannot be removed from a template. These options can only be replaced by specifying either -nrw or -nro before the -remove option.
- All options available for the **createcpg** command can be used with the `settemplate` command to modify an existing template.
- Options preceded with -remove are removed from an existing template.
- When adding options to a template, specify the option flag and its value (if any).
- When removing options from a template, specify the option flag only.

Examples

In the following example, template vvtemp1 is modified to support the availability of data should a drive magazine fail (mag) and to use the the stale_ss policy:

```
cli% settemplate -ha mag -pol stale_ss vvtemp1
```

In the following example, the -nrw and -ha mag options are added to the template template1, and the -t option is removed:

```
cli% settemplate -nrw -ha mag -remove -t template1
```

setuser

Syntax

```
setuser [options] <user>
```

Description

The `setuser` command sets the user properties.

Options

-f

Specifies that the command is forced. No confirmation is requested before executing the command.

-adddomain <domain>:<role>[,<domain>:<role>...]

Adds a specified user (`<user>`) to the specified domain (`<domain>`) at the specified role (`<role>`). Super, audit, and service roles are **not** allowed.

-rmdomain <domain_name_or_pattern>[,<domain_name_or_pattern>...]

Removes the user from each domain with a name that matches one or more of the `<domain_name_or_pattern>` options.

-defaultdomain <domain>

Changes the default domain of the user to a specified domain. You must already have permission set in the domain. Specify "-unset" as the domain name to unset the default domain. The Browse and Edit authorities for this command and option can be used by any user with their own user name.

Specifiers

<user>

Specifies the name of the user.

Authority

- Super, Edit, Browse
- Any role granted the user_setany right can add, remove, and set default domains for any user.
- Any role granted the user_setown right can only set their own default domain.

Usage

- The -adddomain and -rmdomain options cannot be used on a user that is currently logged in. Existing sessions can be terminated with the **removeuserconn** command.
- "-defaultdomain all" is not valid to reset the default domain of the user.
- The setuser command cannot be used with audit users.

Examples

In the following example, user 3paruser is permitted edit role in the domain Engineering:

```
cli% setuser -adddomain Engineering:edit 3paruser
```

setuseracl

Syntax

```
setuseracl [options] <user_name> <operation> <name_or_pattern>...
```

Description

The setuseracl command sets the Access Control List (ACL).

Options

-add

Adds names or patterns of objects at the end of an existing ACL.

-remove

Removes names or patterns of objects from an ACL.

Specifiers

<user_name>

Specifies the name of the user whose ACL is being set.

<operation>

Specifies the operation for which the ACL is being defined.

The only value currently accepted is updatevv, which updates a snapshot with a new snapshot.

<name_or_pattern>...

Specifies a space-separated list of names or patterns of objects on which the operation (as specified by the <operation> argument) is performed.

The object type is dependent on the specified operation. For example, the objects or names specified for the `updatevv` operation are virtual volume names.

This specifier is required. If an empty string is specified without any option, then any existing ACL for the <user_name>, <operation> is removed.

Authority

- Super, Edit
- Any role granted the `useracl_set` right

Usage

- The -add and -remove options are mutually exclusive.
- If the names or pattern of objects to add already exist, the request is ignored.
- If the names or pattern of objects to remove do not exist, the request is ignored.
- `setuseracl` cannot be used with audit users.

Examples

The following example sets the ACL for user testuser1 and allows testuser1 to update the snapshot virtual volumes of virtual volumes vv1 and vv2:

```
cli% setuseracl testuser1 updatevv vv1 vv2
```

The following example shows how to add vv3 and vv4 to the ACL for user testuser1:

```
cli% setuseracl -add testuser1 updatevv vv3 vv4
```

The following example shows how to remove vv3 from the ACL for user testuser1:

```
cli% setuseracl -remove testuser1 updatevv vv3
```

The following example shows how to clear the ACL for user testuser1:

```
cli% setuseracl testuser1 updatevv ""
```

setvasa

Syntax

```
setvasa [options]
```

Description

The `setvasa` command sets properties of the VASA Provider server, including the option to reset the VASA SSL signed certificate to a self-signed one so that the VASA Provider can be reregistered with a different VMware vCenter.

Options

-f

Forces the operation of the `setvasa` command, bypassing the typical confirmation message.

-reset

If the VASA Provider certificate management mode (see -certmgmt option) is set to client, this option will reset VASA SSL certificate to a new self-signed certificate. This option will also, regardless of certificate management mode, remove all trusted vSphere client certificates. The VASA provider server will be restarted after the self-signed certificate is created and/or trusted certificates are removed. This option will cause all operations between the VASA Provider and the vSphere environment to cease, until the VASA Provider is reregistered in vCenter.

-certmgmt {server|client|locked_client} [-release]

Specifies if the VASA provider must allow management of its SSL certificate by the vSphere infrastructure, or be locally managed on the array. If set to `client` or `locked_client`, vSphere manages the VASA provider's certificate. However, only one vSphere infrastructure may use this VASA Provider. To allow additional vCenters to connect in `client` or `locked_client` modes, distribute the VMCA certificate of the vCenter that first registered the VASA provider, to the additional vCenters.

NOTE: The `client` mode is deprecated. Use `locked_client` instead. If set to `locked_client`, only the first vCenter that registers with the VASA provider has permission to manage the VASA provider's certificate going forward. Using `-release` will open up permission for another vCenter to take over management of the VASA provider's certificate, transferring certificate management from the original vCenter.

WARNING: Once the second vCenter registers, the original vCenter's infrastructure (including associated ESXi hosts) will not have access to the VASA provider, and will lose access to the VVol datastores. The `-release` option must be used only if the first vCenter's infrastructure does not need access to the array.

If the second vCenter is within the same certificate management infrastructure, `setvasa -reset` command is required for a successful take over. If set to "server", the array certificate management tools may be used. See **createcert** for additional information.

-tlsmethod {all|tlsv1|tlsv1_1|tlsv1_2}

Specifies the security method for connections between vSphere clients and VASA Provider server. Defaults to `tlsv1_2`. The `all` option uses secure renegotiation to automatically determine the highest `tls`/`ssl` security method supported between the client and VASA Provider. Setting this value to other than `all` restricts the supported security method to that particular version.

-copypri {low|med|high}

Specifies the priority of internally controlled copy operations, including VM cloning and migration. Defaults to `high`.

Specifiers

None.

Authority

- Super, Service
- Any role granted the `vasa_set` right.

Usage

- Access to all domains is required to run this command.
- With the VASA Provider server enabled, the system prompts for confirmation before continuing.

Examples

To reset VASA Provider certificate:

```
cli% setvasa -reset
WARNING: The VASA Provider server is enabled and will be restarted.
Are you sure you want to continue (y/n)? y
```

setvfs

Syntax

```
setvfs [options] <vfs>
```

Description

The `setvfs` command allows modification of the specified virtual file server (VFS).

Options

Only one of the following certificate options can be specified:

`-certfile <certfile>`

Uses the certificate data contained in this file.

`-certdata <certificate string>`

Uses the certificate data contained in this string.

`-certgen`

Generates and sets a certificate for the VFS.

`-rmcert`

Removes the certificate from the VFS.

`-fpg <fpgname>`

The name of the file provisioning group (FPG) to which the VFS belongs.

`-comment <comment_text>`

Specifies any additional textual information.

`-bgrace <bgrace>`

Specifies the block grace time in seconds for quotas within the VFS.

`-igrace <igrace>`

Specifies the inode grace time in seconds for quotas within the VFS.

Specifiers

`<vfs>`

The name of the VFS to be modified.

Authority

- Super, Edit
- Any role granted `vfs_set` right

Usage

- Access to all domains is required to run this command.
- Grace times are specified in seconds.
- Certificates must be in PEM format, containing both public and private keys.

Examples

The following example modifies the comment assigned to examplevfs:

```
cli% setvfs -comment "An example comment." examplevfs
```

setvv

Syntax

```
setvv [options <arg>] <VV_name|pattern>...
```

Description

Changes the properties associated with a virtual volume. Use the `setvv` to modify volume names, volume policies, allocation warning and limit levels, and the volume's controlling common provisioning group (CPG).

Options

`-name <newname>`

Specifies that the name of the virtual volume be changed to a new name (as indicated by the `<new_name>` specifier) that uses up to 31 characters.

`-wwn <new_wwn>`

Specifies that the WWN of the virtual volume be changed to a new WWN as indicated by the `<new_wwn>` specifier. If `<new_wwn>` is set to "auto", the system will automatically choose the WWN based on the system serial number, the volume ID, and the wrap counter. This option is not allowed for the admitted volume before it is imported, or while the import process is taking place.

Only one of the following options can be specified:

`-udid <Udid_Val>`

Specifies the user defined identifier for VVs for OpenVMS hosts. Udid value should be between 0 to 65535 and can be identical for several VVs.

`-clrrsv`

Specifies that all reservation keys (i.e. registrations) and all persistent reservations on the virtual volume are cleared.

`-clralua`

Restores ALUA state of the virtual volume to ACTIVE/OPTIMIZED state. In ACTIVE/OPTIMIZED state, hosts will have complete access to the volume.

`-exp <time>{d|D|h|H|m|M}`

Specifies the relative time from the current time that volume will expire. `<time>` is a positive integer value and in the range of 1 minute - 1825 days. Time can be specified in days, hours, or minutes. Use "d" or "D" for days, "h" or "H" for hours, or "m" or "M" for minutes following the entered time value.

To remove the expiration time for the volume, enter 0 for `<time>`.

`-comment <comment>`

Specifies any additional information up to 511 characters for the volume. Use `-comment ""` to remove the comments.

`-f`

Do not ask for confirmation before setting or modifying volumes with retention time (`-retain`) or modifying the host DIF policy (`-pol xxx_host_dif`).

-retain <time>{d|D|h|H|m|M}

Specifies the amount of time, relative to the current time, that the volume will be retained. <time> is a positive integer value and in the range of 1 minute - 1825 days. Time can be specified in days, hours, or minutes. Use "d" or "D" for days, "h" or "H" for hours, or "m" or "M" for minutes following the entered time value.

NOTE: If the volume is not in any domain, then its retention time cannot exceed the value of the system's VVRetentionTimeMax. The default value for the system's VVRetentionTimeMax is 14 days. If the volume belongs to a domain, then its retention time cannot exceed the value of the domain's VVRetentionTimeMax, if set. The retention time cannot be removed or reduced once it is set. If the volume has its retention time set, it cannot be removed within its retention time. If both expiration time and retention time are specified, then the retention time cannot be longer than the expiration time.

This option requires the HPE 3PAR Virtual Lock license. Contact your local service provider for more information.

-pol <policy>[,<policy>...]

Specifies the following policies that the created virtual volume follows.

If an argument is not specified, the option defaults to stale_ss.

- stale_ss—Specifies that invalid snapshot volumes are permitted. Failure to update snapshot data does not affect the write to the base volume, but the snapshot is considered invalid.
- no_stale_ss—Specifies that invalid snapshot volumes are not permitted. Failure to update a snapshot is considered a failure to write to the base volume.
- one_host—This constrains the export of a volume to one host or one host cluster (when cluster names may be used as a host name).
- no_one_host—This policy should only be used when exporting a virtual volume to multiple hosts for use by a cluster-aware application, or when "port presents" VLUNs are used. This is the default policy setting.
- 3par_host_dif—This policy enables support for 3PAR implementation of host based DIF.
- std_host_dif—This policy enables support for standard SCSI implementation of host based DIF.
- no_host_dif—This policy disables support for host DIF.
- zero_detect—This policy enables the storage system to scan for zeros in the incoming write data. This feature when used during physical copy to a TPVV will avoid allocating space for blocks containing zero. When used with a Thin Persistence license, this feature reclaims allocated space when zero blocks are written to the TPVV. This policy is only applicable for the base TPVV.
- no_zero_detect—This policy disables the storage system to scan for zeros in the incoming write data to reclaim allocated space on the volume. This is the default policy setting.

-snp_cpg <snp_cpg>

Specifies provisioining the volume snapshot spacefrom the specified CPG. With no snp_cpg defined, or when no snapshots exist for the volume, you can set the snp_cpg to any CPG. If snapshots exist the snp_cpg may only be changed if the current snp_cpg is part of an AOCFG, and only to another CPG in the AOCFG.

-snp_aw <percent>

Enables a snapshot space allocation warning. A warning alert is generated when the reserved snapshot space of the VV exceeds the indicated percentage of the VV size.

-snp_al <percent>

Sets a snapshot space allocation limit. The snapshot space of the VV is prevented from growing beyond the indicated percentage of the virtual volume size.

The following options can only be used on **fully provisioned** volumes:

-usr_cpg <usr_cpg>

Specifies that the volume user space is to be provisioned from the specified CPG. The `usr_cpg` of a VV can only be changed if its existing `usr_cpg` is in an AOCFG, and the target CPG is in the same AOCFG.

The following options can only be used on **thinly provisioned** volumes:

-usr_aw <percent>

This option enables user space allocation warning. Generates a warning alert when the user data space of the TPVV exceeds the specified percentage of the virtual volume size.

-usr_al <percent>

Indicates the user space allocation limit. The user space of the TPVV is prevented from growing beyond the indicated percentage of the virtual volume size. After this limit is reached, any new writes to the virtual volume will fail.

-spt <sectors_per_track>

Defines the virtual volume geometry sectors per track value that is reported to the hosts through the SCSI mode pages. The valid range is between 4 to 8192 and the default value is 304.

-hpc <heads_per_cylinder>

Allows you to define the virtual volume geometry heads per cylinder value that is reported to the hosts though the SCSI mode pages. The valid range is between 1 to 255 and the default value is 8.

Specifiers

<vvname|pattern>

Specifies the virtual volume name or all virtual volumes that match the pattern specified, using up to 31 characters. The patterns are glob-style patterns. (For more information, enter `clihelp sub,globpat`.) Valid characters include alphanumeric characters, periods, dashes, and underscores.

Authority

- Super, Edit
- Any role granted the `vv_set` right can edit the properties of base volumes.
- Any role granted the `vvcopy_set` right can edit the properties of physical copies of volumes.
- Any role granted the `sv_set` right can edit the properties of virtual copies of volumes.

Usage

- At least one option must be specified.
- To view policies assigned to the system's virtual volumes, issue the `showvv -pol` command (see, **showvv**).
- It is not possible to rename a virtual volume that is already associated with a Remote Copy group.
- Changing the usr_cpg or snp_cpg of a VV only changes where new allocations occur, and does not move any existing allocations.
- The snp_cpg and usr_cpg can only be set for a base VV.
- See the *HPE 3PAR Command Line Interface Administrator Guide* for instructions on creating CPGs and volumes.

Examples

The following example sets the policy of virtual volume vv1 to no_stale_ss:

```
cli% setvv -pol no_stale_ss vv1
```

Use the following command to change the name:

```
cli% setvv -name newtest test
```

The following example modifies the WWN of virtual volume vv1:

```
cli% setvv -wwn 50002AC0001A0024 vv1
```

The following example modifies the udid value for virtual volume vv1:

```
cli% setvv -udid 1715 vv1
```

setvvolsc

Syntax

```
setvvolsc [options] set:<vvset>
```

Description

`setvvolsc` can be used to create and remove storage containers for VMware Virtual Volumes (vVols). vVols are managed by the vSphere environment, and storage containers are used to maintain a logical collection of them. No physical space is pre-allocated for a storage container. In the HPE 3PAR OS, special VV sets (see **showvvset**) are used to manage vVol storage containers.

Options

-create

An empty existing <vvset> not already marked as a vVol Storage Container will be updated. The VV set should not contain any existing volumes (see -keep option below), must not be already marked as a storage container, nor may it be in use for other services, such as for remote copy groups, QoS, etc.

-remove

If the specified VV set is a VVol storage container, this option will remove the VV set storage container and remove all of the associated volumes. The user will be asked to confirm that the associated volumes in this storage container should be removed.

-keep

Used only with the -create option. If specified, allows a VV set with existing volumes to be marked as a vVol storage container. However, this option should only be used if the existing volumes in the VV set are VVols.

-f

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

Specifiers

<vvset>

The Virtual Volume set (VV set) name, which is used, or to be used, as a vVol storage container.

Authority

- Super, Edit
- Any role granted the `vvset_set` right

Examples

The following example can be used to create a new vVol storage container, known as FinanceTeam:

```
cli% createvvset FinanceTeam
cli% setvvolsc -create set:FinanceTeam
```

The following example can be used to remove an existing vVol storage container, and all its associated volumes:

```
cli% setvvolsc -remove set:FinanceTeam
Are you sure you wish to remove the FinanceTeam storage container and
all its associated volumes? (y/n): y
```

setvvset

Syntax

```
setvvset [options <arg>] <setname>
```

Description

The `setvvset` command sets the parameters and modifies the properties of a Virtual Volume(VV) set.

Options

`-comment <comment>`

Specifies any comment or additional information for the set. The comment can be up to 255 characters long. Unprintable characters are not allowed.

`-name <newname>`

Specifies a new name for the VV set using up to 27 characters.

Specifiers

`<setname>`

Specifies the name of the vv set to modify.

Authority

- Super, Edit
- Any role granted the `vvset_annotate` right can only annotate the comment of a volume set
- Any role granted the `vvset_set` right can set any volume set property

Usage

None.

Examples

To rename a set from foo to bar:

```
cli% setvvset -name bar foo
```

To change the comment on a set bar:

```
cli% setvvset -comment "This used to be set foo" bar
```

setwsapi

Syntax

```
setwsapi [options]
```

Description

Sets properties of the Web Services API server, including options to enable or disable the HTTP and HTTPS ports.

Options

-f

Forces the operation of the `setwsapi` command, bypassing the typical confirmation message.

At least one of the following options are required:

-timeout *<value>*

Specifies the value for the idle session timeout for a WSAPI session. The *<value>* is a positive integer in the range of 3-1440 minutes (3 minutes to 24 hours). Changing the session timeout takes effect immediately and affects already opened and subsequent WSAPI sessions.

-http enable | disable

Enables or disables the HTTP port.

-https enable | disable

Enables or disables the HTTPS port.

-pol <policy>[,<policy>...]

Sets the WSAPI server policy:

tls_strict

Accepts only TLS connections using TLS 1.2 with secure ciphers. `tls_strict` is the default policy setting.

no_tls_strict

Accepts only TLS connections using TLS 1.0 - 1.2, if HTTPS is enabled.

per_user_limit

The maximum number of sessions allowed per user is 80% of the system resource usage.

no_per_user_limit

The maximum number of sessions allowed per user is the system resource usage. This is the default setting.

-evtstream {enable | disable}

Enables or disables the event stream feature in support of the Server Sent Event (SSE) protocol. Defaults to enable.

The default timeout value is 15 minutes.

Specifiers

None.

Authority

- Super, Service
- Any role granted the `wsapi_set` right

Usage

- Access to all domains is required to run this command.
- When the Web Services API server is active, a warning message showing the current status of the Web Services API server is displayed and you will be prompted for confirmation before continuing. The -f option forces the action without a warning message and prompt.
- Setting the session timeout alone is not service affecting and will not restart the WSAPI server. However, if the timeout option is specified along with service affecting options like `-pol` the WSAPI server will restart.

Examples

To disable the HTTPS port:

```
cli% setwsapi -https disable
        WARNING: The Web Services API server is active and will restart.
        Are you sure you want to continue (y/n)? y
```

The following sets the `tls_strict` policy:

```
cli% setwsapi -pol tls_strict
WARNING: The Web Services API server and the
3PAR GUI are enabled and will restart.
Are you sure you want to continue (y/n)? y
```

showalert

Syntax

```
showalert [options]
```

Description

The `showalert` command displays the status of system alerts. When issued without options, all new customer alerts are displayed.

Authority

Any role in the system

Options

The selection of alerts to be displayed is controlled by selecting one of the following options:

-n

Specifies that only new customer alerts are displayed.

This is the default.

-a

Specifies that only acknowledged alerts are displayed.

-f

Specifies that only fixed alerts are displayed.

-all

Specifies that all customer alerts are displayed.

The format of the alert display is controlled by the following options:

-d

Specifies that detailed information is displayed. Cannot be specified with the -oneline option.

-oneline

Specifies that summary information is displayed in a tabular form with one line per alert. For customer alerts, the message text will be truncated if it is too long unless the -wide option is also specified.

-svc

Specifies that only service alerts are displayed. This option can only be used with the -d or -oneline formatting options.

-wide

Do not truncate the message text. Only valid for customer alerts and if the-oneline option is also specified.

Specifiers

None.

Usage

- Without any options, the `showalert` command displays all customer alerts in the New state.
- Set the status of alerts by issuing the **`setalert`** command.
- Alerts can be removed by issuing the **`removealert`** command.

Examples

The following example displays new customer alerts on a system:

```
cli% showalert -n

 Id           : 1
 State        : New
 Message Code : 0x2200de
 Time         : 2015-07-17 20:14:29 PDT
 Severity     : Degraded
 Type         : Component state change
 Message      : Node 0, Power Supply 1, Battery 0 Degraded (Unknown)

 Id           : 1
 State        : New
 Message Code : 0x2200de
 Time         : 2008-07-17 20:14:29 PDT
 Severity     : Degraded
 Type         : Component state change
 Message      : Node 0, Power Supply 1, Battery 0 Degraded (Unknown)

 2 alerts
```

showaocfg

Syntax

`showaocfg [options <arg>] [<AOCFG_name>...|<pattern>...]`

Description

The `showaocfg` command shows Adaptive Optimization (AO) configurations in the system.

Options

`-domain <domain_name_or_pattern,...>`

Shows only AO configurations that are in domains with names matching one or more of the <domain_name_or_pattern> argument. This option does not allow listing objects within a domain of which the user is not a member. Patterns are glob-style (shell-style) patterns. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.)

Specifiers

`[<AOCFG_name>...|<pattern>...]`

Specifies that AO configurations matching either the specified AO configuration name or those AO configurations matching the specified pattern are displayed. This specifier can be repeated to display information for multiple AO configurations. If not specified, all AO configurations in the system are displayed.

Authority

Any role in the system

Usage

AO will limit the space utilization of a CPG to the lowest of: max, warn, or limit. If none of these values is set for the AOCFG tier or CPG, then AO will only be bounded by the available raw space of the CPG characteristics.

Examples

Use the following to get a description of each column:

```
clihelp -col showaocfg
```

showauthparam

Syntax

```
showauthparam
```

Description

The `showauthparam` command shows authentication parameters.

Options

None.

Specifiers

None.

Authority

Any role in the system

Usage

The authentication parameters are set by `setauthparam` command. See the **`setauthparam`** command for more information about the authentication parameters.

Examples

The following is example output from the `showauthparam` command:

```
cli% showauthparam
Param               -----------------Value-----------------
ldap-server         xxx.xxx.xx.xx
ldap-server-hn      domaincontroller.work.com
binding             sasl
sasl-mechanism      GSSAPI
kerberos-realm      NTDOM1.work.COM
accounts-dn         OU=Users,DC=work,DC=com
account-obj         user
account-name-attr   sAMAccountName
memberof-attr       memberOf
edit-map            CN=Software,CN=Users,DC=work,DC=com
browse-map          CN=Eng,CN=Users,DC=work,DC=com
domain-name-attr    description
group-obj           group
domain-name-prefix !InServDomain=
```

showbanner

Syntax

```
showbanner [-cli|-ssh|-all]
```

`showsshbanner` (deprecated)

Description

The `showbanner` command displays the banner that was set with the **`setbanner`** command.

Options

-cli

Shows CLI banner.

-ssh

Shows SSH banner (default).

-all

Shows all banners.

Specifiers

None.

Authority

Any role in the system

Usage

None.

Examples

The following example displays the CLI banner that was set by using the `setbanner -cli` command:

```
cli% showbanner -cli
WARNING : Unauthorized access to this system is forbidden and will be
prosecuted by law. By accessing this system, you agree that your actions
may be monitored if unauthorized usage is suspected.
```

showbattery

Syntax

```
showbattery [options] [<node_ID> ...]
```

```
showbattery -listcols
```

Description

Displays battery status information such as serial number, expiration date and battery life, which could be helpful in determining battery maintenance schedules.

Options

-listcols

Lists the columns available to be shown with the `-showcols` option described below (issue '`clihelp -col showbattery`' for help on each column).

-showcols <column>[,<column>...]

Selects the columns to be shown using a comma-separated list of column names. For this option, the full column names are shown in the header.

Issue '`showbattery -listcols`' to list the available columns.

-d

Specifies that detailed battery information, including battery test information, serial numbers, and expiration dates, is displayed.

-log

Shows battery test log information. This option is not supported on HPE 3PAR 7000, nor on HPE 3PAR 8000 series systems.

-i

Shows battery inventory information.

-state

Shows detailed battery state information.

-s

This is the same as `-state`. This option is deprecated and will be removed in a future release.

-svc

Displays inventory information with HPE serial number, spare part, etc.

This option must be used with `-i` option, and it is not supported on HPE 3PAR 10000 systems

Specifiers

<node_ID>...

Displays the battery information for the specified node ID(s). This specifier is not required. Node_ID is an integer from 0 through 7.

Authority

Any role in the system

Usage

Battery information is set by issuing the **setbattery** command.

Examples

The following example displays battery status information:

```
cli% showbattery
Node Assem_Serial -State- -Service_LED- ChrgLvl(%) -ExpDate-- Expired Testing
   0 00000028     OK      Off                 100 03/08/2020 No      No
   1 00000046     OK      Off                 100 03/09/2020 No      No
   2 00000120     OK      Off                 100 03/12/2020 No      No
   3 00000123     OK      Off                 100 03/12/2020 No      No
```

The following example displays detailed battery information:

```
cli% showbattery -d
--------------Node 0 PS 0 Battery 0---------------
Node ID                 :   0
Power Supply ID         :   0
Battery ID              :   0
Manufacturer            :   MAG
Model                   :   0800-0016-50.0B
Serial Number           :   70315366
State                   :   OK
Charge State            :   FullyCharged
Charge Level(%)         :   100
Max Battery Life(mins)  :   25
Expired                 :   No
Test in Progress        :   No
Expiration Date         :   03/08/2020
```

The following example displays the inventory information for the battery using the `showbattery -i` command:

```
cli% showbattery -i
Node PS Bat -Manufacturer- -Model- -Serial-
   0  0   0  --             --      FFFFFFFF
   0  1   0  --             --      FFFFFFFF
   1  0   0  --             --      --
   1  1   0  --             --      --
```

The following example displays the battery test log information using the `showbattery -log` command:

```
cli% showbattery -log
Node Test Result Dur(mins) ---------Time----------
   0    0 Passed         1 2015-04-14 12:12:23 MDT
   0    1 Passed         3 2015-04-15 09:18:36 MDT
   0    2 Passed         3 2015-04-16 08:02:57 MDT
   0    3 Passed         3 2015-04-21 07:03:32 MDT
   0    4 Passed         3 2015-04-22 07:04:34 MDT
   0    5 Passed         1 2015-04-24 07:53:04 MDT
   0    6 Passed         1 2015-07-01 08:45:23 MDT
   0    7 Passed         1 2015-05-07 10:32:13 MDT
   0    8 Passed         1 2015-05-22 06:39:54 MDT
   0    9 Passed         1 2015-06-05 06:41:14 MDT
   1    0 Passed         1 2015-04-14 12:13:24 MDT
   1    1 Passed         3 2015-04-15 09:18:37 MDT
   1    2 Passed         3 2015-04-16 08:02:57 MDT
   1    3 Passed         3 2015-04-21 07:03:31 MDT
   1    4 Passed         3 2015-04-22 07:04:34 MDT
   1    5 Passed         1 2015-04-24 07:53:14 MDT
   1    6 Passed         1 2015-06-19 08:44:07 MDT
   1    7 Passed         1 2015-05-07 10:09:12 MDT
   1    8 Passed         1 2015-05-22 06:40:55 MDT
   1    9 Passed         1 2015-06-05 06:42:15 MDT
```

The following example displays the detailed status of the battery using the `showbattery -s` command:

```
cli% showbattery -state
Node PS Bat -State- -Detailed_State-
   0  0   0 OK      Normal
   0  1   0 OK      Normal
   1  0   0 OK      Normal
   1  1   0 OK      Normal
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showbattery
```

showblock

Syntax

```
showblock [options <arg>] <dev> <block> [<eblock>]
```

Description

The `showblock` command displays block mapping information for virtual volumes, logical disks, and physical disks.

Options

`-d`

Specifies that detailed information is displayed for the specified device and block.

Specifiers

<dev>

The <dev> specifier is specified as one of the following arguments:

vv <VV_name> usr|snp|adm

Specifies the virtual volume name and the area (usr, snp, or adm) of that volume for the block mapping information to be displayed.

ld <LD_name>

Specifies the logical disk name.

pd <PD_ID>

Specifies the ID of the physical disk.

<block>

Specifies the 512 byte block number on the specified device.

[<eblock>]

Specifies an end range when used with the <block> specifier. Additional mapping for blocks at the start of each device mapping boundary for the range indicated by <block> and <eblock> is displayed. This specifier is optional.

Authority

Any role in the system

Usage

None.

Examples

The following example displays block mapping information for block 0 in the user space of virtual volume admin:

```
cli% showblock vv admin usr 0
VVname Spc    VVBlock         LDname    LDBlock PDid Chnk    PDblock
 admin usr 0x00000000    admin.usr.0 0x00000000   88    0 0x00080000
     =   =          =              = 0x00000000   58    0 0x00080000
```

The following example displays detailed block mapping information for block 0 in the user space of virtual volume admin:

```
cli% showblock -d vv admin usr 0
 VV/LD                   Name      Block                Region
VV usr:                   admin 0x00000000 0x00000000-0x0007FFFF (0-256MB)
LD    :             admin.usr.0 0x00000000 0x00000000-0x0007FFFF (0-256MB)

  PD Chnk               PDStep      Block                LDStep
  88    0 0x00080000-0x000801FF 0x00080000 0x00000000-0x000001FF
  58    0 0x00080000-0x000801FF 0x00080000 0x00000000-0x000001FF
```

The following example displays block mapping information from blocks 0x100 through 0x400 in the user space of virtual volume admin:

```
cli% showblock vv admin usr 0x100 0x400
VVname Spc    VVBlock         LDname    LDBlock PDid Chnk    PDblock
 admin usr 0x00000100    admin.usr.0 0x00000100   88    0 0x00080100
     =   =          =              = 0x00000100   58    0 0x00080100
 admin usr 0x00000200    admin.usr.0 0x00000200   87    0 0x00080000
     =   =          =              = 0x00000200   56    0 0x00080000
 admin usr 0x00000400    admin.usr.0 0x00000400   70    1 0x00100000
     =   =          =              = 0x00000400   54    0 0x00080000
```

The following example displays detailed block mapping information from blocks 0x100 through 0x400 in the user space of virtual volume admin:

```
cli% showblock -d vv admin usr 0x100 0x400
 VV/LD                    Name      Block                 Region
VV usr:                   admin 0x00000100 0x00000000-0x0007FFFF (0-256MB)
LD    :             admin.usr.0 0x00000100 0x00000000-0x0007FFFF (0-256MB)

  PD Chnk                PDStep      Block                 LDStep
  88    0 0x00080000-0x000801FF 0x00080100 0x00000000-0x000001FF
  58    0 0x00080000-0x000801FF 0x00080100 0x00000000-0x000001FF
-----------------------------------------------------------------
 VV/LD                    Name      Block                 Region
VV usr:                   admin 0x00000200 0x00000000-0x0007FFFF (0-256MB)
LD    :             admin.usr.0 0x00000200 0x00000000-0x0007FFFF (0-256MB)

  PD Chnk                PDStep      Block                 LDStep
  87    0 0x00080000-0x000801FF 0x00080000 0x00000200-0x000003FF
  56    0 0x00080000-0x000801FF 0x00080000 0x00000200-0x000003FF
-----------------------------------------------------------------
 VV/LD                    Name      Block                 Region
VV usr:                   admin 0x00000400 0x00000000-0x0007FFFF (0-256MB)
LD    :             admin.usr.0 0x00000400 0x00000000-0x0007FFFF (0-256MB)

  PD Chnk                PDStep      Block                 LDStep
  70    1 0x00100000-0x001001FF 0x00100000 0x00000400-0x000005FF
  54    0 0x00080000-0x000801FF 0x00080000 0x00000400-0x000005FF
...
```

showcage

Syntax

`showcage [options] [<cagename> ...]`

`showcage -sfp [<cagename> ...]`

`showcage -ddm [<cagename> ...]`

`showcage -i [<cagename> ...]`

`showcage -svc -i [<cagename> ...]`

`showcage -state [<cagename> ...]`

Description

The `showcage` command displays information about drive cages.

Options

`-sfp`

Specifies information about the SFP(s) attached to a cage. Currently, additional SFP information can only be displayed for DC2 and DC4 cages.

`-ddm`

Displays the DDM threshold information if available, e.g., reading, warning (high/low), alarm (high/low).

`-i`

Specifies that inventory information about the drive cage is displayed. If this option is not used, then only summary information about the drive cages is displayed.

`-svc`

Specifies that inventory information, including HPE serial number, spare part number, and model name, is displayed. This option must be used with `-i` option.

-state

Specifies that detailed state information is displayed. If the `State` is Normal, `DetailedState` is also Normal. If the `State` is non-Normal, `DetailedState` displays a list of conditions in the cage which map to cage alerts and status in the -all option.

-e

Displays error information. This option is not valid for DCS1 or DCS2 cages.

-c

Uses cached information to display drive cage information. Displays information faster because the cage does not need to be probed, but information could be out of date.

Specifiers

<cagename>...

Specifies a drive cage name for which information is displayed. This specifier can be repeated to display information for multiple cages. If no specifiers are used, the command defaults to displaying information about all cages in the system.

Authority

Any role in the system

Examples

The following example displays summary information about a system's drive:

```
cli% showcage
Id Name  LoopA Pos.A LoopB Pos.B Drives Temp  RevA RevB Model FormFactor State
0 cage0 0:0:1     0 1:0:1     1     12 21-24 2.88 2.88 DCS6  SFF        Normal
1 cage1 0:0:1     1 1:0:1     0     12 20-22 2.88 2.88 DCS6  SFF        Normal
2 cage2 2:0:1     0 3:0:1     1     12 19-22 2.88 2.88 DCS6  SFF        Normal
3 cage3 2:0:1     1 3:0:1     0     12 19-22 2.88 2.88 DCS6  SFF        Normal
```

The following example displays inventory information about a system's drive cages:

```
cli% showcage -i
-----------------------------------------Midplane-----------------------------------------
Cage -Name- -Manufacturer- --Assem_Part--- -Assem_Serial- -Firmware- -Type- ------WWN-------
   0 cage0  HPE            G3J45-60301     n/a            n/a        DCS6   5001438032F9CA80
   1 cage1  HPE            G3J45-60301     n/a            n/a        DCS6   5001438035176900
   2 cage2  HPE            G3J45-60301     n/a            n/a        DCS6   5001438032FADD40
   3 cage3  HPE            G3J45-60301     n/a            n/a        DCS6   5001438032FB15C0

-------------------------------Interface Card--------------------------------
Cage -I/F_card- -Name- -Manufacturer- --Assem_Part--- -Assem_Serial- -Firmware-
   0          1 SAS B  HPE            E7W10-60402     PDXBPB1LM7X013 2.88
   0          0 SAS A  HPE            E7W10-60402     PDXBPB1LM7X00C 2.88
  ...
   3          1 SAS B  HPE            E7W10-60402     PDXBPB1LM7X0CE 2.88
   3          0 SAS A  HPE            E7W10-60402     PDXBPB1LM7X05C 2.88

--------------------------Power Supply---------------------------
Cage PS -----Name----- -Manufacturer- -Assem_Part- -Assem_Serial-
   0  0 Power Supply 0 DELTA          511777-001   5ANLD0C4D7U1SA
   0  1 Power Supply 1 DELTA          511777-001   5ANLD0C4D7U1SD
  ...
   3  0 Power Supply 0 DELTA          511777-001   5ANLD0C4D7U1XM
   3  1 Power Supply 1 DELTA          511777-001   5ANLD0C4D7U1XJ
```

The following example displays all the components status information about a system's drive cages:

```
cli% showcage -all
-------------------------Cooling--------------------------
...
------------------------Connector-------------------------
...
```

```
----------------------Enclosure-------------------------
...
----------------------Expander--------------------------
...
-----------------------IOModule-------------------------
...
----------------------Magazine--------------------------
...
-------------------------Power--------------------------
...
--------------------------SEP---------------------------
...
----------------------Temperature-----------------------
...
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showcage
++++++++++++++++++++++++Column Help++++++++++++++++++++++++
col,showcage - Explanation of column headers for showcage
Id                  Numeric cage identifier.
Name                Cage name.
LoopA               Port for A loop.
Pos.A               Position on A loop.
LoopB               Port for B loop.
Pos.B               Position on B loop.
Drives              Drive count in the cage.
Temp                Temperature range.
RevA                Controller firmware revision on A side.
RevB                Controller firmware revision on B side.
Model               Cage model.
FormFactor          Supported form factor for the drives, LFF (large) or SFF (small).
State               Cage status.
```

showcert

Syntax

The syntax for the `showcert` command can be one of the following:

```
showcert [-showcols <column>,[<column>...]] [-service <SSL_service_name(s)>]
[-type <certificate_type(s)>]
```

```
showcert {-pem|-text} [-service <SSL_service_name(s)>] [-type
<certificate_type(s)>] [-file <filename>]
```

```
showcert -listcols
```

Description

The `showcert` command has two forms. The first is a table with a high level overview of the certificates used by the SSL Services. This table is customizable with the `-showcols` option. The second form provides detailed certificate information in either human readable format or in PEM (Privacy Enhanced Mail) format. It can also save the certificates in a specified file. With both forms the user is able to select the certificates with the `-service` and `-type` options.

Options

`-listcols`

Displays the valid table columns.

`-showcols <column>,[<column>...]`

Changes the columns displayed in the table.

-service <SSL_service_name>

Displays only the certificates used by the service(s). Multiple services must be delimited by a comma.

Valid service names are cim, cli, ekm-client, ekm-server, ldap, qw-client, qw-server, syslog-gen-client, syslog-gen-server, syslog-sec-client, syslog-sec-server, wsapi, vasa, and unified-server.

-type <certificate_type>

Displays only certificates of the specified type, e.g., only root CA. Multiple types must be delimited by a comma.

Valid types are csr, cert, intca, and rootca.

-pem

Displays the certificates in PEM format. When a filename is specified the certificates are exported to the file.

-text

Displays the certificates in human readable format. When a filename is specified the certificates are exported to the file.

-file <filename>

Specifies the export file of the `-pem` or `-text` option.

Specifiers

None.

Authority

Any role in the system

Usage

- Use the **createcert** command to create self-signed certificates and **importcert** to import signed certificates.
- An asterisk following the service of a CA indicates that it is used by multiple services.

Examples

The following example shows how to display the certificate in table format.

```
cli% showcert
```

The following example shows how to display the certificate used by the cli in PEM format.

```
cli% showcert -service cli -type cert -pem
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showcert
```

showcim

Syntax

```
showcim [options]
```

Description

The `showcim` command displays the CIM server service state being configured, either enabled or disabled. It also displays the server current running status, either active or inactive. It displays the current status of the HTTP and HTTPS ports and their port numbers. In addition, it shows the current status of the SLP port, that is either enabled or disabled.

Options

-pol

Shows CIM server policy information.

Specifiers

None.

Authority

Any role in the system

Examples

The following example shows the current CIM status:

```
cli% showcim
-Service- -State-- --SLP-- SLPPort -HTTP-- HTTPPort -HTTPS- HTTPSPort PGVer CIMVer
Enabled   Active   Enabled     427 Enabled     5988 Enabled      5989 2.14.1 3.3.1
```

The following example shows the current CIM policy:

```
cli% showcim -pol
--------------Policy---------------
replica_entity,one_hwid_per_view,use_pegasus_interop_namespace,no_tls_strict
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showcim
```

showclienv

Syntax

```
showclienv
```

Description

The `showclienv` command displays the CLI environment parameters.

Options

None.

Specifiers

None.

Authority

Any role in the system

Usage

See the **`setclienv`** command for a complete description of the environment parameters.

Examples

The following example displays the currently set CLI environment variables for sample system TestSystem:

```
cli% showclienv
Parameter Value
csvtable  0
nohdtot   0
hafter    -1
listdom   0
editor    emacs
```

showcorequest

Syntax

```
showcorequest [-d] [-request <requestid>]
showcorequest [-state {approved|pending|rejected|executed|inprogress|
cancelled}]
showcorequest -queuesize
```

Description

The `showcorequest` command displays the status of CO approval requests or queue size.

Options

`-request <requestid>`

Limits the output to `<requestid>`.

`-queuesize`

Displays the queue size. The maximum size of queue is 99.

`-d`

Provides the detailed output. This option can be used only with `-request` option.

Specifiers

None.

Authority

- CO, super and command requester
- CO and super user can view all the requests in queue while other users can see only their own requests
- Any user can view the queue size

Usage

Options `-request`, `-state`, and `-queuesize` are mutually exclusive.

Examples

The following example shows status of request ID 12345:

```
cli% showcorequest -request 12345
```

The following example displays list of all requests and their status:

```
cli% showcorequest
```

The following example displays the queue size of the CO list:

```
cli% showcorequest -queuesize
```

The following example displays the approved requests:

```
cli% showcorequest -state approved
```

showcpg

Syntax

```
showcpg [options <arg>] [<CPG_name>...|<pattern>...]
```

Description

The showcpg command displays common provisioning groups (CPGs) in the system.

Options

-listcols

Lists the columns available to be shown in the -showcols option described below (see "clihelp -col showcpg" for help on each column).

-showcols <column>[,<column>...]

Selects the columns to be shown using a comma-separated list of column names. For this option the full column names are shown in the header. Run "showcpg -listcols" to list the available columns. Run "clihelp -col showcpg" for a description of each column.

The following options cannot be used together with other options except -sortcol, -hist and -domain:

-d

Displays detailed information about the CPGs. The following columns are shown:

```
Id Name Warn% VVs TPVVs TDVVs UsageUsr UsageSnp Base SnpUsed Free Total
LDUsr LDSnp RC_UsageUsr RC_UsageSnp DDSType
```

-r

Specifies that raw space used by the CPGs is displayed. The following columns are shown:

```
Id Name Warn% VVs TPVVs TDVVs UsageUsr UsageSnp Base RBase SnpUsed SnpRUsed
Free RFree Total RTotal
```

-alert

Indicates whether alerts are posted. The following columns are shown:

```
Id Name Warn% UsrTotal DataWarn DataLimit DataAlertW% DataAlertW DataAlertL
DataAlertF
```

-alerttime

Shows times when alerts were posted (when applicable). The following columns are shown:

```
Id Name DataAlertW% DataAlertW DataAlertL DataAlertF
```

-sag

Specifies that the snapshot admin space auto-growth parameters are displayed. The following columns are displayed:

```
Id Name AdmWarn AdmLimit AdmGrow AdmArgs
```

-sdg

Specifies that the snapshot data space auto-growth parameters are displayed. The following columns are displayed:

```
Id Name DataWarn DataLimit DataGrow DataArgs
```

-space (-s)

Shows the space saving of CPGs. The following columns are displayed:

```
Id Name Warn% PrivateBase PrivateSnp Shared Free Total Compaction Dedup
DataReduce Overprov
```

The following options can be used together with other options:

```
-sortcol <col>[,<dir>][:<col>[,<dir>]...]
```

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

-hist

Specifies that current data from the CPG, as well as the CPG's history data is displayed.

-domain <domain_name_or_pattern,...>

Shows only CPGs that are in domains with names matching one or more of the <domain_name_or_pattern> argument. This option does not allow listing objects within a domain of which the user is not a member.

Patterns are glob-style (shell-style) patterns. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.)

-aocfg <aocfg_name>

Shows only CPGs that are in the AOCFG <aocfg_name>.

Specifiers

[<CPG_name>...|<pattern>...]

Specifies that CPGs matching either the specified CPG name or those CPGs matching the specified pattern are displayed. This specifier can be repeated to display information for multiple CPGs. If not specified, all CPGs in the system are displayed.

Authority

Any role in the system

Usage

- When using the `createvv` command, the size of the logical disk space created is the first integer multiple of the RAID set size that is large enough to accommodate the requested virtual volume size.
- For example, with the RAID-5 layout with a set size of 3 GiB, a requested virtual volume size of 8 GiB causes the creation of logical disks with a total size rounded up to an integer multiple of 3 GiB that is 9 GiB. The growth increment of CPGs is similarly rounded up because the growth is done by creating logical disks that must be created in units of the logical disk RAID set size.
- For this command, 1 MiB = 1048576 bytes.
- A Domain column may be included by using the **setclienv** command to set listdom to 1. The listdom option is also set to 1 if the CLI was started with the -listdom option, or if the CLI was started with the TPDLISTDOM environment variable set. Issue "`cli -h`" and "`setclienv -h`" for details about environment variables.

Examples

The following example displays a system's CPGs:

```
cli% showcpg
                ----Volumes---- -Usage- ----------(MiB)-----------
Id Name  Warn% VVs TPVVs TDVVs Usr Snp  Base    Snp   Free  Total
 2 cpg1      -   2     0     0   2   1 24576    512  36352  60928
 0 FC_r1     -   2     2     0   2   2  1024   1024  30720  32768
 1 FC_r6     -   2     0     2   2   1  1536    512  34816  36864
-------------------------------------------------------------------
 3 total                     6   4 27136 103936 101888 130560
```

The following example displays the historical data for the CPGs:

```
cli% showcpg -hist
CPG FC_r6
                    ----Volumes---- -Usage- -------(MiB)--------
Time            Warn% VVs TPVVs TDVVs Usr Snp Base Snp  Free Total
Sep 21 11:56:33     -   2     0     2   2   1 1536 512 34816 36864
Sep 21 03:37:01     -   0     0     0   0   0    0   0     0     0
Sep 20 03:37:01     -   0     0     0   0   0    0   0     0     0
-------------------------------------------------------------------
CPG cpg1
                    ----Volumes---- -Usage- --------(MiB)--------
Time            Warn% VVs TPVVs TDVVs Usr Snp  Base Snp  Free Total
Sep 21 11:56:33     -   2     0     0   2   1 24576 512 36352 60928
Sep 21 03:37:01     -   0     0     0   0   0     0   0     0     0
-------------------------------------------------------------------
CPG FC_r1
                    ----Volumes---- -Usage- --------(MiB)--------
Time            Warn% VVs TPVVs TDVVs Usr Snp Base  Snp  Free Total
Sep 21 11:56:33     -   2     2     0   2   2 1024 1024 30720 32768
Sep 21 03:37:01     -   0     0     0   0   0    0    0     0     0
Sep 20 03:37:01     -   0     0     0   0   0    0    0     0     0
-------------------------------------------------------------------
```

The following example displays the common provisioning group's snapshot data space's autogrowth parameters:

```
cli% showcpg -sdg
         -----(MiB)------
Id Name  Warn Limit  Grow Args
 2 cpg1     -     - 32768 -ha mag -t r6 -p -devtype FC
 1 FC_r6    -     - 32768 -ssz 8 -ha mag -t r6 -p -devtype FC
```

The following example displays the alerts posted for the CPGs:

```
cli% showcpg -alert
               -------------Data--------------
               -(MiB)- -Setting(MiB)-   Alerts
Id Name  Warn%   Total   Warn   Limit W% W L F
 2 cpg1      -   25088   4096    5120  - - Y -
 1 FC_r6     -    2048      -       -  - - - -
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showcpg
```

showdate

Syntax

`showdate`

Description

The `showdate` command displays the date and time for each system node.

Options

None.

Specifiers

None.

Authority

Any role in the system

Usage

- Set date and time information on nodes by issuing the **setdate** command.
- The output of the `showdate` command include an abbreviated timezone which may be non-unique. A unique and parenthesized long format timezone name is appended at the end of each row. The long format timezone name is the official timezone string set via `setdate`.

Examples

The following example displays the date and time for the system node:

```
cli% showdate
Node Date
0    2013-04-17 17:07:44 PDT (US/Pacific)
1    2013-04-17 17:07:44 PDT (US/Pacific)
```

showdomain

Syntax

```
showdomain [options]
```

Description

The `showdomain` command displays a list of domains in a system.

Options

-d

Specifies that detailed information is displayed.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

None.

Authority

Any role in the system

Usage

- If the VVRetentionTimeMax is "--", then the maximum volume retention time for the system is used instead.
- If the VVRetentionTimeMax is 0, then the volume retention time in the domain is disabled.

Examples

The following example displays detailed information about a system's domains:

```
cli% showdomain -d
ID Domain      -----CreationTime------ --Comments--- -VVRetentionTimeMax-
 1 TestDomain1 2015-06-22 14:23:30 PDT Beef Test                  0 Hours
```

```
 2 TestDomain2 2015-06-22 14:23:30 PDT Chicken Test                0 Hours
 3 TestDomain3 2015-06-22 14:23:30 PDT Pork Test                   0 Hours
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showdomain
```

showdomainset

Syntax

```
showdomainset [options] [<setname_or_pattern>...]
```

```
showdomainset -domain [options] [<domainname_or_pattern>...]
```

Description

The `showdomainset` command lists the domain sets defined on the system and their members.

Options

`-d`

Shows a more detailed listing of each set.

`-domain`

Shows domain sets that contain the supplied domains or patterns

Specifiers

`<setname_or_pattern>...`

An optional list of setnames or patterns. If no setnames or patterns are specified all sets are displayed, otherwise only sets with names matching one or more of the setnames or patterns are displayed.

The patterns are glob-style patterns. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.)

`<domainname_or_pattern>...`

Specifies that the domain sets containing domains with the specified names or matching the glob-style patterns should be displayed.

Authority

Any role in the system

Usage

None.

Examples

The “Total” line in the output specifies the total number of unique objects.

Use the following to show all domain sets defined to the system:

```
cli% showdomainset
Id Name      Members
 7 domainset domainset.1
             domainset.2
23 newset    testdomain
-----------------------------------
 2 total      3
```

Use the following to show details of a specific set:

```
cli% showdomainset -d newset
Id Name   Members    Comment
23 newset testdomain this set has been renamed and now has a comment
-------------------------------------------------
 1 total   1
```

Use the following to show domain sets containing domains matching the pattern domainset.*:

```
cli% showdomainset -domain domainset.*
Id Name      Members
 7 domainset domainset.1
             domainset.2
------------------------------
 1 total       2
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showdomainset
```

showeeprom

Syntax

```
showeeprom [options] [<node_ID>...]
```

Description

The `showeeprom` command displays node EEPROM log information.

Options

`-dead`

Specifies that an EEPROM log for a node that has not started or successfully joined the cluster be displayed. If this option is used, it must be followed by a non empty list of nodes.

Specifiers

`<node_ID>...`

Specifies the node ID for which EEPROM log information is retrieved.

Multiple node IDs are separated with a single space (0 1 2). If no specifiers are used, the EEPROM log for all nodes is displayed.

Authority

Any role in the system

Usage

None.

Examples

The following example displays the EEPROM log for all nodes:

```
cli% showeeprom
Node: 0
--------
      Board revision: 0920-1053-03.01
            Assembly: FLH 2007/50 Serial 0039
       System serial: 1000183
        BIOS version: 4.8.34
          OS version: 3.2.2.204
        Reset reason: Unknown
           Last boot: 2015-06-24 02:43:13 MDT
```

```
   Last cluster join: 2015-06-24 02:45:43 MDT
          Last panic: 2015-06-23 16:35:43 MDT
  Last panic request: 2015-06-16 20:33:22 MDT
   Error ignore code: 00
         SMI context: 00
       Last HBA mode: 2a000000
          BIOS state: ff 23 26 27 28 29 2b 80
           TPD state: 34 40 ff 2a 2c 2e 30 32
Code 128 (BIOS update) - Subcode 0x2020103 (2020102) 2015-06-12 17:34:37 MDT
Code 128 (BIOS update) - Subcode 0x2020102 (2020009) 2015-06-12 12:37:45 MDT

Node: 1
--------
...
```

showencryption

Syntax

```
showencryption [options]
```

Description

The `showencryption` command shows Data Encryption information.

Options

`-d`

Provides details on the encryption status.

Specifiers

None.

Authority

- Super, Service
- Any role granted the `control_encryption_status` right

Usage

If the state shows recovery_needed, the enable or latest rekey operation failed. Re-issue the operation to fix this state.

Examples

Use the following to get a description of each column:

```
clihelp -col showencryption
```

showeventlog

Syntax

```
showeventlog [options <arg>]
```

Description

The `showeventlog` command displays the current system event log.

Options

-min <number>

Specifies that only events occurring within the specified number of minutes are shown. The <number> is an integer from 1 through 2147483647.

-more

Specifies that you can page through several events at a time.

-oneline

Specifies that each event is formatted as one line.

-d

Specifies that detailed information is displayed.

-startt <time>

Specifies that only events after a specified time are to be shown. The time argument can be specified as either <timespec>, <datespec>, or both. If you would like to specify both a <timespec> and <datespec>, you must place quotation marks around them; for example, -startt "2012-10-29 00:00".

<timespec>

Specified as the hour (hh), as interpreted on a 24 hour clock, where minutes (mm) and seconds (ss) can optionally be specified. Acceptable formats are hh:mm:ss or hhmm.

<datespec>

Specified as the month (mm or month_name) and day (dd), where the year (yy) can be optionally specified. Acceptable formats are mm/dd/yy, month_name dd, dd month_name yy, or yy-mm-dd. If the syntax yy-mm-dd is used, the year must be specified.

-endt <time>

Specifies that only events before a specified time are to be shown. The time argument can be specified as either <timespec>, <datespec>, or both.

See -startt option for descriptions of <timespec> and <datespec>.

The <pattern> argument in the following options is a regular expression pattern that is used to match against the events each option produces. (For more information, enter the clihelp sub,regexpat command.) For each option, the pattern argument can be specified multiple times by repeating the option and <pattern>. For example:

showeventlog -type Disk.* -type <tpdtcl client> -sev Major The -sev Major displays all events of severity Major and with a type that matches either the regular expression Disk.* or <tpdtcl client>.

-sev <pattern>

Specifies that only events with severities that match the specified pattern(s) are displayed. The supported severities include Fatal, Critical, Major, Minor, Degraded, Informational, and Debug.

-nsev <pattern>

Specifies that only events with severities that do not match the specified pattern(s) are displayed. The supported severities include Fatal, Critical, Major, Minor, Degraded, Informational, and Debug.

-class <pattern>

Specifies that only events with classes that match the specified pattern(s) are displayed.

-nclass <pattern>

Specifies that only events with classes that do not match the specified pattern(s) are displayed.

-node <pattern>

Specifies that only events from nodes that match the specified pattern(s) are displayed.

-nnode <pattern>

Specifies that only events from nodes that do not match the specified pattern(s) are displayed.

-type <pattern>

Specifies that only events with types that match the specified pattern(s) are displayed.

-ntype <pattern>

Specifies that only events with types that do not match the specified pattern(s) are displayed.

-msg <pattern>

Specifies that only events, whose messages match the specified pattern(s), are displayed.

-nmsg <pattern>

Specifies that only events, whose messages do not match the specified pattern(s), are displayed.

-comp <pattern>

Specifies that only events, whose components match the specified pattern(s), are displayed.

-ncomp <pattern>

Specifies that only events, whose components do not match the specified pattern(s), are displayed.

Specifiers

None.

Authority

Any role in the system

Usage

The HPE 3PAR OS stores 44MiB of event logs in segments of 4MiB each by default. There are 11 segments, the active event log and the most recent 10 generations. See also `setsys -h` under EventLogSize and `setsys -h` under EventLogNum (**setsys**). When one segment is filled, events are written to the next segment. When all segments are filled, the oldest is overwritten.

Examples

The following example displays the system events occurred during the last 20 minutes:

```
cli% showeventlog -min 20
```

The following example displays the system events same as above using the `-oneline` option:

```
cli% showeventlog -min 20 -oneline
```

The following example displays CRC errors in the last 10 minutes using the `-msg` option:

```
cli% showeventlog -debug -min 10 -msg CRC
```

The following example displays Remote Copy messages in the last 5 minutes using the `-comp` option:

```
cli% showeventlog -min 5 -comp rmm
```

showfirmwaredb

Syntax

```
showfirmwaredb [options]
```

Description

The `showfirmwaredb` command displays the current database of firmware levels for possible upgrade. If issued without any options, the firmware for all vendors is displayed.

Options

`-n <vendor_name>`

Specifies that the firmware vendor from the SCSI database file is displayed.

`-l`

Reloads the SCSI database file into the system.

`-all`

Specifies current and past firmware entries are displayed. If not specified, only current entries are displayed.

`-sortcol <col>[,<dir>][:<col>[,<dir>]...]`

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

None.

Authority

Any role in the system

Usage

The firmware information displayed by the `showfirmwaredb` command is used when issuing the **`upgradecage`** and **`upgradepd`** commands.

Examples

The following example displays the current database of firmware levels and prints firmware data:

```
cli% showfirmwaredb
Vendor Prod_rev Dev_Id Fw_status Cage_type
Firmware_File
HITACHI [C1C1] DK..DJ-18FC Current DC4
/opt/tpd/fw/drive/MDJFC1C1.BIN
...
SEAGATE [0004] ST318203FC Current DC4
/opt/tpd/fw/drive/0004.lod
...
Source file: /var/opt/tpd/scsi_db.cfg
```

showflashcache

Syntax

```
showflashcache [options] [<name|pattern>]
```

Description

The `showflashcache` command displays the status of the flash cache for each node or displays the flash cache policy for virtual volumes.

Options

`-vvset`

Specifies that the virtual volume set with flash cache enabled will be displayed instead of the status for the nodes. Only virtual volume sets matching the SPECIFIER name or pattern will be displayed.

`-vv`

Specifies that the virtual volumes with flash cache enabled will be displayed instead of the status for the nodes. Only virtual volumes matching the SPECIFIER name or pattern will be displayed.

Specifiers

`<name|pattern>`

Specifies that information is displayed for virtual volumes or virtual volume sets matching the specified name or pattern. This is valid only if -vv or -vvset is also specified. If not specified, the default with -vv is for all virtual volumes and for -vvset is for all virtual volume sets.

Authority

Any role in the system

Usage

- Patterns are glob-style (shell-style) patterns. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.)
- If name or pattern is not specified then all virtual volumes or virtual volume sets will be returned.

Examples

The following examples display the status of the flash cache for each node:

```
cli% showflashcache
                        -(MB)-
   Node  Mode   State     Size    Used%
      0  SSD    Normal  131072       30
      1  SSD    Normal  131072       15
   ----------------------------------------------------
      2  total          262144

   cli% showflashcache -vvset

   Id  VVSetName   AFCPolicy
    0  vs1         enabled
    1  vs2         enabled
    2  vs3         enabled
   ---------------------------------------
    3  total

   cli% showflashcache -vvset vs1

   Id  VVSetName  AFCPolicy
    0  vs1        enabled
   ---------------------------------------
    1  total

   cli% showflashcache -vv

   VVid    VVName         AFCPolicy
   1777    VV20720_0001   enabled
   1778    VV20786_0000   enabled
```

```
  1779    VV20855_0000   enabled
  1780    VV20720_0001   enabled
  1781    VV20786_0000   enabled
  1782    VV20855_0000   enabled
  1783    VV20720_0001   enabled
  1784    VV20786_0000   enabled
  1785    VV20855_0000   enabled
  ------------------------------
    10    total

  cli% showflashcache -vv VV20720_0001

  VVid    VVName         AFCPolicy
  1777    VV20720_0001   enabled
  ------------------------------
     1    total
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showflashcache
```

showfpg

Syntax

```
showfpg [options] [<fpgname>]
```

Description

Displays information on file provisioning groups.

Options

-d

Detailed output.

Specifiers

<fpgname>

Limits output to the specified file provisioning group.

Authority

Any role in the system

Usage

- Access to all domains is required to run this command.
- A version displayed with "*" indicates current version is upgradable.
- The default CPG reflects the CPG used for any additional VVs which may be created during a **growvv**. It does not reflect the current CPG of the VVs in the FPG. This can be changed using the `setfpg -defaultcpg` command (**setfpg**). This should also be run after a **tunevv** on any FPG VVs.

Examples

The following example shows a verbose listing of the file provisioning group, examplefpg:

```
cli% showfpg -d examplefpg
```

showfs

Syntax

```
showfs [options]

showfs -usermap

showfs -usermap [-d] {-username <username> | -groupname <groupname> | -userid
<uid|sid> | -groupid <gid|sid>} [-mapped <true|false>]

showfs -usermap -exportconf -file <client_filepath>

showfs -usermap -export {users|groups} -provider <provider> -file
<client_filepath>
```

Description

Displays information on File Persona nodes.

Options

-obj

Displays the Object default profile settings.

-net

Displays network configuration.

-ad

Displays Active Directory domain information. See, `clihelp -col showfs` for more detail.

-ldap

Displays LDAP configuration. See, `clihelp -col showfs` for more detail.

-auth

Displays the user authentication settings.

-idmap

Displays NFSv4 domain name information.

-pdc

Displays current Preferred Domain Controller configuration.

-rfc2307

Displays RFC2307 configuration.

-smb

Displays values of various SMB tunable parameters.

-supportedversion

Displays all upgradable On-Disk versions.

-usermap

Displays the user mapping status, mapped user/group profiles. Copies the exported users/group entries or mapping configuration to the client.

The following options are only allowed with `-usermap`:

-username <username>

Displays the mapped information for the specified username. `<username>` should be in UPN format.

-userid <uid|sid>

Displays the mapped information for the specified user uid/sid.

-groupname <groupname>

Displays the mapped information for the specified groupname. <groupname> should be in UPN format.

-groupid <gid|sid>

Displays the mapped information for the specified group gid/sid.

-d

Displays the detailed information of mapped users. This option is allowed only for users with super role.

-exportconf

Exports and copies the mapping configuration to the client. If this option is specified, "-file <filepath_on_client>" must also be specified.

-export {users|groups}

Specifies the type of exported entries to be copied to the client. If this option is specified, "-provider <provider>" and "-file <filepath_on_client>" must also be specified. See "setfs usermap" for the command to create the exports.

-provider {Local|Ldap|ActiveDirectory}

Specifies the type of provider to be used to copy the exported entries. This option can only be used with -export.

-file <filepath_on_client>

Specifies the file path on client to which the mapping configuration or exported users/groups are to be copied to. This option can be used with only -export and -exportconf.

-mapped {true|false}

Specifies whether mapped user details are to be displayed or its own details are to be displayed. By default, the mapped user details will be displayed. This option can only be used with -username, -userid, and -d.

-nfs

Displays global health message for NFS shares and NFS configuration.

Specifiers

None.

Authority

- Any role in the system
- Super role is required to use -file and -d options.

Usage

- Access to all domains is required to run this command.
- If no options are specified a summary of the nodes status is displayed.
- The options -export and -exportconf are not available with SSH access to CLI.

Examples

```
cli% showfs
Node FSNode State   Active InCluster -----Version----- ---N:S:P--- BondMode  MTU
```

```
   0 Yes    Running Yes    Yes       1.4.0.78-20170713 0:2:1,0:2:2        6 1500
   1 Yes    Running No     Yes       1.4.0.78-20170713 1:2:1,1:2:2        6 1500
   2 Yes    Running No     Yes       1.4.0.78-20170713 2:2:1,2:2:2        6 1500
   3 Yes    Running No     Yes       1.4.0.78-20170713 3:2:1,3:2:2        6 1500
-------------------------------------------------------------------------------
   4 total
```

```
cli% showfs -obj
Non-SSL Port                        : 80
SSL Port                            : 443
Allow Persistent Connections        : true
Timeout for Persistent Connections : 5
Maximum Clients                     : 50
Read Block Size (KB)                : 8.00
Write Block Size (KB)               : 8.00
```

```
cli% showfs -net
IP_Address   Node        Subnet VLAN_Tag Network
10.135.4.221    0 255.255.255.0        0 user
10.135.4.222    1 255.255.255.0        0 user
Default Route: 10.135.4.254 user
DNS servers:   10.35.0.1 10.10.0.6 10.10.15.8
DNS suffixes:  ibftc.vpi.example.com 2003ad.lab
```

```
cli% showfs -ad
Domain Name  : 2003AD.LAB
NetBIOS Name : 2003AD
Forest       : 2003ad.lab
Status       : Online
problemDescription : -
correctiveAction   : -
```

```
cli% showfs -pdc
 Enabled: true

    Domain      Hostname            IPAddress      Site
    sample.com  america.sample.com  10.35.155.222  america
    ------------------------------------------------------
               1 total
```

```
cli% showfs -ldap
LDAP Server       : 100.10.10.10
Administrator DN  : cn=Manager,dc=example,dc=com
Search Base       : dc=example,dc=com
Netbios Name      : hostname
SSL/TLS           : disabled
Status            : OK
Clone Servers     : 2.2.2.0,ldap1.example.com
```

```
cli% showfs -auth
Auth Order                   : ActiveDirectory Local
Domain Prefix Enabled        : true
Supported AD Encryption Type : AES
```

```
cli% showfs -idmap
NFSv4 domain name: example.com
```

```
cli% showfs -rfc2307
RFC2307: enabled
```

```
cli% showfs -smb
SMB TUNABLES PARAMETERS
enableoplocks              : true
```

```
signingenabled             : false
signingrequired            : false
ignorewritethroughrequests : true
supportpersistenthandles   : true
smb3dialectenable          : true
enablesmb1ad               : true
enablesmb2ad               : true
enablesmbleases            : true
enabledirleases            : true
enablesmb1                 : true
enablesmb2                 : true
```

```
cli% showfs -usermap
Mapping Enabled    : true
State              : OK
Health Description : Component is healthy.
Corrective Action  : --
```

```
cli% showfs -usermap -username trk1@2008ad
----------Mapping-----------
From         Type To
2008AD\trk1 =>   2008ad\trk2
```

```
cli% showfs -usermap -groupid 1176528715
----------------Mapping-----------------
From              Type To
2008AD\trk1_group =>   2008ad\trk2_group
```

The following two commands are allowed only for users with super role:

```
cli% showfs -usermap -userid 1176528713 -d
SID                      : S-1-5-21-2943099029-2375420575-3763763779-550730
UID                      : 1176528714
GID                      : 1175978497
UPN                      : trk2@2008AD.LAB
NetBIOS Name             : 2008AD
SAM Account Name         : trk2
Domain                   : CN=trk2,DC=2008ad,DC=lab
Primary Group Name       : 2008AD\domain^users
Alias Name               : --
Map Found                : true
Primary Group SID        : S-1-5-21-2943099029-2375420575-3763763779-513
Additional Info          : trk2
Password Expired         : false
Password Never Expires   : true
Prompt Password Change   : false
User Can Change Password : true
Account Disabled         : false
Account Expired          : false
Account Locked           : false
Mapping From             : 2008AD\trk1
Mapping To               : 2008ad\trk2
Mapping Type             : =>
```

```
cli% showfs -usermap -groupname trk1_group@2008ad -d
GID              : 1176528716
SID              : S-1-5-21-2943099029-2375420575-3763763779-550732
NetBIOS Name     : 2008AD
UNIX Name        : 2008AD\trk2_group
SAM Account Name : trk2_group
Domain           : CN=trk2_group,DC=2008ad,DC=lab
Alias Name       : --
Map Found        : true
```

```
Mapping From     : 2008AD\trk1_group
Mapping To       : 2008ad\trk2_group
Mapping Type     : =>
```

```
cli% showfs -supportedversion
Supported Versions:    11.0 11.1 12.0 12.1 12.2
```

```
cli% showfs -nfs
    Overall NFS Health
    Node    State ----------Detailed Description----------
    0       OK    NFS Component on node node0fs is healthy
    1       OK    NFS Component on node node1fs is healthy
    2       OK    NFS Component on node node2fs is healthy
    3       OK    NFS Component on node node3fs is healthy

    NFS Configuration Settings
    Delegations Enabled : true
    Server GID Enabled  : true
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showfs
```

showfsarchive

Syntax

The syntax for the `showfsarchive` command can be one of the following:

```
showfsarchive pol [-fstore <fstorename>] [-fpg <fpgname>] [<vfs>]
```

```
showfsarchive files [-basepath <basepath>] {-files
<filepath>[,<filepath>,...] | -importfile <source_path> | -inputfile
<pathoffile>} [-fpg <fpgname>] <vfs>
```

```
showfsarchive scan [-jobid <jobid>] [-exportsummary | -export <target_path>]
-fstore <fstorename> [-fpg <fpgname>] <vfs>
```

```
showfsarchive auditlogs [-export {all|<logname>}] [-fpg <fpgname>] <vfs>
```

```
showfsarchive export <taskid> <target_path>
```

Description

The `showfsarchive` command displays retention policy, retention setting for specified files and information of already running validation scans.

Options

`-fpg <fpgname>`

Specifies the FPG that contains the VFS. This is required if VFS exists under multiple FPGs.

`-fstore <fstorename>`

Files store name.

`-basepath <basepath>`

Base path of files for which retention settings are to be displayed.

`-files <filepath>[,<filepath>,]...`

Relative path of files from the above base path for which retention settings are to be displayed.

-inputfile <pathoffile>

Path of the input file containing list of files for which retention settings are to be displayed. The input file should be placed in file persona namespace (.admin file store). It should be absolute path of the input file, which is visible to the file persona namespace.

-importfile <source_path>

Path of the input file containing list of files for which retention settings are to be displayed. The input file should be present on client machine. The max size of import file can be up to 2GB.

-jobid <jobid>

Job Id of the running data validation scan process.

-exportsummary

Exports the summary of validation scan to .admin file store.

-export {all|<logname>}

Exports either all or a specific audit log to the .admin file store inside the VFS. If the string "all" is specified, it will export all of the log files and if the name of a log file is specified, it will export the given log file.

-export <target_path>

Exports the summary file of validation scan to the specified target path.

Specifiers

<vfs>

Specifies the name of the VFS.

<taskid>

Specifies the task id which mentions the result file.

<target_path>

Specifies the target path on client machine to export the result file.

Subcommands

pol

Displays the retention policy configured at VFS or FSTORE level for specific FPG.

files

Lists the retention setting for the specified file at the given path or for all the files at the given path (Including wildcard characters in filenames). This subcommand will spawn a task and return the taskid.

scan

This command inquires about the retention validation scan process with the provided job id. If no job Id is given, all validation scans are listed.

Inquiry of validation scan contains the following fields:

- Validated - File has been compared against a stored checksum of content and metadata.
- Skipped - File could not be validated and was skipped.
- ContInconsistent - Content checksum is corrupted or inconsistent.
- MetaInconsistent - metadata checksum is corrupted or inconsistent.
- JobError - The validation job ended with errors.

Once the job is completed, the summary file of the validation job can be exported to .admin file store using the "-exportsummary" option.

auditlogs

This subcommand will list all the audit logs at the VFS level.

export

This subcommand is used to export the result file that is generated upon admin retention command execution to the specified target location.

Authority

- Any role in the system
- Only the `-export` option requires "Edit" role.
- The "`export`" subcommand, and the "`-importfile <source_path>`" and "`-export <target_path>`" options can be accessed only by users with Super role.

Usage

- Access to all domains is required to run this command.
- Columns displayed with "-" indicate not applicable.
- Inquiry of validation scan contains the following fields:
  - Validated - File has been compared against a stored checksum of content and metadata.
  - Skipped - File could not be validated and was skipped.
  - ContInconsistent - Content checksum is corrupted or inconsistent.
  - MetaInconsistent - metadata checksum is corrupted or inconsistent.
  - JobError - The validation job ended with errors.
- The following subcommand and options are not available with SSH access to CLI:
  - `export`
  - `-importfile <source_path>`
  - `-export <target_path>`

Examples

The following example displays the properties of retention policy configured on VFS vfs1:

```
cli% showfsarchive pol -fpg fpg1 vfs1
```

The following example displays the properties of retention policy configured on fstore fstore1:

```
cli% showfsarchive pol -fstore fstore1 -fpg fpg1 vfs1
```

The following example will list the retention setting for the file(s) specified in filepath:

```
cli% showfsarchive files -basepath /fpg1/vfs1
                    -files /fstore/abcd.txt,/fstore2/xyz.txt
                    -fpg fpg1 vfs1
cli% showfsarchive files -basepath /fpg1/
                   -inputfile /fpg1/vfs1/.admin/s2.txt -fpg fpg1 vfs1
cli% showfsarchive files -importfile /home/worm/file.txt -fpg fpg1 vfs1
```

The following example inquires about data validation scan process:

```
cli% showfsarchive scan -fstore fstore1 -fpg fpg1 vfs1
```

The following example inquires about data validation scan process of a particular job:

```
cli% showfsarchive scan -jobid 10 -fstore fstore1 -fpg fpg1 vfs1
```

The following example exports summary of a particular job to .admin fstore:

```
cli% showfsarchive scan -exportsummary -jobid 10 -fstore fstore1
                                       -fpg fpg1 vfs1
```

The following example shows admin audit logs:

```
cli% showfsarchive auditlogs vfs1
```

The following example exports all audit logs to the .admin folder of VFS:

```
cli% showfsarchive auditlogs -export all vfs1
```

The following example exports specified audit log to .admin folder of VFS:

```
cli% showfsarchive auditlogs -export audit_fpg1_vfs1.log vfs1
```

The following example exports summary of a particular job to /home/archive_results.txt on client machine:

```
cli% showfsarchive scan -export /home/archive_results.txt -jobid 10
                                    -fstore fstore1 -fpg fpg1 vfs1
```

The following example exports result file of CLI task 1234 to /home/WORM/output.txt on client machine:

```
cli% showfsarchive export 1234 /home/WORM/output.txt
```

showfsaudit

Syntax

The syntax for `showfsaudit` can be one of the following:

```
showfsaudit pol [-fpg <fpgname>] [-vfs <vfsname>]

showfsaudit log [-fpg <fpgname>] <vfs>

showfsaudit log -export [-fpg <fpgname>] <vfs>

showfsaudit log -export -file <client_filepath> [-fpg <fpgname>] <vfs>
<logfile>

showfsaudit protobuf -export <client_filepath>

showfsaudit smb [-fpg <fpgname>] [-vfs <vfsname>]

showfsaudit tz
```

Description

Displays File Access Audit settings.

Options

The following options are specific to the **pol** and **smb** subcommands:

`-fpg <fpgname>`

Limits the display to virtual file servers (VFS) contained within the FPG.

-vfs <vfsname>

Limits the display to the specified virtual file server.

The following options are specific to the **log** and **protobuf** subcommands:

-export

Exports specified logfile to .admin folder on file store or path on client's machine.

The following options are specific to the **log** subcommand:

-file <client_filepath>

Specifies the target path on client machine to export the log file. If `-file` is not specified, exports log file to .admin directory on file store.

Specifiers

The following specifiers are specific to the **log** subcommand:

<logfile>

Logs file to be exported.

<vfs>

The VFS for which file access audit log policy file details to be displayed.

The following specifiers are specific to the **protobuf** subcommand:

<client_filepath>

Specifies the target path on client machine to export the protocol specific event definition file. Extension of client_filepath needs to be ".tar".

Subcommands

pol

Displays file access audit policy related details.

log

Displays list of file access audit policy log files generated.

protobuf

Exports file access audit policy related protocol specific event definition files.

smb

Displays file access audit global settings and admin audit share settings for SMB protocol.

tz

Displays the time zone set for audit log files.

Authority

- Any role in the system.
- The `showfsaudit log -export` command requires Edit role to access .admin directory.
- The `showfsaudit log -export -file <client_filepath>` and `showfsaudit protobuf -export` commands can be executed by Super role only.

Usage

None.

Examples

The following example displays file access audit policy defined for specified VFS in a cluster:

```
cli% showfsaudit pol -vfs vfs1
VFS name              : vfs1
FPG name              : fpg1
Audit host            : 079d888f-8a2f-41a3-b0de-151b06ee192c_5
Queue names           : nfs,smb
Authorized user       : testuser
Exchange name         : hpe.exchange
Replica Enabled       : true
ClientType            : internal
LogRotationCriteria   : 100mb
RetentionPeriod       : 10d
Format                : json
```

The following example displays list of compressed log files:

```
cli% showfsaudit log vfs1
Audit Logs
-------------
nfs.log_2017-01-17-10-37-23.zip
```

Exporting the active log file to the sub-directory "fileauditing" under the .admin file store:

```
cli% showfsaudit log -export -fpg fpg1 vfs1
```

Copying log file to client path:

```
cli% showfsaudit log -export -file /home/smblog.zip -fpg fpg1 vfs1 activeLogs_2017-02-10-09-00-28.zip
```

Copying protocol specific event definition files to client path:

```
cli% showfsaudit protobuf -export /home/cluster1_protofiles.tar
```

The following example displays information for vfs1:

```
cli% showfsaudit smb -vfs vfs1
```

The following example displays the time zone set for audit log files:

```
# showfsaudit tz America/Denver
```

showfsav

Syntax

The syntax for the `showfsav` command can be one of the following:

```
showfsav
showfsav svc
showfsav scan [-fpg <fpgname>] [-fstore <fstore>] <vfs>
showfsav pol [-fpg <fpgname>] [-fstore <fstore>] [<vfs>]
showfsav quar [-fpg <fpgname>] <vfs>
```

Description

The `showfsav` command displays antivirus properties for File Persona. When no options are specified, the virus scan engine (VSE) information is displayed. With appropriate subcommands, the policies or scans associated with a virtual file server (VFS) or file store may be displayed.

Options

-fpg <fpgname>

Limits the display to virtual file servers (VFS) contained within the FPG.

-fstore <fstore>

Specifies the file store for which to display the policies or active scans.

Specifiers

<vfs>

Specifies the VFS for which to display the policies or active scans.

Subcommands

svc

Specifies that status of AV service to be displayed.

scan

Specifies that scan information to be displayed.

pol

Specifies that policy information to be displayed.

quar

Specifies that statistics of quarantined files to be displayed.

Authority

Any role in the system

Usage

- Access to all domains is required to run this command.
- The `stopfsav scan <vfs_name> <scan_id>` command can be used to stop the running scan.
- The `stopfsav scan -pause <vfs_name> <scan_id>` command can be used to pause the running scan (see **stopfsav**).
- `[o]` represents the setting has been overridden in file store.
- The `quar` subcommand is only applicable for VFS.

Examples

The following example displays various antivirus properties for a VFS named "testvfs":

```
cli% showfsav pol testvfs
```

The following example displays various antivirus properties for the file store named "engineering" under VFS "testvfs":

```
cli% showfsav pol -fstore engineering testvfs
```

showfsgroup

Syntax

```
showfsgroup [options] [<groupname|pattern> ...]
```

Description

The `showfsgroup` command displays information about groups associated with the File Persona subsystem.

Options

-d

Displays full details on the local groups.

Specifiers

<groupname>

Limits the output to the group specified.

Authority

Any role in the system

Usage

Access to all domains is required to run this command.

Examples

The following example displays information about group admin:

```
cli% showfsgroup admin
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showfsgroup
```

showfshare

Syntax

```
showfshare
showfshare {smb|nfs|obj|ftp} [options <arg>] [<sharename>|<pattern>]
showfshare {smb|nfs|obj|ftp} -dirperm [-fpg <fpgname>] -vfs <vfs> -fstore
<fstorename> <sharename>
```

Description

The `showfshare` command displays file share information for supported protocols.

Options

The following options are for **all** subcommands:

-dirperm

Displays ACL permissions, UNIX permissions and owner group permissions of a share directory. If this option is used, -fstore and share name must also be specified.

-fpg <fpgname>

Specifies the file provisioning group name. This limits the share output to those shares associated with the specified file provisioning group.

-vfs <vfs>

Specifies the virtual file server name. This limits the share output to those shares associated with the specified virtual file server. If this option is specified, but -fpg is not specified, the command will find out the file provisioning group based on <vfs>.

However, if <vfs> exists under multiple file provisioning groups, -fpg must be specified.

-fstore <fstore>

Specifies the file store name. This limits the share output to only those shares associated with the specified file store. If this is specified, option -vfs must be specified.

-pat

Specifies the file share names using the glob-style pattern. Shares which have the name matching any of the specified glob-style patterns will be displayed. The -pat option can specify a list of patterns, and it must be used if specifier <pattern> is used.

-d

Shows share attributes which are not displayed by other options.

Specifiers

<pattern|sharename>

Displays only shares with names matching the specified <sharename> or one of the glob-style patterns.

Subcommands

smb

Displays file shares information for SMB.

nfs

Displays file shares information for NFS.

obj

Displays file shares information for Object.

ftp

Displays file shares information for FTP.

If no subcommand is specified, this command displays file shares for all protocols.

Authority

Any role in the system

Usage

- Access to all domains is required to run this command.
- To specify VFS or file store filters, the parent components must be specified.
- When displaying allow/deny permissions for SMB shares, this command displays users/groups configured on Active Directory or the LDAP server in the format of "<netbios>\<name>". The <netbios> is the Active Directory or LDAP server NetBIOS name, which can be found by running "showfs -ad" or "showfs -ldap".

Examples

The following example displays all SMB file shares with all attributes:

```
cli% showfshare smb -d
```

The following example displays all SMB file shares matching the pattern "share*":

```
cli% showfshare smb -pat "share*"
```

The following example displays SMB file shares associated with the VFS "myvfs":

```
cli% showfshare smb -vfs myvfs
```

The following example displays all NFS file shares:

```
cli% showfshare nfs
```

The following example displays NFS file shares associated with the virtual file server "myvfs":

```
cli% showfshare nfs -vfs myvfs
```

The following example displays a single Object file share:

```
cli% showfshare obj -vfs myvfs myshare
```

The following example displays a single FTP file share:

```
cli% showfshare ftp -vfs myvfs myshare
```

The following example displays share directory permissions of NFS file share "myshare":

```
cli% showfshare nfs -dirperm -vfs myvfs -fstore myfstore myshare
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showfshare
```

showfsip

Syntax

```
showfsip [options] <vfs>
```

Description

The `showfsip` command shows the network config of a Virtual File Server.

Options

`-fpg <fpgname>`

Specifies the File Provisioning Group in which the Virtual File Server was created.

Specifiers

`<vfs>`

Specifies the Virtual File Server which is to have its network config modified.

Authority

Any role in the system

Usage

Access to all domains is required to run this command.

Examples

The following example displays the network config of Virtual File Server vfs_1:

```
cli% showfsip -fpg fpg_1 vfs_1
```

showfsnap

Syntax

```
showfsnap [options <arg>] [<snapname>|<pattern>]
```

Description

The `showfsnap` command lists the snapshots for File Persona.

Options

-d

Displays detailed output that displays creation time of the snapshot in addition to other fields.

-fpg <fpgname>

Specifies the file provisioning group (FPG) name. This option limits the snapshot output to those associated snapshots with the specified FPG.

-vfs <vfs>

Specifies the virtual file server (VFS) name. This option limits the snapshot output to those snapshots associated with the specified VFS. If this option is specified, but `-fpg` is not specified, the command will find out the FPG based on `<vfs>`. However, if `<vfs>` exists under multiple FPGs, `-fpg` must be specified.

-fstore <fstore>

Specifies the file store name. This option limits the snapshot output to only those snapshots associated with the specified file store. If this is specified, option `-vfs` must be specified.

-pat

Specifies the snapshot names using glob-style patterns. Snapshots which have the name matching any of the specified glob-style patterns will be displayed. Patterns can be repeated using a comma-separated list. The `-pat` option must be used if `<pattern>` specifier is used.

Specifiers

<snapname|pattern>

Displays only snapshots with names matching the specified `<snapname>` or one of glob-style patterns.

Authority

Any role in the system

Usage

- Access to all domains is required to run this command.
- If no arguments are specified, this command displays all snapshots created in the File Persona cluster.

Examples

The following example displays all snapshots for "myfstore":

```
cli% showfsnap -vfs vfs1 -fstore fstore1
SnapshotName            FPG   VFS  FileStore
2017-09-07T113322_snap1 fpg1  vfs1 fstore1
-----------------------------------------
                     1 total

cli% showfsnap -d
----------------------showfsnap-----------------------
Snapshot Name              : 2017-09-07T113322_snap1
File Provisioning Group    : fpg1
Virtual File Server        : vfs1
File Store                 : fstore1
Creation Time              : 2017-09-07 05:33:22
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showfsnap
```

showfsnapclean

Syntax

The syntax for the `showfsnapclean` command is as follows:

```
showfsnapclean [options] [<fpgname>]
```

Description

The `showfsnapclean` command displays the details of an on-demand snapshot reclamation task active on a file provisioning group.

Options

`-d`

Displays detailed output.

Specifiers

`<fpgname>`

Specifies the name of the file provisioning group.

Authority

Any role in the system.

Usage

Access to all domains is required to run this command.

Examples

The following example displays the running reclamation task on FPG fpg1:

```
cli% showfsnapclean fpg1
---------------ID--------------- FPG   State     -------Start------- --------End-------- CumSpaceRecovered
3fc6487f0f374358a06b21ec743bec0b fpg1  COMPLETED 2017-09-07 05:26:07 2017-09-07 05:26:37                 0
-------------------------------------------------------------------------------------------------------------
                              1 total

cli% showfsnapclean -d fpg1
-----------------------Showfsnapclean-------------------------
Task Id                    : 3fc6487f0f374358a06b21ec743bec0b
File Provisioning Group    : fpg1
Task State                 : COMPLETED
Start Time                 : 2017-09-07 05:26:07
End Time                   : 2017-09-07 05:26:37
Verbose Mode               : NA
Log Level                  : INFO
Strategy                   : maxspace
Entries Scanned            : 0
Entries Reclaimed          : 0
Inodes Reclaimed           : 0
Inodes Skipped             : 0
Average File Size (Kb)     : 0
Cumulative Space Recovered : 0
Errors                     : 0
Exit Status                : OK
```

If the task is not running and there is no history of completed tasks, the following output is displayed:

```
cli% showfsnapclean fpg1
No reclamation task running on FPG fpg1
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showfsnapclean
```

showfsndmp

Syntax

The syntax for the `showfsndmp` command can be:

```
showfsndmp
showfsndmp -conf
showfsndmp -vtl {vtldevices|vtltapes}
showfsndmp -sessions {active|completed} [-d]
```

Description

The `showfsndmp` command displays NDMP properties for File Persona. When no options are specified, this command displays the global service status of NDMP and cluster-wide IQN. With appropriate options specified, VTL devices/tapes information, configuration of DMA, and session status are displayed.

Options

`-conf`

Displays the current value of all NDMP configurable parameters.

`-vtl vtldevices|vtltapes`

Displays list of all target VTL devices when set to `vtl devices`. Lists all target VTL tapes when set to `vtl tapes`.

`-sessions active [-d]`

Displays active/running NDMP sessions. Using -d displays all the details.

`-sessions completed [-d]`

Displays all completed NDMP sessions. Using `-d` displays all the details.

Specifiers

None.

Authority

Any role in the system

Usage

- Access to all domains is required to run this command.
- If no option is specified after `showfsndmp`, global service status and cluster-wide IQN is displayed.
- The command `showfsndmp -conf` displays the NDMP parameters for the File Persona.
- VTL information can be displayed by using the `-vtl` option.
- NDMP sessions-related information can be displayed by using the `-sessions` option.

Examples

The following example will display the global service status of all NDMP services and cluster-wide IQN:

```
cli% showfsndmp
```

The following example will display various NDMP parameters for File Persona:

```
cli% showfsndmp -conf
```

The following example will display a list of all target VTL devices or VTL tapes:

```
cli% showfsndmp -vtl vtldevices
cli% showfsndmp -vtl vtltapes
```

The following example will display the active NDMP sessions:

```
cli% showfsndmp -sessions active
```

The following example will display the active NDMP sessions with details:

```
cli% showfsndmp -sessions active -d
```

The following example will display the completed NDMP sessions:

```
cli% showfsndmp -sessions completed
```

The following example will display completed NDMP sessions with details:

```
cli% showfsndmp -sessions completed -d
```

showfspn

Syntax

```
showfspn [options]
```

Description

The `showfspn` command displays information about SPN and VFS association.

Options

`-fpg <fpgname>`

Specifies the name of FPG that contains VFS.

`-vfs <vfs>`

Specifies the name of the VFS

Specifiers

None.

Authority

Any role in the system

Usage

Access to all domains is required to run this command.

Examples

The following example displays an SPN details matching specified VFS:

```
cli% showfspn -vfs vfs1
```

The following example displays SPN details for all VFS contained in specified FPG:

```
 cli% showfspn -fpg fpg1
```

The following example displays all SPN details:

```
cli% showfspn
```

showfsquota

Syntax

```
showfsquota [options]
```

Description

Displays the quotas for File Persona.

Options

-username *<uname>*

Displays the quotas for the user name

For users configured on Active Directory, use *<domain>**<uname>* or *<ad-netbios>**<uname>* to specify the user (`-username example.com\aduser`). The *<ad-netbios>* variable is the Active Directory NetBIOS name, which you can find by running `showfs -ad`. To join the Active Directory domain with the *<domain>*, use `setfs ad`

For users configured on the LDAP server, use *<ldap-netbios>**<username>* to specify the user (`-username ldaphost\\ldapuser`). The *<ldap-netbios>* variable is the LDAP server NetBIOS name. To find it, run `showfs -ldap`. To create the LDAP configuration with `<ldap-netbios>` use `setfs ldap`.

-groupname *<gname>*

Displays the quotas for the group name.

For groups configured on Active Directory, use *<domain>**<gname>* or *<ad-netbios>**<uname>* to specify the user (`-groupname example.com\adgroup`). The *<ad-netbios>* variable is the Active Directory NetBIOS name. To find it, run `showfs -ad`. To join the Active Directory domain with the *<domain>* use `setfs ad`.

For groups configured on the LDAP server, use *<ldap-netbios>**<gname>* to specify the user (`-groupname ldaphost\\ldapgroup`). To create the LDAP configuration with *<ldap-netbios>* use `setfs ldap`.

-fstore *<fstore>*

The file store of the quotas to be displayed.

-vfs *<vfs>*

Specifies the name of the virtual file server associated with the quotas.

-fpg *<fpgname>*

Specifies the name of the file provisioning group hosting the virtual file server.

Specifiers

None.

Authority

Any role in the system.

Usage

- Access to all domains is required to run this command.
- Only one set of quotas may be specified per call.

- To apply to all users or groups, use the `-all` specifier.
- Users/groups configured on Active Directory or the LDAP server are displayed in the format of <netbios>\<name>. The <netbios> is the Active Directory or LDAP server NetBIOS name, which can be found by running `showfs -ad` or `showfs -ldap`.

Examples

The following example displays the quotas for the virtual file server `examplevfs`:

```
cli% showfsquota -vfs examplevfs
```

showfsroute

Syntax

```
showfsroute [-d] [-target <targetaddr>] [-vlan <vlantag>] [-gateway
<gatewayaddr>]
```

Description

The `showfsroute` command displays all routes including default or created with `createfsroute`.

Options

`-d`

Displays the detailed information for each route.

`-target <targetaddr>`

Takes an IPv4/IPv6 address and lists all routes for this address.

`-vlan <vlantag>`

Takes an integer value and lists routes configured on this vlan.

`-gateway <gatewayaddr>`

Displays all routes with this gateway.

Specifiers

None.

Authority

Any role in the system

Usage

- Access to all domains is required to run this command.
- If no option is specified after `showfsroute` all routes will be listed.

Examples

The following example will display all the routes:

```
cli% showfsroute
```

The following example displays all routes for specified target address:

```
cli% showfsroute -target 10.16.26.111
```

The following example will display the route for specified target address at given vlan tag and gateway:

```
cli% showfsroute -target 10.16.26.111 -vlan 10 -gateway 10.16.26.1
```

The following example will display a detailed list of all routes:

```
cli% showfsroute -d
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showfsroute
```

showfstore

Syntax

The syntax for the `showfstore` command is as follows:

Description

The `showfstore` command displays information on the file stores.

```
showfstore [-fpg <fpgname> [-vfs <vfs> [-fstore <fstore>]]]
```

Options

`-fpg <fpgname>`

Limits the display to virtual file servers contained within the FPG.

`-vfs <vfs>`

Limits the display to the specified virtual file server.

`-fstore <fstore>`

Limits the display to the specified file store.

Specifiers

None.

Authority

Any role in the system.

Usage

- Access to all domains is required to run this command.
- To specify VFS or fstore filters, the parent components must be specified.

Examples

The following example displays all file stores on FPG named "fpg1" under VFS named "vfs1":

```
cli% showfstore -fpg fpg1 -vfs vfs1
                          --------Security--------
Fstore  VFS   FPG  State   Mode    ErrorSuppression  -Comment-

.admin  vfs1  fpg1 normal  legacy  false
fstore1 vfs1  fpg1 normal  ntfs    false
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showfstore
```

showfsuser

Syntax

```
showfsuser [options] [<user|pattern> ...]
```

Description

The `showfsuser` command displays information about local users associated with the File Persona subsystem.

Options

`-d`

Displays full details on the local users.

Specifiers

`<username>`

Limits the output to the user specified.

Authority

Any role in the system

Usage

Access to all domains is required to run this command.

Examples

The following example displays information about user joe:

```
cli% showfsuser joe
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showfsuser
```

showhost

Syntax

```
showhost [options <arg>] [<host_name>... | <pattern> ... | <host_set> ... ]
```

Description

The `showhost` command displays information about defined hosts and host paths in the system.

Options

`-d`

Shows a detailed listing of host and path information. This option can be used with -agent and -domain options.

`-verbose`

Shows a verbose listing of all host information. This option cannot be used with -d.

`-chap`

Shows the CHAP authentication properties. This option cannot be used with -d.

`-desc`

Shows the host descriptor information. This option cannot be used with -d.

-agent

Shows information provided by Host Explorer.

-pathsum

Shows summary information about hosts and paths. This option cannot be used with -d.

-persona

Shows the host persona settings in effect. This option cannot be used with -d.

-listpersona

Lists the defined host personas. This option cannot be used with -d.

-noname

Shows only host paths (WWNs and iSCSI names) not assigned to any host.

This option cannot be used with -d.

-domain <domainname_or_pattern,...> | <domain_set>

Shows only hosts that are in domains or domain sets that match one or more of the specifier <domainname_or_pattern> or set <domainset> arguments. The set name <domain_set> must start with "set:". This specifier does not allow listing objects within a domain of which the user is not a member.

-lesb

Shows the CRC error counts for the host/port.

Specifiers

<host_name>...

Name of the host up to 31 characters in length. This specifier can be repeated to specify multiple hosts.

<pattern>...

Specifies that information is shown for all hosts matching the specified glob-style pattern. This specifier can be repeated to show information for multiple hosts using different patterns.

<host_set>...

Specifies that information about all the hosts that are members of set <host_set> should be displayed. The set name <host_set> must start with "set:". Acts as if all members were individually specified as parameters to the comment. May be repeated to specify multiple host sets.

Authority

Any role in the system

Usage

- If hostnames, host sets or patterns are specified, then hosts with names that match any of the patterns or are members of the sets are listed.
- Otherwise all hosts are listed. Patterns are glob-style (shell-style) patterns. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.)
- Host descriptor information is available only for hosts that have been assigned a name via the **createhost** command.
- A Domain column may be included by using the `setclienv` command to set listdom to 1. The listdom option is also set to 1 if the CLI was started with the -listdom option, or if the CLI was started with the TPDLISTDOM environment variable set. Issue "`cli -h`" and "`setclienv -h`" for details of the environment variables.

Examples

The following example displays detailed host and path information:

```
cli% showhost -d
```

The following will show all hosts in host set, host_set1:

```
cli% showhost set:host_set1
Id Name  Persona      -WWN/iSCSI_Name- Port
 0 host1 Generic-ALUA
 1 host2 Generic-ALUA
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showhost
```

showhostset

Syntax

```
showhostset [options] [<setname_or_pattern>...]
```

```
showhostset -host [options] [<hostname_or_pattern>...]
```

Description

The `showhostset` command lists the host sets defined on the storage system and their members.

Options

`-d`

Shows a more detailed listing of each set.

`-host`

Shows host sets that contain the supplied hostnames or patterns

Specifiers

`<setname_or_pattern>...`

An optional list of setnames or patterns. If no setnames or patterns are specified all sets are displayed, otherwise only sets with names matching one or more of the setnames or patterns are displayed.

The patterns are glob-style patterns. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.)

`<hostname_or_pattern>...`

Specifies that the sets containing hosts with the specified names or matching the glob-style patterns should be displayed.

Authority

Any role in the system

Usage

- A Domain column may be included by using the **`setclienv`** command to set listdom to 1. The listdom option is also set to 1 if the CLI was started with the -listdom option, or if the CLI was started with the TPDLISTDOM environment variable set. Issue "`cli -h`" and "`setclienv -h`" for details of the environment variables.
- Total line in the output specifies the total number of unique objects.

Examples

The following shows all host sets defined to the system:

```
cli% showhostset
Id Name       Members
22 myset      -
18 seta       -
 3 sunv40z-09 sunv40z-09-0
              sunv40z-09-1
              sunv40z-09-2
              sunv40z-09-3
----------------------------
 3 total 4
```

The following shows the details of myset only:

```
cli% showhostset -d myset
Id Name  Members Comment
22 myset -       This is an empty set
-------------------------------------
 1 total 0
```

The following shows the host sets containing host sun40z-09-0:

```
cli% showhostset -host sun40z-09-0
Id Name       Members
 3 sunv40z-09 sunv40z-09-0
              sunv40z-09-1
              sunv40z-09-2
              sunv40z-09-3
----------------------------
 1 total 4
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showhostset
```

showinventory

Syntax

```
showinventory [options]
```

Description

Shows information about all the hardware components in the system.

Options

-svc

Displays inventory information with HPE serial number, spare part number, and so on. It is not supported on HPE 3PAR 10000 systems.

Specifiers

None.

Authority

Any role in the system

Usage

The inventory is divided into sections and inventory for each of these sections can be shown separately with the following commands:

Section	Command
Node	`shownode -i`
Battery	`showbattery -i`
Port	`showport -i`
Cage	`showcage -i`
Disk	`showpd -i`

Examples

None.

showiscsisession

Syntax

```
showiscsisession [options <arg>] [<node:slot:port>|<pattern>...]
```

Description

The `showiscsisession` command shows the iSCSI sessions.

Options

`-sortcol <col>[,<dir>][:<col>[,<dir>]...]`

Sorts command output based on column number (`<col>`). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (`<dir>`) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

`-d`

Specifies that more detailed information about the iSCSI session is displayed. If this option is not used, then only summary information about the iSCSI session is displayed.

`-state`

Specifies the connection state of current iSCSI sessions. If this option is not used, then only summary information about the iSCSI session is displayed.

Specifiers

`<node:slot:port>`

Requests that information for a specified port is displayed. This specifier can be repeated to display session information about multiple ports. If not specified, session information for all ports in the system is displayed.

`node`

Specifies the node using a number from 0 through 7.

slot

Specifies the PCI slot in the specified node. Valid range is 0 through 9.

port

Specifies the port using a number from 1 through 4.

<pattern>

Specifies that sessions matching the specified glob-style pattern are displayed. This specifier can be repeated to display session information about multiple ports and iSCSI names. If not specified, session information for all ports and iSCSI names in the system is displayed.

Authority

Any role in the system

Usage

None.

Examples

The following example demonstrates iSCSI session information with VLAN support:

```
N:S:P --IPAddr--- TPGT TSIH Conns ------------iSCSI_Name------------ -------StartTime------- VLAN State  QID
0:4:1 10.10.31.29   41    0     1 iqn.1994-05.com.redhat:dbd942a820f 2019-10-15 15:21:16 PDT    - ONLINE 0
1:4:1 10.20.31.29  141    0     1 iqn.1994-05.com.redhat:dbd942a820f 2019-10-15 15:15:55 PDT    - ONLINE 0
-------------------------------------------------------------------------------------------------------------
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showiscsisession
```

showld

Syntax

```
showld [options <arg>] [<LD_name> |<pattern>...]
```

Description

The `showld` command displays configuration information about the system's LDs.

Options

-cpg <CPG_name|pattern>...

Requests that only LDs in common provisioning groups (CPGs) that match the specified CPG names or patterns be displayed. Multiple CPG names or patterns can be repeated using a comma-separated list (for example, `-cpg <CPG_name>,<CPG_name>...`).

-vv <VV_name|pattern>...

Requests that only LDs mapped to virtual volumes that match and of the specified names or patterns be displayed. Multiple volume names or patterns can be repeated using a comma-separated list (for example, `-vv <VV_name>,<VV_name>...`).

-domain <domainname|pattern>...

Only shows LDs that are in domains with names that match any of the names or specified patterns. Multiple domain names or patterns can be repeated using a comma separated list (for example, `-domain <domainname_name>,<domainname_name>...`).

-degraded

Only shows LDs with degraded availability.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

-d

Requests that more detailed layout information is displayed.

-ck

Requests that checkld information is displayed.

-p

Requests that policy information about the LD is displayed.

-state

Requests that the detailed state information is displayed. This is the same as -s.

-s

Requests that the detailed state information is displayed. This option is deprecated and will be removed in a subsequent release.

Specifiers

<LD_name>...

Requests that information for a specified LD is displayed. This specifier can be repeated to display configuration information about multiple LDs. If not specified, configuration information for all LDs in the system is displayed.

<pattern>...

Specifies that the LD matching the specified glob-style pattern is displayed. This specifier can be repeated to display configuration information about multiple LDs. If not specified, configuration information for all LDs in the system is displayed.

Authority

Any role in the system

Usage

- For this command, KB=1024 bytes and MB=1048576 bytes.
- A Domain column may be included by using the **setclienv** command to set listdom to 1. The listdom option is also set to 1 if the CLI was started with the -listdom option, or if the CLI was started with the TPDLISTDOM environment variable set. Issue "cli -h" and "setclienv -h" for details of the environment variables.

Examples

The following example displays information for system LDs:

```
cli% showld
Id Name          RAID -Detailed_State- Own SizeMB UsedMB Use  Lgct LgId WThru MapV
 2 .srdata.usr.0    1 normal           0/1  40960  40960 C,V     0  ---     N    Y
```

```
 3 .srdata.usr.1    1 normal           1/0  40960  40960 C,V     0  ---     N    Y
 0 admin.usr.0      1 normal           0/1   5120   5120 V       0  ---     N    Y
 1 admin.usr.1      1 normal           1/0   5120   5120 V       0  ---     N    Y


 8 log0.0           1 normal           0/-  20480      0 log     0  ---     Y    N
 9 log1.0           1 normal           1/-  20480      0 log     0  ---     Y    N
10 pdsld0.0         1 normal           1/0   1024      0 P,F     0  ---     Y    N
11 pdsld0.1         1 normal           0/1   4096      0 P       0  ---     Y    N
12 pdsld0.2         1 normal           1/0   4096      0 P       0  ---     Y    N
...
```

The following example displays information for LDs in a four node system:

```
cli% showld
Id Name           RAID -Detailed_State- Own      SizeMB  UsedMB Use  Lgct LgId WThru MapV
 4 .srdata.usr.0    1 normal           0/1/2/3   14336   14336 V       0  ---     N    Y
 5 .srdata.usr.1    1 normal           0/1/3/2    6144    6144 V       0  ---     N    Y
 6 .srdata.usr.2    1 normal           1/0/2/3   14336   14336 V       0  ---     N    Y
 7 .srdata.usr.3    1 normal           1/0/3/2    6144    6144 V       0  ---     N    Y
 8 .srdata.usr.4    1 normal           2/3/0/1   14336   14336 V       0  ---     N    Y
 9 .srdata.usr.5    1 normal           2/3/1/0    6144    6144 V       0  ---     N    Y
10 .srdata.usr.6    1 normal           3/2/0/1   14336   14336 V       0  ---     N    Y
11 .srdata.usr.7    1 normal           3/2/1/0    6144    6144 V       0  ---     N    Y
 0 admin.usr.0      1 normal           0/1/2/3    3072    3072 V       0  ---     N    Y
 1 admin.usr.1      1 normal           1/0/3/2    3072    3072 V       0  ---     N    Y
 2 admin.usr.2      1 normal           2/3/0/1    3072    3072 V       0  ---     N    Y
 3 admin.usr.3      1 normal           3/2/1/0    3072    3072 V       0  ---     N    Y
12 log0.0           1 normal           0/-/-/-   20480       0 log     0  ---     Y    N
13 log1.0           1 normal           1/-/-/-   20480       0 log     0  ---     Y    N
14 log2.0           1 normal           2/-/-/-   20480       0 log     0  ---     Y    N
15 log3.0           1 normal           3/-/-/-   20480       0 log     0  ---     Y    N
...
```

The following example displays detailed information about the system LDs:

```
cli% showld -d
Id Name          CPG RAID Own SizeMB RSizeMB RowSz StepKB SetSz Refcnt Avail CAvail -----CreationTime------ ---CreationPattern----
 7 .srdata.usr.0 ---    1 0/1  30720   61440    15    256     2      0 cage  port   2016-07-13 10:06:06 PDT --
 8 .srdata.usr.1 ---    1 1/0  30720   61440    15    256     2      0 cage  port   2016-07-13 10:06:06 PDT --
 0 admin.usr.0   ---    1 0/1   5120   10240     5    256     2      0 cage  port   2016-07-13 10:02:26 PDT --
 1 admin.usr.1   ---    1 1/0   5120   10240     5    256     2      0 cage  port   2016-07-13 10:02:26 PDT --
 2 log0.0        ---    1 0/-  20480   40960     4    256     2      0 cage  port   2016-07-13 10:03:17 PDT -p -nd 0
 3 log1.0        ---    1 1/-  20480   40960     4    256     2      0 cage  port
-------------------------------------------------------------------------------------------------------------------------------------
 6                       112640  225280
```

The following example displays policy information about system LDs:

```
cli% showld -p
Id Name           Keep_Preserved Dev_Type RPM
 7 .srdata.usr.0                1 FC         10
 8 .srdata.usr.1                1 FC         10
 0 admin.usr.0                  1 FC         10
 1 admin.usr.1                  1 FC         10
 2 log0.0                       0 FC         10
 3 log1.0                       0 FC         10
 ...
```

The following example displays detailed states of each LD:

```
cli% showld -state
Id Name              -State- -Detailed_State-
 7  .srdata.usr.0      normal  normal
 8  .srdata.usr.1      normal  normal
 0  admin.usr.0        normal  normal
 1  admin.usr.1        normal  normal
 2  log0.0             normal  normal
 3  log1.0             normal  normal
 ...
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showld
```

showldch

Syntax

```
showldch [options <arg>] <LD_name>
```

Description

The `showldch` command displays configuration information about the chunklet mapping for one logical disk (LD).

Options

`-degraded`

Shows only the chunklets in sets that cause the logical disk availability to be degraded. For example, if the logical disk normally has cage level availability, but one set has two chunklets in the same cage, then the chunklets in that set are shown. This option cannot be specified with option -lformat or -linfo.

`-lformat <form>`

Shows the logical disk's row and set layout on the physical disk, where the line format <form> is one of:

- row—One line per logical disk row.
- set—One line per logical disk set.

`-linfo <info>[,<info>...]`

Specifies the information shown for each logical disk chunklet, where <info> can be one of:

- pdpos—Shows the physical disk position (default).
- pdid—Shows the physical disk ID.
- pdch—Shows the physical disk chunklet.

If multiple <info> fields are specified, each corresponding field will be shown separately by a dash (-).

Specifiers

`<LD_name>`

Specifies the logical disk name.

Authority

Any role in the system

Usage

None.

Examples

The following example displays information about chunklets for logical disk r1.usr.0:

```
cli% showldch r1.usr.0
Ldch Row Set PdPos Pdid Pdch  State Usage Media Sp From  To
   0   0   0 1:3:2   30    0 normal    ld valid  N  ---  ---
   1   0   0 1:3:3   31    3 normal    ld valid  N  10:0 ---
   2   0   1 2:3:1   45    0 normal    ld valid  N  ---  ---
   3   0   1 1:3:0   28    0 normal    ld valid  N  ---  ---
   4   0   2 0:1:2    6    0 normal    ld valid  N  ---  ---
   5   0   2 2:1:3   39    0 normal    ld valid  N  ---  ---
   6   0   3 0:2:0    8    0 normal    ld valid  N  ---  ---
   7   0   3 1:2:2   26    0 normal    ld valid  N  ---  ---
```

```
   8   0   4 2:3:3   47    0 normal    ld valid  N  ---  ---
   9   0   4 1:1:2   22    0 normal    ld valid  N  ---  ---
   10  0   5 2:2:3   43    0 normal    ld valid  N  ---  ---
   11  0   5 0:3:0   12    0 normal    ld valid  N  ---  ---
   12  0   6 0:0:0    0    0 normal    ld valid  N  ---  ---
   13  0   6 1:0:2   18    0 normal    ld valid  N  ---  ---
   14  0   7 2:0:1   33    0 normal    ld valid  N  ---  ---
   15  0   7 2:0:2   34    5 normal    ld valid  N  16:1 ---
```

The following example displays information about degraded chunklets for logical disk r1.usr.0:

```
cli% showldch -degraded r1.usr.0
Ldch Row Set PdPos Pdid Pdch  State Usage Media Sp From  To
   0   0  0  1:3:2   30    0 normal    ld valid  N  --- ---
   1   0  0  1:3:3   31    3 normal    ld valid  N 10:0 ---
  14   0  7  2:0:1   33    0 normal    ld valid  N  --- ---
  15   0  7  2:0:2   34    5 normal    ld valid  N 16:1 ---
cli% showldch -degraded r1.usr.1
No degraded sets in LD r1.usr.1
```

The following example displays row and set layout including pdid and pdch for chunklets for logical disk test.usr.0.

```
cli% showldch -lformat row -linfo pdid,pdch test.usr.0
    ---------Set0---------- ---------Set1----------
row    Ch0   Ch1   Ch2   Ch3   Ch0   Ch1   Ch2   Ch3
  0 2:3:1 1:2:2 0:3:2 1:3:0 0:0:0 2:0:1 1:2:0 0:1:0
  1 2:3:3 1:0:2 0:0:2 2:0:3 0:2:2 1:1:0 2:1:1 1:3:0
  2 0:3:0 2:2:1 1:2:2 2:1:3 0:1:2 1:3:2 2:3:1 0:2:2
  3 0:2:0 1:2:0 2:1:3 1:0:0 2:2:3 0:0:0 1:1:2 0:1:0
  4 2:3:3 1:0:2 0:3:2 2:0:1 0:0:2 1:3:0 2:1:1 1:1:2
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showldch
```

showldmap

Syntax

```
showldmap <LD_name>
```

Description

The `showldmap` command displays the mapping from a logical (LD) disk to virtual volumes (VVs).

Options

None.

Specifiers

<LD_name>

Specifies the logical disk name.

Authority

Any role in the system

Usage

None.

Examples

The following example displays the region of logical disk v0.usr.0 that is used for a virtual volume:

```
cli% showldmap v0.usr.0
Area Start(MB) Length(MB) VVId VVName      VVSp VVOff(MB)
0           0       512     0    v0          usr  0
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showldmap
```

showlicense

Syntax

```
showlicense [options]
```

Description

The `showlicense` command displays the currently installed license key and current license information.

Options

`-raw`

Specifies that the license key originally entered (the raw license) be displayed. The license key is displayed in a manner that is acceptable input for the `setlicense` command.

Specifiers

None.

Authority

Any role in the system

Usage

Without options, a description of which features are enabled is displayed, When the license was generated and any expiration dates are also shown. If the -raw option is specified, the license key will be printed in a manner that is acceptable input for the **`setlicense`** command.

Examples

The following example displays the currently installed license information:

```
cli% showlicense
License key was generated on Wed Dec  4 07:28:26 2019

License features currently enabled:
All features excluding encryption
```

The following example displays the original license key entered on the system:

```
cli% showlicense -raw
60R3-0C1G-60R3-2C1G-60R3-0C9G-70R3-0C1G
60RK-0C0A-FSXZ-8YZ4-Z884-84DW-7CD6-JLKB...
```

showmaint

Syntax

```
showmaint [option]
```

Description

The `showmaint` command displays maintenance window records.

Options

-all

Display all maintenance window records, including active and expired ones. If this option is not specified, only active window records will be displayed.

-sortcol <col>

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- `inc`: Sorts in increasing order. (default)
- `dec`: Sorts in decreasing order.

Specifiers

None.

Authority

Any role in the system

Examples

To display both active and expired maintenance window records:

```
cli% showmaint -all
```

To display maintenance window records with sorted output based on column number:

```
cli% showmaint -sortcol 2
```

To display active maintenance window records:

```
cli% showmaint
```

shownet

Syntax

```
shownet [option]
```

Description

The `shownet` command displays the configuration and status of the administration network interfaces, including the configured gateway and network time protocol (NTP) server.

Options

-d

Shows detailed information.

Specifiers

None.

Authority

Any role in the system

Usage

If the -d option is specified, information for administration, File Persona, node rescue, and Remote Copy over IP (RCIP) interfaces is displayed. This information is useful for debugging possible network issues.

Examples

The following example displays the status of the system administration network interfaces:

```
cli% shownet -d
IP Address: 192.168.5.191 Netmask 255.255.252.0
    Assigned to nodes: 01
    Connected through node 0
    Status: Active

    Admin interface on node 0
    MAC Address: 00:02:AC:43:00:34
    RX Packets:             534389        TX Packets:             121669
    RX Bytes:             67828134        TX Bytes:             60638375
    RX Errors:                   0        TX Errors:                   0
    RX Dropped:                  0        TX Dropped:                  0
    RX FIFO Errors:              0        TX FIFO Errors:              0
    RX Frame Errors:             0        TX Collisions:               0
    RX Multicast:                0        TX Carrier Errors:           0
    RX Compressed:               0        TX Compressed:               0

    Remote copy interface in slot 1 on node 0
    ...

    Admin interface on node 1
    ...
```

Use the following to get a description of each column in the above examples:

```
clihelp -col shownet
```

shownode

Syntax

`shownode [option] [<node_ID> ...]`

Description

The `shownode` command displays an overview of the node-specific properties and its component information. Various command options can be used to display the properties of PCI cards, CPUs, Physical Memory, IDE drives, and power supplies.

Options

The following options are for node summary and inventory information:

`-listcols`

Lists the columns available to be shown with the `-showcols` option described below (see '`clihelp -col shownode`' for help on each column).

By default (if none of the information selection options below are specified) the following columns are shown:

`Node Name State Master InCluster LED Control_Mem Data_Mem Available_Cache`

To display columns pertaining to a specific node component use the `-listcols` option in conjunction with one of the following options: `-pci`, `-cpu`, `-mem`, `-drive`, `-fan`, `-ps`, `-mcu`, `-uptime`.

`-showcols <column>[,<column>...]`

Selects the columns to be shown using a comma-separated list of column names. For this option, the full column names are shown in the header.

Run '`shownode -listcols`' to list Node component columns. Run '`shownode -listcols <node_component>`' to list columns associated with a specific `<node_component>`.

<node_component> can be one of the following options: -pci, -cpu, -mem, -drive, -fan, -ps, -mcu, -uptime.

If a specific node component option is not provided, then -showcols expects Node columns as input.

If a column (Node or specific node component) does not match either the Node columns list or a specific node component columns list, then 'shownode -showcols <cols>' request is denied.

If an invalid column is provided with -showcols, the request is denied.

The -showcols option can also be used in conjunction with a list of node IDs.

Run 'clihelp -col shownode' for a description of each column.

-i

Shows node inventory information in table format.

-d

Shows node and its component information in table format.

The following options are for node component information. These options cannot be used together with options, -i and -d:

-verbose

Displays detailed information in verbose format. It can be used together with the following component options.

-fan

Displays the node fan information.

-pci

Displays PCI card information.

-cpu

Displays CPU information.

-mem

Displays physical memory information.

-drive

Displays the disk drive information.

-ps

Displays power supply information.

-mcu

Displays MicroController Unit information.

-state

Displays the detailed state information for node or power supply (-ps). This is the same as -s.

-s

Displays the detailed state information for node or power supply (-ps). This option is deprecated and will be removed in a subsequent release. Use -state instead.

-uptime

Shows the amount of time each node has been running since the last shutdown.

-svc

Displays inventory information with HPE serial number, spare part, etc. This option must be used with `-i` option, and it is not supported on HPE 3PAR 10000 systems.

Specifiers

<node_ID>...

Displays the node information for the specified node ID(s). This specifier is not required. Node_ID is an integer from 0 through 7.

Authority

Any role in the system

Usage

None.

Examples

The following example displays the operating environment status for all nodes in the system:

```
cli% shownode
                                                    Control Data           Cache
Node -Name-- -State-- Master InCluster ---LED--- Mem(MB) Mem(MB) Available(%)
   0 nodeh7  Degraded No     Yes       GreenBlnk    2048    2048          100
   1 nodeh3c Degraded Yes    Yes       GreenBlnk    2048    2048          100
```

The following examples display detailed information (`-d` option) for the nodes including their components in a table format. The `shownode -d` command can be used to display the tail information of the nodes, including their components in name and value pairs.

```
cli% shownode -mem
Node SlotID --Name--- -Usage- ---Type--- -Manufacturer- -Serial- --Latency-- Size(GiB)
           0     J0101  DIMM0.0.0 Control DDR4_SDRAM Samsung       356AE535
CL10.0/20.0        16
   ...
   1 J0101  DIMM0.0.0 Control DDR4_SDRAM Samsung        350516D5 CL10.0/20.0        16
   ...
   2 J0101  DIMM0.0.0 Control DDR4_SDRAM Samsung        350517E6 CL10.0/20.0        16
   ...
   3 J0101  DIMM0.0.0 Control DDR4_SDRAM Samsung        350507E8 CL10.0/20.0        16
   ...
```

```
cli% shownode -drive
----------------------------Internal Drives----------------------------
Node Drive -Manufacturer- ---Model--- -Serial- -Firmware- Size(MB) Type
   0     0 Seagate        ST9100821AS 5NJ09DF3 3.AAB         95396 SATA
   1     0 Seagate        ST9100821AS 5NJ08NA4 3.AAB         95396 SATA
```

```
cli% shownode -pci
--------------------------PCI Cards--------------------------
Node Slot Type -Manufacturer- -Model- -Serial- -Rev- Firmware
   0    0 FC   QLOGIC         2302    D13503   1     3.3.16
   0    1 FC   QLOGIC         2302    D44171   1     3.3.16
   1    0 FC   QLOGIC         2302    C31037   1     3.3.16
```

```
cli% shownode -cpu
-----------------------------------------CPUs-----------------------------------------
    Node CPUs Cores Threads -Manufacturer- --------------------Model---------------------
    0    1    10      20 GenuineIntel   85 (Intel(R) Xeon(R) Silver 4114 CPU @ 2.20GHz)
    1    1    10      20 GenuineIntel   85 (Intel(R) Xeon(R) Silver 4114 CPU @ 2.20GHz)
```

```
cli% shownode -ps
-------------------------Power Supplies-------------------------
Node PS -PSState-- FanState ACState DCState -BatState- ChrgLvl(%)
```

```
   0  0 OK          OK        OK       OK       OK                  100
   0  1 OK          OK        OK       OK       OK                  100
   1  0 OK          OK        OK       OK       OK                  100
   1  1 OK          OK        OK       OK       OK                  100
```

```
cli% shownode -state
Node -State- -Detailed_State-
   0 OK      OK
   1 OK      OK
   2 OK      OK
   3 OK      OK
```

Use the following to get a description of each column in the above examples:

```
clihelp -col shownode
```

shownodeenv

Syntax

```
shownodeenv [options <arg>...]
```

Description

The `shownodeenv` command displays the node operating environment status, including voltages and temperatures.

Options

`-n <node_ID>...`

Specifies the ID of the node whose environment status is displayed.

Multiple node IDs can be specified as a series of integers separated by a space (1 2 3). If no option is used, then the environment status of all nodes is displayed.

Specifiers

None.

Authority

Any role in the system

Usage

None.

Examples

The following example displays the operating environment status for all nodes in the system:

```
cli% shownodeenv
    Node 0
    ---------
          Measurement   Reading  Lo Limit Hi Limit  Status
          CPU0 1.32V:    1.31 V    1.28 V   1.36 V  Within Tolerance
          CPU1 1.32V:    1.31 V    1.28 V   1.36 V  Within Tolerance
         82563 1.20V:    1.22 V    1.13 V   1.26 V  Within Tolerance
         31154 1.30V:    1.29 V    1.22 V   1.37 V  Within Tolerance
         82563 1.90V:    1.92 V    1.79 V   2.00 V  Within Tolerance
         ...
    Node 1
    ---------
        Measurement     Reading  Lo Limit Hi Limit  Status
        ...
```

Use the following to get a description of each column in the above examples:

```
clihelp -col shownodeenv
```

showpatch

Syntax

```
showpatch [option <arg>]
```

Description

The `showpatch` command displays patches applied to a system.

Options

`-hist`

Provides an audit log of all patches and updates that have been applied to the system.

`-d <ID>`

When used with the `-hist` option, shows detailed history information including the username who installed each package. If `-d` is used with a patch specification, it shows detailed patch information. Otherwise it shows detailed information on the currently installed patches.

Specifiers

None.

Authority

Any role in the system

Usage

This command displays all the patches currently affecting the system if options are not used.

Examples

The following example shows all patches currently installed on the system, with additional detail:

```
cli% showpatch -d
```

The following example shows all updates that have been applied to the system over time, with all detail:

```
cli% showpatch -hist -d
```

The `showpatch`command with a specific individual installed patch number displays the fields below when used with the optional `-d` option:

```
cli% showpatch P### -d
```

```
  Patch ID.          Specifies the patch ID.
  Release Version.   Specifies TPD or UI release affected by the patch.
  Synopsis.          Specifies the purpose of the patch.
  Date.              Specifies the build date of the patch.
  Bugs fixed.        Specifies the bugs fixed.
  Description.       Specifies a detailed description of the problem or fix.
  Affected Packages. Specifies the new packages being changed.
  Obsoletes.         Specifies the patch IDs deleted by this patch.
  Requires.          Specifies the patch IDs of any other patches required by this
 patch.
  Notes.             Specifies any special instructions for the patch.
```

showpd

Syntax

```
showpd [options] [<PD_ID> ...]
showpd -listcols
```

Description

The `showpd` command displays configuration information about a system's physical disks.

Options

`-listcols`

Lists the columns available to be shown in the -showcols option described below (see '`clihelp -col showpd`' for help on each column).

The [options] are generally of two kinds: those that select the type of information that is displayed, and those that filter the list of PDs that are displayed.

By default (if none of the information selection options below are specified) the following columns are shown:

```
Id CagePos Type RPM State Size_MB Free_MB Port_A0 Port_B0 Port_A1 Port_B1
Capacity
```

Options that select the type of information shown include the following:

```
-showcols <column>[,<column>...]
```

Explicitly select the columns to be shown using a comma-separated list of column names. For this option the full column names are shown in the header. For other options that select the type of information, the column names may not be exactly the same and there may be spanning headers on top of the column names.

Run '`showpd -listcols`' to list the available columns.

Run '`clihelp -col showpd`' for a description of each column.

`-i`

Shows disk inventory (inquiry) data.

The following columns are shown:

```
Id CagePos State Node_WWN MFR Model Serial FW_Rev Protocol MediaType
AdmissionTime.
```

`-e`

Shows disk environment and error information. Note that reading this information places a significant load on each disk.

The following columns are shown:

```
Id CagePos Type State Rd_CErr Rd_UErr Wr_CErr Wr_UErr Temp_DegC
LifeLeft_PCT.
```

`-c`

Shows chunklet usage information. Any chunklet in a failed disk will be shown as "Fail".

The following columns are shown:

```
Id CagePos Type State Total_Chunk Nrm_Used_OK Nrm_Used_Fail Nrm_Unused_Free
Nrm_Unused_Uninit Nrm_Unused_Unavail Nrm_Unused_Fail Spr_Used_OK
Spr_Used_Fail Spr_Unused_Free Spr_Unused_Uninit Spr_Unused_Fail.
```

-s

Shows detailed state information.

This option is deprecated and will be removed in a subsequent release.

-state

Shows detailed state information. This is the same as -s.

The following columns are shown:

```
Id CagePos Type State Detailed_State SedState.
```

-path

Shows current and saved path information for disks.

The following columns are shown:

```
Id CagePos Type State Path_A0 Path_A1 Path_B0 Path_B1 Order.
```

-space

Shows disk capacity usage information (in MB).

The following columns are shown:

```
Id CagePos Type State Size_MB Volume_MB Spare_MB Free_MB Unavail_MB
Failed_MB.
```

```
-sortcol <col>[,<dir>][:<col>[,<dir>]...]
```

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

The PD filtering options include:

-failed

Specifies that only failed physical disks are displayed.

-degraded

Specifies that only degraded physical disks are displayed. If both-failed and -degraded are specified, the command shows failed disks and degraded disks.

-p <pattern>

Physical disks matching the specified pattern are displayed.

The following arguments can be specified as patterns for this option:

An item is specified as an integer, a comma-separated list of integers, or a range of integers specified from low to high.

-nd <item>

Specifies one or more nodes. Nodes are identified by one or more integers (item). Multiple nodes are separated with a single comma (e.g. 1,2,3). A range of nodes is separated with a hyphen (e.g. 0-7). The primary path of the disks must be on the specified node(s).

-st <item>

Specifies one or more PCI slots. Slots are identified by one or more integers (item). Multiple slots are separated with a single comma (e.g. 1,2,3). A range of slots is separated with a hyphen (e.g. 0-7). The primary path of the disks must be on the specified PCI slot(s).

-pt <item>

Specifies one or more ports. Ports are identified by one or more integers (item). Multiple ports are separated with a single comma (e.g. 1,2,3). A range of ports is separated with a hyphen (e.g. 0-4). The primary path of the disks must be on the specified port(s).

-cg <item>

Specifies one or more drive cages. Drive cages are identified by one or more integers (item). Multiple drive cages are separated with a single comma (e.g. 1,2,3). A range of drive cages is separated with a hyphen (e.g. 0-3). The specified drive cage(s) must contain disks.

-mg <item>

Specifies one or more drive magazines. The "1." or "0." displayed in the CagePos column of showpd output indicating the side of the cage is omitted when using the -mg option. Drive magazines are identified by one or more integers (item). Multiple drive magazines are separated with a single comma (e.g. 1,2,3). A range of drive magazines is separated with a hyphen(e.g. 0-7). The specified drive magazine(s) must contain disks.

-pn <item>

Specifies one or more disk positions within a drive magazine. Disk positions are identified by one or more integers (item). Multiple disk positions are separated with a single comma (e.g. 1,2,3). A range of disk positions is separated with a hyphen(e.g. 0-3). The specified position(s) must contain disks.

-dk <item>

Specifies one or more physical disks. Disks are identified by one or more integers (item). Multiple disks are separated with a single comma (e.g. 1,2,3). A range of disks is separated with a hyphen (e.g. 0-3). Disks must match the specified ID(s).

-tc_gt <number>

Specifies that physical disks with total chunklets greater than the number specified be selected.

-tc_lt <number>

Specifies that physical disks with total chunklets less than the number specified be selected.

-fc_gt <number>

Specifies that physical disks with free chunklets greater than the number specified be selected.

-fc_lt <number>

Specifies that physical disks with free chunklets less than the number specified be selected.

-devid <model>

Specifies that physical disks identified by their models be selected. Models can be specified in a comma-separated list. Models can be displayed by issuing the "`showpd -i`" command.

-devtype <type>

Specifies that physical disks must have the specified device type (FC for Fast Class, NL for Nearline, SSD for Solid State Drive, SCM for Storage Class Memory) to be used. Device types can be displayed by issuing the "`showpd`" command.

-rpm <number>

Drives must be of the specified relative performance metric, as shown in the "RPM" column of the "showpd" command. For FC and NL drives, the number corresponds to both a performance measure and actual rotational speed. For SSD & SCM drives, this field is N/A.

Disks that satisfy all of the specified characteristics are used.

For example -p -fc_gt 60 -fc_lt 230 -nd 2 specifies all the disks that have greater than 60 and less than 230 free chunklets and that are connected to node 2 through their primary path.

-nodes <node_list>

Specifies that the display is limited to specified nodes and physical disks connected to those nodes. The node list is specified as a series of integers separated by commas (e.g. 1,2,3). The list can also consist of a single integer. If the node list is not specified, all disks on all nodes are displayed.

-slots <slot_list>

Specifies that the display is limited to specified PCI slots and physical disks connected to those PCI slots. The slot list is specified as a series of integers separated by commas (e.g. 1,2,3). The list can also consist of a single integer. If the slot list is not specified, all disks on all slots are displayed.

-ports <port_list>

Specifies that the display is limited to specified ports and physical disks connected to those ports. The port list is specified as a series of integers separated by commas (e.g. 1,2,3). The list can also consist of a single integer. If the port list is not specified, all disks on all ports are displayed.

-w <world-wide_name>...

Specifies the WWN of the physical disk. This option and argument can be specified if the <PD_ID> specifier is not used. This option should be the last option in the command line.

Specifiers

<PD_ID>

Specifies a physical disk ID. This specifier can be used if the -w option is not specified.

Authority

Any role in the system

Usage

- In the showpd output, when the position of the disk is not valid (for example, in the case of a missing disk), the most recent position might be displayed, followed by a question mark (?).
- To see the device type for a physical disk (Fast Class, Nearline, or Solid State Drive), use showpd.
- Without the -i, -c, -e, -s, -path or -space options, basic information about the physical disk is printed. Note that the primary path to the disk is shown by an asterisk (*) in either the APort or BPort column, and a minus (-) sign following a path indicates the port is disabled.
- If the <PD_ID> specifier is not specified and -w is not specified, all disks will be displayed.

Examples

The following example displays chunklet use information for all disks:

```
cli% showpd -c
```

The following example displays chunklet use information for the Nearline disks only:

```
cli% showpd -c -p -devtype NL
```

The following example displays disk information for the Nearline disks on cage 0 and 3 only:

```
cli% showpd -p -cg 0,3 -devtype NL
```

The following example displays disk information for the Nearline disks on cage 0 and sort the output by the cage position:

```
cli% showpd -sortcol 1 -p -cg 0,0 -devtype NL
```

The following example displays all the FC disks in magazine 0 of cage 4 and 5:

```
cli% showpd -p -mg 0 -cg 4,5 -devtype FC
```

The following example displays all the FC disks in magazine 0 of all cages and all FC disks in cages 4 and 5:

```
cli% showpd -p -mg 0 -devtype FC -p -cg 4,5 -devtype FC
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showpd
```

showpdata

Syntax

```
showpdata [<LD_name>]
```

Description

The `showpdata` command displays information about the preserved data in the system.

Options

None.

Specifiers

`<LD_name>`

Requests that preserved data for a specific LD is displayed. This specifier is not required on the command line. If not specified, then the amount of preserved data by LD is displayed.

Authority

Any role in the system

Usage

Access to all domains is required to run this command if no LD name is specified. If an LD name is specified, then access to the domain of the LD is required.

Examples

The following example displays preserved data information for LD admin.usr.3:

```
cli% showpdata admin.usr.3
showpdata v0.usr.0
Preserved LD raid sets
  Ldname   Set     PreservedCnt
     admin.usr.3   0  1
No preserved chunklets
```

showpdch

Syntax

```
showpdch [options]
```

```
showpdch [options] <pdid>|<pattern>...
```
(deprecated usage)

Description

The `showpdch` command displays the status of selected physical disk (PD) chunklets.

Options

The following filtering options specify conditions that a chunklet should meet in order for the chunklet to be displayed.

Multiple conditions can be specified and a chunklet will be displayed if any of the specified conditions are met. By default (if no filtering options are specified) only those chunklets that are mapped to an LD are displayed.

`-a`

Specifies that information about all chunklets is displayed.

`-fail`

Specifies that information about failed chunklets(media-failed chunklets, disk-failed chunklets, or chunklets marked as failed by the operating system) is displayed.

`-mov`

Specifies that information about chunklets that have moved, are scheduled to move, or are moving, is displayed.

`-from <pdid,...>`

Specifies chunklets that have moved or are to be moved from the matching PD IDs are displayed.

`-cln`

Specifies that information for clean chunklets is displayed.

`-cng`

Specifies that information for chunklets that are being cleaned by the system, is displayed.

`-tgt`

Specifies that information for chunklets marked as targets of relocation is displayed.

`-src`

Specifies that information about chunklets that are marked as sources of relocation, is displayed.

`-spr`

Specifies that information about chunklets that are marked as spares, is displayed.

`-log`

Specifies that information about chunklets that are logging is displayed.

`-sync`

Specifies that information about chunklets that are synchronizing with their RAID sets is displayed.

`-sortcol <col>[,<dir>][:<col>[,<dir>]...]`

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

The following option is used to filter the disks on which the chunklets are to be displayed.

-p <pattern>

Specifies a pattern for disks whose chunklets are to be displayed.

If no pattern is specified, all disks are included.

If specified multiple times, each instance of the specified pattern adds additional candidate disks matching the pattern.

The following arguments can be specified as patterns for this option:

An item is specified as an integer, a comma-separated list of integers, or a range of integers specified from low to high.

-nd <item>

Specifies one or more nodes. Nodes are identified by one or more integers (item). Multiple nodes are separated with a single comma (e.g. 1,2,3). A range of nodes is separated with a hyphen (e.g. 0-7). The primary path of the disks must be on the specified node(s).

-st <item>

Specifies one or more PCI slots. Slots are identified by one or more integers (item). Multiple slots are separated with a single comma (e.g. 1,2,3). A range of slots is separated with a hyphen (e.g. 0-7). The primary path of the disks must be on the specified PCI slot(s).

-pt <item>

Specifies one or more ports. Ports are identified by one or more integers (item). Multiple ports are separated with a single comma (e.g. 1,2,3). A range of ports is separated with a hyphen (e.g. 0-4). The primary path of the disks must be on the specified port(s).

-cg <item>

Specifies one or more drive cages. Drive cages are identified by one or more integers (item). Multiple drive cages are separated with a single comma (e.g. 1,2,3). A range of drive cages is separated with a hyphen (e.g. 0-3). The specified drive cage(s) must contain disks.

-mg <item>

Specifies one or more drive magazines. The "1." or "0." displayed in the CagePos column of showpd output indicating the side of the cage is omitted when using the -mg option. Drive magazines are identified by one or more integers (item). Multiple drive magazines are separated with a single comma (e.g. 1,2,3). A range of drive magazines is separated with a hyphen (e.g. 0-7). The specified drive magazine(s) must contain disks.

-pn <item>

Specifies one or more disk positions within a drive magazine. Disk positions are identified by one or more integers (item). Multiple disk positions are separated with a single comma (e.g. 1,2,3). A range of disk positions is separated with a hyphen(e.g. 0-3). The specified position(s) must contain disks.

-dk <item>

Specifies one or more physical disks. Disks are identified by one or more integers (item). Multiple disks are separated with a single comma (e.g. 1,2,3). A range of disks is separated with a hyphen (e.g. 0-3). Disks must match the specified ID(s).

-tc_gt <number>

Specifies that physical disks with total chunklets greater than the number specified be selected.

-tc_lt <number>

Specifies that physical disks with total chunklets less than the number specified be selected.

-fc_gt <number>

Specifies that physical disks with free chunklets greater than the

number specified be selected.

-fc_lt <number>

Specifies that physical disks with free chunklets less than the number specified be selected.

-devid <model>

Specifies that physical disks identified by their models be selected. Models can be specified in a comma-separated list. Models can be displayed by issuing the "`showpd -i`" command.

-devtype <type>

Specifies that physical disks must have the specified device type (FC for Fast Class, NL for Nearline, SSD for Solid State Drive, SCM for Storage Class Memory) to be used. Device types can be displayed by issuing the "`showpd`" command.

-rpm <number>

Drives must be of the specified relative performance metric, as shown in the `RPM` column of the `showpd` command. For FC and NL drives, the number corresponds to both a performance measure and actual rotational speed. For SSD & SCM drives, this field is "N/A".

Disks that satisfy all of the specified characteristics are used.

For example -p -fc_gt 60 -fc_lt 230 -nd 2 specifies all the disks that have greater than 60 and less than 230 free chunklets and that are connected to node 2 through their primary path.

Specifiers

<pdid>|<pattern>... (deprecated usage)

Specifies the PD IDs or the PD ID pattern for disks whose chunklets are to be displayed.

Authority

Any role in the system

Usage

If no option is specified, the `showpdch` command defaults to displaying information about chunklets used for logical disks.

Examples

The following example displays information about chunklets on PD 1:

```
cli% showpdch 1
Pdid Chnk     LdName LdCh   State  Usage  Media Sp Cl From  To
   1    0   vvfromcpg2.usr.1    0  normal     ld  valid  N  N  --- ---
   1    1   vvfromcpg2.usr.1    4  normal     ld  valid  N  N  --- ---
   1    2   vvfromcpg2.usr.1    8  normal     ld  valid  N  N  --- ---
-------------------------------------------------------------------------
Total chunklets: 3
```

Displaying chunklets currently on physical disk 1 that have moved from other disks:

```
cli%showpdch -mov -p -dk 1
```

Displaying chunklets on all physical disks that were moved from physical disk 4 and 5:

```
cli%showpdch -from 4,5
```

Displaying spare chunklets on physical disks on cage 3:

```
cli%showpdch -spr -p -cg 3
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showpdch
```

showpdvv

Syntax

```
showpdvv [options] [<PD_ID>[:<chunklet>]...]
```

Description

The showpdvv command displays the virtual volumes that are mapped to a particular physical disk.

Options

-sum

Shows number of chunklets used by virtual volumes for different space types for each physical disk.

-p <pattern>

Specifies a pattern to select <PD_ID> disks.

The following arguments can be specified as patterns for this option:

An item is specified as an integer, a comma-separated list of integers, or a range of integers specified from low to high.

-nd <item>

Specifies one or more nodes. Nodes are identified by one or more integers (item). Multiple nodes are separated with a single comma (e.g. 1,2,3). A range of nodes is separated with a hyphen (e.g. 0-7). The primary path of the disks must be on the specified node(s).

-st <item>

Specifies one or more PCI slots. Slots are identified by one or more integers (item). Multiple slots are separated with a single comma (e.g. 1,2,3). A range of slots is separated with a hyphen (e.g. 0-7). The primary path of the disks must be on the specified PCI slot(s).

-pt <item>

Specifies one or more ports. Ports are identified by one or more integers (item). Multiple ports are separated with a single comma (e.g. 1,2,3). A range of ports is separated with a hyphen (e.g. 0-4). The primary path of the disks must be on the specified port(s).

-cg <item>

Specifies one or more drive cages. Drive cages are identified by one or more integers (item). Multiple drive cages are separated with a single comma (e.g. 1,2,3). A range of drive cages is separated with a hyphen (e.g. 0-3). The specified drive cage(s) must contain disks.

-mg <item>

Specifies one or more drive magazines. The "1." or "0." displayed in the CagePos column of showpd output indicating the side of the cage is omitted when using the -mg option. Drive magazines are identified by one or more integers (item). Multiple drive magazines are separated with a single comma (e.g. 1,2,3). A range of drive magazines is separated with a hyphen (e.g. 0-7). The specified drive magazine(s) must contain disks.

-pn <item>

Specifies one or more disk positions within a drive magazine. Disk positions are identified by one or more integers (item). Multiple disk positions are separated with a single comma (e.g. 1,2,3). A range of disk positions is separated with a hyphen (e.g. 0-3). The specified position(s) must contain disks.

-dk <item>

Specifies one or more physical disks. Disks are identified by one or more integers (item). Multiple disks are separated with a single comma (e.g. 1,2,3). A range of disks is separated with a hyphen (e.g. 0-3). Disks must match the specified ID(s).

-tc_gt <number>

Specifies that physical disks with total chunklets greater than the number specified be selected.

-tc_lt <number>

Specifies that physical disks with total chunklets less than the number specified be selected.

-fc_gt <number>

Specifies that physical disks with free chunklets greater than the

number specified be selected.

-fc_lt <number>

Specifies that physical disks with free chunklets less than the number specified be selected.

-devid <model>

Specifies that physical disks identified by their models be selected. Models can be specified in a comma-separated list. Models can be displayed by issuing the "`showpd -i`" command.

-devtype <type>

Specifies that physical disks must have the specified device type (FC for Fast Class, NL for Nearline, SSD for Solid State Drive) to be used. Device types can be displayed by issuing the "`showpd`" command.

-rpm <number>

Drives must be of the specified relative performance metric, as shown in the "RPM" column of the "`showpd`" command. For FC and NL drives, the number corresponds to both a performance measure and actual rotational speed. For SSD & SCM drives, this field is "N/A".

Disks that satisfy all of the specified characteristics are used.

For example -p -fc_gt 60 -fc_lt 230 -nd 2 specifies all the disks that have greater than 60 and less than 230 free chunklets and that are connected to node 2 through their primary path.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

<PD_ID>

Specifies the physical disk ID using an integer. This specifier is not required if -p option is used, otherwise it must be used at least once on the command line.

[:<chunklet>]

Specifies the chunklet number in a physical disk to which virtual volumes are mapped. This specifier is not required. This specifier cannot be used along with -p option.

Authority

Any role in the system

Usage

None.

Examples

The following example displays that the user space virtual volumes v0 and v1, respectively, are mapped to physical disk 44:

```
cli% showpdvv 44
PDId CagePos Type      RPM VVId VVName VVSp
  44 2:3:2   FC 10  109 v0     usr
  44 2:3:2   FC 10  110 v1     usr
```

The following example displays the summary output for volumes mapped to physical disk 55:

```
cli% showpdvv -sum 55
                                    ----Chunklets----
PDId CagePos Type      RPM VVId VVName Adm Snp Usr Total
  55 4:9:0   FC          15    2 v0       0   0   2     2
  55 4:9:0   FC          15    3 v1       0   0  20    20
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showpdvv
```

showport

Syntax

```
showport [options] [-failed] [-sortcol <col>[,<dir>]] [<node:slot:port> |
<pattern>...]
```

Description

The `showport` command displays information about ports in the system.

Options

-i

Shows port hardware inventory information.

-c

Displays all devices connected to the port. Such devices include cages (for initiator ports), hosts (for target ports) and ports from other storage system (for RCFC and peer ports).

-par

Displays a parameter listing such as the configured data rate of a port and the maximum data rate that the card supports. Also shown is the type of attachment (Direct Connect or Fabric Attached) and whether the unique_nwwn and VCN capabilities are enabled.

-rc

Displays information that is specific to the Remote Copy ports.

-rcfc

Displays information that is specific to the Fibre Channel Remote Copy ports.

-peer

Displays information that is specific to the Fibre Channel ports for Data Migration.

-rcip

Displays information specific to the Ethernet Remote Copy ports.

-iscsi

Displays information about iSCSI ports.

-iscsiname

Displays iSCSI names associated with iSCSI ports.

-iscsivlans

Displays information about VLANs on iSCSI ports.

-sfp

Displays information about the SFPs attached to ports.

-ddm

Displays Digital Diagnostics Monitoring (DDM) readings from the SFPs if they support DDM. This option must be used with the `-sfp` option.

-d

Displays detailed information about the SFPs attached to ports. This option is used with the `-sfp` option.

-failed

Shows only failed ports.

-state

Displays the detailed state information. This is the same as `-s`.

-s

Displays the detailed state information. This option is deprecated and will be removed in a subsequent release. Use `-state`.

-ids

Displays the identities hosted by each physical port.

-fs

Displays information specific to the Ethernet File Persona ports.

To see IP address, netmask and gateway information on File Persona, run "`showfs -net`".

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

<node:slot:port>...

Requests that information for a specified port is displayed. This specifier can be repeated to display configuration information about multiple ports. If not specified, configuration information for all ports in the system is displayed.

node—Specifies the node using a number from 0 through 7.

slot—Specifies the PCI slot in the specified node. Valid range is 0 through 9.

port—Specifies the port using a number from 1 through 4.

<pattern>...

Specifies that the port matching the specified glob-style pattern is displayed. This specifier can be repeated to display configuration information about multiple ports. If not specified, configuration information for all ports in the system is displayed.

Authority

Any role in the system

Usage

None.

Examples

The following example displays information about all ports in the system:

```
cli% showport
N:S:P      Mode     State ----Node_WWN---- -Port_WWN/HW_Addr- Type Protocol Label Partner FailoverState
0:0:1 initiator     ready 50002ACFF707E87D   50002AC00107E87D disk      SAS     -       -             -
0:0:2 initiator     ready 50002ACFF707E87D   50002AC00207E87D disk      SAS     -       -             -
0:0:3 initiator     ready 50002ACFF707E87D   50002AC00307E87D disk      SAS     -       -             -
0:0:4 initiator     ready 50002ACFF707E87D   50002AC00407E87D disk      SAS     -       -             -
0:1:1      peer   offline                -       20677C610291 free       IP      -       -             -
0:2:1      peer   offline                -       20677C610292 free       IP      -       -             -
0:3:1    target loss_sync 2FF70002AC07E87D   20310002AC07E87D free       FC      -       -             -
0:3:2    target loss_sync 2FF70002AC07E87D   20320002AC07E87D free       FC      -       -             -
0:3:3    target     ready 2FF70002AC07E87D   20330002AC07E87D host       FC      -       -             -
0:3:4    target     ready 2FF70002AC07E87D   20340002AC07E87D host       FC      -       -             -
0:4:1    target loss_sync 2FF70002AC07E87D   20410002AC07E87D free       FC      -       -             -
0:4:2    target loss_sync 2FF70002AC07E87D   20420002AC07E87D free       FC      -       -             -
0:4:3    target loss_sync 2FF70002AC07E87D   20430002AC07E87D free       FC      -       -             -
0:4:4    target loss_sync 2FF70002AC07E87D   20440002AC07E87D free       FC      -       -             -
0:5:1    target   offline                -       1402ECD2D016 free    iSCSI      -   1:5:1          none
0:5:2    target   offline                -       1402ECD2D017 free    iSCSI      -   1:5:2          none
0:5:3    target     ready                -       1402ECD2D028 host    iSCSI      -   1:5:3          none
0:5:4    target     ready                -       1402ECD2D029 host    iSCSI      -   1:5:4          none
...
```

The following example displays hardware and connection information about the ports settings:

```
cli% showport -i
N:S:P Brand  Model    Rev Firmware    Serial          HWType
0:0:1 LSI    9205-8e 01  17.11.00.00 SP12430085      SAS
0:0:2 LSI    9205-8e 01  17.11.00.00 SP12430085      SAS
0:1:1 QLOGIC QLE2672 02  8.1.1       RFE1228G50820   FC
```

```
0:1:2 QLOGIC QLE2672 02  8.1.1       RFE1228G50820  FC
0:2:1 QLOGIC QLE8242 58  4.15.2      PCGLTX0RC1G3PX CNA
0:2:2 QLOGIC QLE8242 58  4.15.2      PCGLTX0RC1G3PX CNA
```

The following example displays devices connected to ports and information about the ports settings on node 0 slot 4:

```
cli% showport -par 0:4:*
N:S:P Connmode ConnType CfgRate MaxRate Class2   UniqNodeWwn VCN      IntCoal TMWO
0:4:1 host     point    auto    4Gbps   disabled disabled    disabled enabled enabled
0:4:2 host     point    auto    4Gbps   disabled disabled    disabled enabled enabled
0:4:3 disk     loop     auto    4Gbps   disabled disabled    disabled enabled n/a
0:4:4 disk     loop     auto    4Gbps   disabled disabled    disabled enabled n/a
-------------------------------------------------------------------------------
4
```

The following example displays each system ports' configuration:

```
cli% showport -par
N:S:P Connmode ConnType CfgRate MaxRate Class2   UniqNodeWwn VCN      IntCoal TMWO     Smart_SAN
0:0:1 disk     point    6Gbps   6Gbps   n/a      n/a         n/a      enabled n/a      n/a
0:0:2 disk     point    6Gbps   6Gbps   n/a      n/a         n/a      enabled n/a      n/a
0:1:1 host     point    auto    16Gbps  disabled disabled    disabled enabled disabled n/a
0:1:2 host     point    auto    16Gbps  disabled disabled    disabled enabled disabled n/a
0:7:1 host     point    auto    16Gbps  disabled disabled    disabled enabled disabled n/a
0:7:2 host     point    auto    16Gbps  disabled disabled    disabled enabled disabled n/a
```

The following example displays information about all Remote Copy ports:

```
cli% showport -rc
N:S:P State            HwAddr    Rate Type
5:3:2 ready 25320002AC000006   2Gbps rcip
4:1:1 ready 24110002AC000006   2Gbps rcip
0:5:1 ready     0002B39B2013 100Mbps rcip
```

The following example displays the detailed state of the port:

```
cli% showport -state
N:S:P --State-- -Detailed_State-
0:0:1 loss_sync loss_sync
0:0:2 loss_sync loss_sync
0:3:1 offline   offline
0:3:1 offline   offline
0:5:1 loss_sync loss_sync
--------------------------------
10
```

The following example displays information about RCIP ports:

```
cli% showport -rcip
N:S:P State ---HwAddr--- IPAddr Netmask Gateway MTU Rate Duplex AutoNeg
0:1:1 offline 000423C21B72 - - - - n/a n/a n/a
0:1:1 offline 000423C21B73 - - - - n/a n/a n/a
1:2:1 offline 000423ADE95E - - - - n/a n/a n/a
1:2:1 offline 000423ADE95F - - - - n/a n/a n/a
```

The following example displays information about iSCSI ports:

```
cli% showport -iscsi
N:S:P State IPAddr             Netmask/PrefixLen Gateway TPGT  MTU   Rate iSNS_Addr iSNS_Port STGT VLAN
0:2:1 ready 1df9:7b7b:790::21  64                ::        21 1500 10Gbps ::             3205   21 Y
0:2:2 ready 10.99.1.3          255.255.255.0     0.0.0.0   22 1500 10Gbps 0.0.0.0        3205   22 Y
1:2:1 ready 1df9:7b7b:790::121 64                ::       121 1500 10Gbps ::             3205  121 Y
1:2:2 ready 10.99.1.4          255.255.255.0     0.0.0.0  122 1500 10Gbps 0.0.0.0        3205  122 Y
-------------------------------------------------------------------------------------------------------
4
```

The following example displays information about iSCSI names associated with iSCSI ports:

```
cli% showport -iscsiname
N:S:P IPAddr ---------------iSCSI_Name----------------
1:3:1 0.0.0.0 iqn.2000-05.com.3pardata:21310002ac00000a
1:3:2 0.0.0.0 iqn.2000-05.com.3pardata:21320002ac00000a
```

The following example displays information about VLANs on iSCSI ports:

```
cli% showport -iscsivlans
N:S:P VLAN IPAddr                Netmask/PrefixLen Gateway  MTU TPGT STGT iSNS_Addr iSNS_Port
0:6:1    - 10.100.31.61          255.255.255.0     0.0.0.0 1500   61   61 0.0.0.0        3205
0:6:2   12 10.100.12.62          255.255.255.0     0.0.0.0 1500 1025 1025 0.0.0.0        3205
1:6:1    - 10.100.31.161         255.255.255.0     0.0.0.0 1500  161  161 0.0.0.0        3205
1:6:2   12 10.100.12.162         255.255.255.0     0.0.0.0 1500 1024 1024 0.0.0.0        3205
2:6:1   13 1df9:7b7b:53b3::13:261 64                ::      1500 1027 1027 ::             3205
2:8:2    - 10.100.31.82          255.255.255.0     0.0.0.0 1500  282  282 0.0.0.0        3205
3:6:1   13 1df9:7b7b:53b3::13:361 64                ::      1500 1026 1026 ::             3205
3:8:2    - 10.100.31.83          255.255.255.0     0.0.0.0 1500  382  382 0.0.0.0        3205
-------------------------------------------------------------------------------------------------
8
```

The following example displays all ports that are configured as peer:

```
cli% showport -peer
N:S:P Mode State ----Node_WWN---- ----Port_WWN---- Rate
0:5:1 initiator ready 2FF70202AC000007 20510202AC000007 2Gbps
```

The following examples display both standard and detailed information about SFPs attached to ports:

```
cli% showport -sfp
N:S:P -State- -Manufacturer- MaxSpeed(Gbps) TXDisable TXFault RXLoss DDM
0:0:1 OK      FINISAR_CORP.           2.10 No         No      No     Yes
0:0:2 OK      FINISAR_CORP.           2.10 No         No      Yes    Yes
0:3:2 OK      SIGMA-LINKS             2.10 No         No      Yes    Yes
1:0:1 OK      FINISAR_CORP.           2.10 No         No      No     Yes
1:0:2 OK      FINISAR_CORP.           2.10 No         No      Yes    Yes

cli% showport -sfp -d
------------Port 0:0:2------------
N:S:P          :   0:0:2
State          :   Degraded
Manufacturer   :   PICOLIGHT
Part Number    :   PL-XPL-VE-S24-31
Serial Number  :   425EF1E6
Revision       :   N/A
MaxSpeed(Gbps):   2.10
Qualified      :   No
TX Disable     :   --
TX Fault       :   --
RX Loss        :   --
RX Power Low   :   No
DDM Support    :   No
----------Port 1:3:1-----------
N:S:P          :   1:3:1
State          :   OK
Manufacturer   :   FINISAR_CORP.
Part Number    :   FTLF8519P2BNL
Serial Number  :   U76031S
Revision       :   A
MaxSpeed(Gbps):   2.10
Qualified      :   Yes
TX Disable     :   No
TX Fault       :   No
RX Loss        :   No
RX Power Low   :   No
DDM Support    :   Yes
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showport
```

showportarp

Syntax

```
showportarp [<node>:<slot>:<port>]
```

Description

The `showportarp` command shows the ARP table for iSCSI ports in the system.

Options

None.

Specifiers

`[<node>:<slot>:<port>]`

Specifies the port for which information about devices on that port are displayed.

`node` Specifies the node.

`slot` Specifies the PCI bus slot in the specified node.

`port` Specifies the iSCSI port number of the PCI card in the specified PCI bus slot. If <N>:<S>:<P> is not specified, the ARP table for all iSCSI ports is displayed.

If <node>:<slot>:<port> is not specified, the ARP table for all iSCSI ports is displayed.

Authority

Any role in the system

Usage

None.

Examples

The following example displays the ARP table for the iSCSI ports in the system.

```
cli% showportarp
N:S:P        HwAddr    IPAddr        Time (mins)
1:3:1 00E07BF8BF87     192.168.9.1           0
1:3:1 001143CD039A     192.168.8.151         0
```

In the example output above:

- HwAddr. The MAC address of a remote host discovered through the MAC address resolution process.
- IPAddr. The IP address of the remote host to which the port is attempting to connect.
- Time (mins). The amount of time (in minutes) that the entry has been in the table. When the entry has been in the table for 20 minutes, it is removed.

Use the following to get a description of each column in the above examples:

```
clihelp -col showportarp
```

showportdev

Syntax

```
showportdev loop|all|ns|fcf|sas|fcswitch|fcfabric <node:slot:port>
```

```
showportdev sas [options] <node:slot:port>
showportdev findport <WWN> <node:slot:port>
showportdev tzone [[options] <node:slot:port>]
showportdev uns [options] <node:slot:port>
showportdev lldp [options] [<node:slot:port>]
showportdev dcbx [options] [<node:slot:port>]
```

Description

The `showportdev` command displays detailed information about devices on a specified port.

Options

`-pel`

Includes the SAS Phy Error Log (PEL) data for each phy in the SAS topology. This option is only valid when using the `sas` subcommand.

`-d`

Includes detailed initiator information: HBA Manufacturer, HBA Model, HBA Firmware Version, HBA OS Name/ Version, the HBA port's supported and current speeds, HBA port's OS device name, hostname, alias name(s), and whether the Smart SAN QoS and Security features are supported when used with the `tzone` or `uns` subcommand.

When used with the `lldp` or `dcbx` subcommand, this option will return relevant detailed information on the LLDP and DCBX information received from the peer device. This option is only valid when using either the `tzone`, `uns`, `lldp` or `dcbx` subcommand.

`-app`

Includes detailed information provided from the DCBX Application Protocol TLV configured on the peer device.

`-pfc`

Includes detailed information from the DCBX Priority Flow Control TLV configured on the peer device.

`-pg`

Includes detailed information from the DCBX Priority Groups TLV configured on the peer device.

Specifiers

`[<node:slot:port>]`

Specifies the port for which information about devices on that port are displayed.

`node`—Specifies the node.

`slot`—Specifies the PCI bus slot in the specified node.

`port`—Specifies the Fibre Channel port number of the PCI card in the specified PCI bus slot.

`[<WWN>]`

Specifies the Fibre Channel worldwide port name of an attached port.

Subcommands

`loop`

Specifies that information is returned for arbitrated loop devices that are attached to the specified port. This subcommand is only for use with Fibre Channel arbitrated loop ports.

`all`

Specifies that information for all devices attached to the specified port is returned.

ns

Specifies that information for the switch name server database is returned. This subcommand is only for use with fabric-attached topologies.

fcf

Specifies that information for all Fibre Channel over Ethernet forwarders (FCFs) known to the specified port is returned. This subcommand is for use only with Fibre Channel over Ethernet (FCoE) ports.

sas

Specifies that information for all devices in the SAS topology attached to the specified port is returned. This subcommand is only for use with SAS ports.

fcswitch

Specifies that a list of all switches in the Fibre Channel fabric is returned. This subcommand is only for use with fabric-attached Fibre Channel ports.

fcfabric

Specifies that a description of the Fibre Channel fabric is returned. This subcommand is only for use with fabric-attached Fibre Channel ports.

findport

Searches the Fibre Channel fabric attached to the specified port for information on the supplied WWN. Supplying the term "this" in place of a WWN indicates that the port WWN of the specified HPE 3PAR storage system host port should be used. This subcommand is only for use with fabric-attached Fibre Channel ports.

tzone

Without the `<node:slot:port>`, this command will return a list of all the current target-driven zones for any port. If `<node:slot:port>` is provided, then detailed information about the target-driven zone for this port will be provided. This command is only used with fabric-attached Fibre Channel ports.

uns

Specifies that information for all initiators from the switch unzoned name server database is returned. This subcommand is only for use with fabric-attached topologies.

lldp

Specifies available Link Layer Discovery Protocol information for each iSCSI port physically connected is returned. If `<node:slot:port>` is provided, then only information for this port will be displayed. This subcommand is only used with iSCSI QLogic 83XX series ports.

dcbx

Specifies available Data Center Bridging Exchange Protocol information for each iSCSI port physically connected is returned. If `<node:slot:port>` is provided, then only information for this port will be displayed. This subcommand is only used with iSCSI QLogic 83XX series ports.

Authority

Any role in the system

Usage

The `loop` subcommand is functional only in a private loop topology.

Examples

The following example displays information about all devices attached to Fibre Channel port 1:0:2:

```
cli% showportdev all 1:0:2
PtId LpID Hadr Node_WWN          Port_WWN          ftrs   svpm   bbct flen   Name
0xd3 0x0c 0x00 2FF70002AC000013 21020002AC000013 0x8800 0x0022 n/a  0x0800 1:0:2
```

```
0xef 0x00 0xef 2000000087002078 2200000087002078 0x8800 0x0012 n/a  0x0800 pd7
0xe8 0x01 0xe8 2000000087002515 2200000087002515 0x8800 0x0012 n/a  0x0800 pd6
...
```

The following example displays the SAS topology for port 1:0:2:

```
cli% showportdev sas 1:0:2
```

The following example displays the SAS topology with PEL data for port 1:0:2:

```
cli% showportdev sas -pel 1:0:2
```

The following example displays the attached Fibre Channel fabric for port 0:2:3:

```
cli% showportdev findport 21002C27D752ABC7 0:2:3
cli% showportdev findport this 0:2:3
```

The following example searches the attached Fibre Channel fabric for the port WWN of port 0:2:3:

```
cli% showportdev findport this 0:2:3
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showportdev
```

showportisns

Syntax

```
showportisns [<node>:<slot>:<port>]
```

Description

The `showportisns` command shows iSNS host information for iSCSI ports in the system.

Authority

Any role in the system

Options

None.

Specifiers

`[<node:slot:port>]`

Specifies the port for which information about devices on that port are displayed.

`node`—Specifies the node.

`slot`—Specifies the PCI bus slot in the specified node.

`port`—Specifies the Fibre Channel port number of the PCI card in the specified PCI bus slot.

If this specifier is not specified, iSNS host information for all iSCSI ports is displayed.

Usage

None.

Examples

The following example displays hosts discovered by the iSCSI port from the iSNS server in the system.

```
cli% showportisns
N:S:P Host_IPAddr ----------Host_iSCSI_Name---------- ------Host_alias------- VLAN
0:2:2 10.99.1.12  iqn.1991-05.com.microsoft:dl360g7-409 <MS SW iSCSI Initiator>    -
```

```
1:2:2 10.99.1.12  iqn.1991-05.com.microsoft:dl360g7-409 <MS SW iSCSI Initiator>    -
1:2:2 10.99.1.11  iqn.1991-05.com.microsoft:dl360g7-409 <MS SW iSCSI Initiator>    1
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showportisns
```

showportlesb

Syntax

```
showportlesb reset [node:slot:port]
showportlesb compare [all | node:slot:port]
showportlesb single|both [-portwwn <wwn>] node:slot:port
showportlesb hist [options] node:slot:port
showportlesb diffhist [options] node:slot:port
```

Description

The `showportlesb` command displays Fibre Channel Link Error Status Block (LESB) counters (the number of errors accumulated for Fibre Channel devices). The LESB is composed of six counters that can measure Fibre Channel signal integrity or status.

Options

`-startt <time>`

Specifies that samples taken of LESB counters should commence after the indicated time (<time>). Time can be specified in hours or as a specific date. When specifying the time in hours, the following formats can be used:

- hh[:mm[:ss]], where hh is the hour, mm is the minute (optional), and ss is the second (optional).
- hhmm, where hh is interpreted as 24 hour clock.

When specifying the time as a date, the following formats can be used:

- mm/dd[/yy], where mm is the month, dd is the day, and yy is the year (optional).
- monthname dd[,yy], where dd is the day and yy is the year (optional).
- dd monthname [yy], where dd is the day and yy is the year (optional).
- yy-mm-dd, where yy is the year, mm is the month, and dd is the day.

`-endt <time>`

Specifies that samples taken of LESB counters cease after the indicated time (<time>).

<time> can be specified in hours or as a specific date.

When specifying in hours, the following can be used:

- hh[:mm[:ss]], where hh is the hour, mm is the minute (optional), and ss is the second (optional).
- hhmm, where hh is interpreted as a 24 hour clock.

When specifying the time as a date, the following formats can be used:

- mm/dd[/yy], where mm is the month, dd is the day, and yy is the year
- (optional).
- monthname dd[,yy], where dd is the day and yy is the year (optional).

- dd monthname [yy], where dd is the day and yy is the year (optional).
- yy-mm-dd, where yy is the year, mm is the month, and dd is the day.

-portwwn <wwn>

Specifies the port WWN of an attached device from which LESB counters should be retrieved.

Specifiers

[<node:slot:port>]

Specifies the port for which information about devices on that port are displayed.

`node` Specifies the node.

`slot` Specifies the PCI bus slot in the specified node.

`port` Specifies the Fibre Channel port number of the PCI card in the specified PCI bus slot.

This specifier is required for the single, both, hist, and diffhist subcommands, and optional for the compare subcommand. If this specifier is not used with the compare subcommand, then all ports are compared.

<wwn>

Specifies the port WWN of an attached device. This specifier may be used with the single and both subcommands.

Subcommands

reset

Specifies that internal counters are checked against current LESB counters and event alerts and are raised as necessary. All ports for the internal counters are reset.

compare

Specifies that internal HPE 3PAR counters are checked against current LESB counters.

single|both

The single subcommand specifies that counters for the indicated port are displayed. The both subcommand specifies that counters for both ports, if the device is dual ported, are displayed. If the port WWN of an attached device is specified using the -portwwn option, only counters for that device are displayed.

hist

Displays the history of LESB raw counters on the specified port of loop devices.

diffhist

Displays differences between historical samples of LESB counters on the specified port of loop devices.

Authority

Any role in the system

Usage

- If the both subcommand is specified, the resulting output text is greater than 80 columns wide.
- Internal counters are checked every 10 minutes.
- Each LESB counter is a 32-bit, unsigned integer.
- LESB counters on Fibre Channel devices cannot be reset.

Examples

The following example displays the comparison of all ports:

```
cli% showportlesb compare
Port <1:0:1>

Loop <1:0:1> Time since last save: 0:00:28
     ID ALPA LinkFail LossSync LossSig PrimSeq InvWord InvCRC
<1:0:1> 0xef        1       19      19       0       0      0
pd7     0x6d        1        5       0       0     270      0
pd6     0x72        1        4       0       0     524      0
pd5     0x73        1        4       0       0     335      0
pd4     0x76        1        4       0       0     334      0
pd3     0x79        1        4       0       0     401      0
pd2     0x80        1        4       0       0     344      0
pd1     0x81        1        4       0       0     270      0
pd0     0x88        1        4       0       0     401      0
Port <0:0:1>

Loop <0:0:1> Time since last save: 0:00:28
     ID ALPA LinkFail LossSync LossSig PrimSeq InvWord InvCRC
<0:0:1> 0xef        1       19      19       0       0      0
pd7     0x6d        1        5       0       0     465      0
pd6     0x72        1        5       0       0     890      0
pd5     0x73        1        4       0       0     969      0
pd4     0x76        1        5       0       0     761      0
pd3     0x79        1        4       0       0     815      0
pd2     0x80        1        4       0       0     925      0
pd1     0x81        1        7       0       0    3283      0
pd0     0x88        1      258       0       0     269      0

LESB compare completed
```

The following example displays the counters for port 1:0:2:

```
cli% showportlesb single 1:0:2
      ID  ALPA LinkFail LossSync LossSig PrimSeq InvWord InvCRC
cage1      0x1        3        4       0       0     755      0
pd12      0xef     1245    39201       0       0  156804      0
pd13      0xe8        1     1608       0       0    6432      0
pd14      0xe4        1     1586       0       0    6344      0
pd15      0xe2        1     1588       0       0    6352      0
pd16      0xe1        2     5088       0       0   20352      0
pd17      0xe0        1     1596       0       0    6384      0
pd18      0xdc        1     1595       0       0    6380      0
pd19      0xda        1     1596       0       0    6384      0
pd20      0xd9        2     5047       0       0   20188      0
pd21      0xd6        1     1604       0       0    6416      0
pd22      0xd5        1     1609       0       0    6436      0
pd23      0xd4        1     1616       0       0    6464      0
<1:0:2>   0xd3        0        1       1       0       0      0
cage0     0x18        1      165       0       0      26      0
pd8       0x67        2   865794       0       0 3463176      0
pd9       0x66        1    63932       0       0  255728      0
pd10      0x65        1    61572       0       0  246288      0
pd11      0x63        1    61525       0       0  246100      0
pd4       0x6c        2    67006       0       0  268024      0
pd5       0x6b        1    63474       0       0  253896      0
pd6       0x6a        1    63471       0       0  253884      0
pd7       0x69        1    63598       0       0  254392      0
pd0       0x72        2    65863       0       0  263452      0
pd1       0x71        1    64024       0       0  256096      0
pd2       0x6e        1    63942       0       0  255768      0
pd3       0x6d        1    63897       0       0  255588      0
```

The following example displays the reset of internal counters and ports:

```
cli% showportlesb reset
LESB reset completed
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showportlesb
```

showportpel

Show Phy Error Log (PEL) information about devices on a SAS port.

Syntax

System collected data:

```
showportpel show [options] [all | <n:s:p>]
showportpel hist [options] [-startt] [-endt] [-day | -hour | -min] [all | <n:s:p]
showportpel diffhist [options] [-startt] [-endt] [-day | -hour | -min] [all | n:s:p]
```

User collected data:

```
showportpel save/reset [all| n:s:p]
showportpel showsaved [options] [all| n:s:p]
showportpel compare [options] [all| n:s:p]
```

Common syntax:

Subcommands

show

Displays the latest system collected PEL data.

hist

Displays the history of the PEL data. The default is `-day`.

diffhist

Displays differences between consecutive samples of PEL data.

save/reset

Saves the current hardware PEL data to an internal table.

showsaved

Displays the last user saved PEL data.

compare

Compares current hardware PEL data against last user saved PEL data.

Description

The `showportpel` command displays SAS Phy Error Log (PEL) counters (the number of errors accumulated for SAS devices). The PEL is composed of four counters that can measure SAS signal integrity or status. The `showportpel` command also displays the link rate of each Phy and which device is connected to it. The status column determines a normal or degraded state based on the data above. To calculate the PHY status, compare it against previous PEL Monitor intervals. Address persistent degraded Phys.

Authority

`save/reset`: Any role granted by the `save_pel` right.

All other options available to all roles in the system.

Options

-listcols

List the columns available to be shown in the -showcols option. (see "clihelp -col showportpel" for help on each column).

-showcols <column>[,<column>...]

Explicitly select the columns to be shown using a comma-separated list of column names. Run "showportpel -listcols" to list the available columns. Run "clihelp -col showportpel" for a description of each column.

-detail

Each table row represents a Phy to Phy connection. Without this option, each row represents the device to device connection, emphasizing the field replaceable unit.

-degraded

Shows only degraded Phys.

Hist Options

-startt <time>

Specifies that samples taken of PEL counters should commence after the indicated time (`<time>`).

-endt <time>

Specifies that samples taken of PEL counters cease after the indicated time (`<time>`).

`<time>` can be specified as a relative time, a specific date, or in hours. When specifying a relative time, the following formats can be used:

- `now`: The current date and time.
- `[-/+ n unit]`: You can use minute, hour, or year as the unit. For example, -1 day.

  When specifying the time as a date, the following formats can be used:

  - `monthname dd[,yy]`, where dd is the day and yy is the year (optional).
  - `dd monthname [yy]`, where dd is the day and yy is the year (optional).
  - `yyyy-mm-dd`, where yyyy is the year, mm is the month, and dd is the day.

  When specifying in hours, the following can be used:

  - `hh[:mm[:ss]]`, where hh is the hour, mm is the minute (optional), and ss is the second (optional).
  - `hhmm`, where hh is interpreted as a 24 hour clock.

-day

Displays day to day data for the last year.

-hour

Displays hour to hour data for the last day.

-min

Displays minute to minute data for the last hour in increments of five.

Specifiers

[all]

Specifies that all ports are selected for the subcommand.

[<n:s:p>]

Specifies the port for which information about devices on that port are displayed.

nSpecifies the node.

sSpecifies the PCI bus slot in the specified node.

pSpecifies the SAS port number of the PCI card in the specified PCI bus slot.

Notes

None.

Examples

The following example displays the last system collected PEL data for port 1:0:2:

```
cli% showportpel show 1:0:2
```

The following example displays a history of PEL data for all ports 10 days ago till today. Without the -startt option, it will show a year worth of data:

```
cli% showportpel hist -startt "-10 days" all
```

The following example displays the change in PEL data every 5 minutes for all ports in the last hour:

```
cli% showportpel diffhist -min all
```

showqos

Syntax

```
showqos [options] [{{vvset|domain}:{<name>|<pattern>}|sys:all_others}]...
```

Description

The showqos command lists the QoS rules configured in a system.

Options

-domain {<domainname|pattern>}...

Shows statistics for rules with target objects in domains with names that match one or more of the specified domains or patterns. If -domain is not specified, statistics for all rules with target objects in the current domain are shown. See the currentdomain parameter in the **setclienv** and **showclienv** commands.

{-on|-off}

List QoS rules that are "on" or "off" only.

-vv {<VV_name>|<pattern>}[,{<VV_name>|<pattern>}]...

Applies only to QoS rules whose targets include virtual volumes with names matching any of the names or patterns specified.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

{{vvset|domain}:{<name>|<pattern>}|sys:all_others}

The QoS rules to list. If none is specified, it lists all configured QoS rules. <name> and <pattern> refer to the target object name of the rule, as listed in show{vvset|domain} commands.

Authority

Any role in the system

Usage

- For all the <pattern> fields, the patterns are glob-style (shell-style) patterns. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.)
- A Domain column may be included by using the **setclienv** command to set listdom to 1. The listdom option is also set to 1 if the CLI was started with the -listdom option, or if the CLI was started with the TPDLISTDOM environment variable set. Run "`cli -h`" and "`setclienv -h`" for details of the environment variables.

Examples

Use the following to show all QoS rules defined in the system:

```
cli% showqos
                            -I/O_per_second- -KBytes_per_sec-
Id Type  Name QoS Priority       Min   Max       Min    Max LatencyGoal
42 vvset vsa  off normal           -   300         -      -           -
43 vvset vsb  on  normal           -  4000         - 204800           -
44 vvset vsc  on  normal           -     -         - 102400           -
45 vvset vsd  on  normal           -  5000         - 256000           -
------------------------------------------------------------------------
4 total
```

The following only show QoS rules that are "on":

```
cli% showqos -on
                             -I/O_per_second- -KBytes_per_sec-
Id Type  Name QoS Priority       Min   Max        Min    Max LatencyGoal
43 vvset vsb  on  normal           -  4000          - 204800           -
44 vvset vsc  on  normal           -     -          - 102400           -
45 vvset vsd  on  normal           -  5000          - 256000           -
------------------------------------------------------------------------
3 total
```

Use the following to show QoS rules containing VV vv1:

```
cli% showqos -vv vv1
                       -I/O_per_second- -KBytes_per_sec-
Id Type  Name QoS Priority       Min   Max       Min    Max LatencyGoal
43 vvset vsb  on  normal           -  4000         - 204800           -
45 vvset vsd  on  normal           -  5000         - 256000           -
------------------------------------------------------------------------
2 total
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showqos
```

showrcopy

Syntax

```
showrcopy [options <arg>] [links]
```

```
showrcopy [options <arg>] [groups [<name_or_pattern> ...]]
showrcopy [options <arg>] [targets [<name_or_pattern> ...]]
```

Description

The `showrcopy` command displays details of the Remote Copy configuration.

Options

`-d`

Displays more detailed configuration information.

`-qw`

Displays additional target specific Peer Persistence related configuration where applicable.

`-domain <domainname_or_pattern>[,<domainname_or_pattern>...]`

Shows only Remote Copy groups whose virtual volumes are in domains with names matching one or more of the specified domain name or pattern. This option does not allow listing objects within a domain of which the user is not a member.

Specifiers

`links`

Specifies all Remote Copy links.

`groups [<name_or_pattern>]`

Specifies either all Remote Copy volume groups or specific Remote Copy volume group(s) by name or glob-style patterns.

`targets [<name_or_pattern>]`

Specifies either all target definitions or specific target(s) definition by name or glob-style patterns.

Authority

Any role in the system

Usage

- This command requires the HPE 3PAR Remote Copy license. Contact your local service provider for more information.
- If the `showrcopy` command is used with no specifiers, all configuration information is displayed.
- If `showrcopy link|group|target` is used without specifying a name, information for all links, groups, or targets is displayed.
- (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.)
- To limit the performance impact of remote copy on the rest of the HPE 3PAR storage system, the number of volumes that are concurrently synchronizing is limited. This limit is not user-configurable and applies to the initial synchronization as well as subsequent resynchronizations for synchronous as well as asynchronous periodic groups.
- A Domain column may be included by using the **`setclienv`** command to set listdom to 1. The listdom option is also set to 1 if the CLI was started with the -listdom option, or if the CLI was started with the TPDLISTDOM environment variable set. Run "`cli -h`" and "`setclienv -h`" for details of the environment variables.

Examples

The following example displays output from the `showrcopy` command:

```
cli% showrcopy
```

The following example displays output from the `showrcopy` groups <pat> command, where <pat> is specified as b* and 1*:

```
cli% showrcopy groups b* l*
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showrcopy
```

showrctransport

Syntax

```
showrctransport [options]
```

Description

The `showrctransport` command shows status and information about end-to-end transport for Remote Copy in the system.

Options

`-rcip`

Shows information about Ethernet end-to-end transport.

`-rcfc`

Shows information about Fibre Channel end-to-end transport.

Specifiers

None.

Authority

Any role in the system

Usage

None.

Examples

The following example displays status and information about the end-to-end transport for all Remote Copy configured ports:

```
cli% showrctransport
```

The following example displays status and information about the end-to-end transport for RCIP ports:

```
cli% showrctransport -rcip
```

The following example displays status and information about the end-to-end transport for RCFC ports:

```
cli% showrctransport -rcfc
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showrctransport
```

showrole

Syntax

```
showrole [<role_name | pattern ...>]
showrole -listrights
```

Description

The `showrole` command displays information about rights assigned to roles in the system. Without any argument, `showrole` displays the roles currently defined.

Options

`-listrights`

Lists the rights available in the system

Specifiers

`<role_name | pattern>...`

Specifies that the role name matching the specified glob-style pattern will have their rights displayed. This specifier can be repeated. This specifier is not required.

Authority

Any role in the system

Usage

- To avoid cluttering the output, show commands rights are not included in the output of the `showrole` because show commands are granted to all roles by default.
- If both `-listrights` and `<role_name>` patterns are provided, the `-listrights` option will take precedence and only information related to rights will be displayed.

Examples

The following example displays the role names in the system:

```
cli% showrole
---Role--- ---------------------------------Comments----------------------------------
3PAR_AO    Used internally by HPE for Adaptive Optimization operations.

3PAR_RM    Used internally by HPE for operations required by Recovery Manager.

audit      For security scanners to perform a scan of the HPE 3PAR OS file system. An
           audit user has no access to the CLI.

basic_edit Rights are similar to Edit role, but more restricted, specifically in the
           ability to remove objects such as volumes, VLUNs, and hosts.

browse     Rights are limited to read-only access.

co         Rights to approve the Compliance WORM changes.

create     Rights are limited to creation of objects such as volumes, CPGs, hosts,
           and schedules.

edit       Rights are granted to most operations, such as for creating, editing, and
           removing virtual volumes.

service    Rights are limited to operations required to service the storage server.

super      Rights are granted to all operations.
       cli% showrole browse
---Role--- -------Right------- ------------------------Right_Description------------------------
browse     password_checkown   Check a user's own password
   password_setown     Set a user's own password
   sshkey_remove       Remove users public SSH keys
   sshkey_set  Set the public SSH key for a user
   user_setown Set a user's own properties
   vv_update   Remove old virtual copy and create a new virtual copy of a volume
```

The following example displays the rights assigned to the browse role:

```
cli% showrole browse
---Role--- -------Right------- ------------------------Right_Description-------------------------
browse     password_checkown   Check a user's own password
           password_setown     Set a user's own password
           sshkey_remove       Remove users public SSH keys
           sshkey_set          Set the public SSH key for a user
           user_setown         Set a user's own properties
           vv_update           Remove old virtual copy and create a new virtual copy of a volume
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showrole
```

showrsv

Syntax

```
showrsv [options <arg>] [<VV_name>]
```

Description

The `showrsv` command displays SCSI reservation and registration information for Virtual Logical Unit Numbers (VLUNs) bound for a specific port.

Options

`-l <scsi3|scsi2>`

Specifies that either SCSI-3 persistent reservation or SCSI-2 reservation information is displayed. If this option is not specified, information about both scsi2 and scsi3 reservations will be shown.

`-host <hostname>`

Displays reservation and registration information only for virtual volumes that are visible to the specified host.

Specifiers

`[<VV_name>]`

Specifies the virtual volume name, using up to 31 characters.

Authority

Any role in the system

Usage

None.

Examples

The following example displays information about reservation and registration information for VLUNs bound with host name w2k_emx1_cisco:

```
cli% showrsv -host w2k_emx1_cisco

          VVname              Host              Owner     Port ReservationType
w2k_clustered.10    w2k_emx1_cisco  10000000C92B9909    1:3:1       SCSI-3(6)
w2k_clustered.11    w2k_emx1_cisco  10000000C92B9909    1:3:1       SCSI-3(6)
w2k_clustered.12    w2k_emx1_cisco  10000000C92B9909    1:3:1       SCSI-3(6)
w2k_clustered.13    w2k_emx1_cisco  10000000C92B9909    1:3:1       SCSI-3(6)
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showrsv
```

showsched

Syntax

The syntax for the `showsched` command can be one of the following:

```
showsched [options] [<schedname>|<pattern>]
```

```
showsched [options] -window <start_window_time> <end_window_time>
[<schedname>|<pattern>]
```

Description

The `showsched` command shows the state of tasks currently scheduled on the system.

Options

`-all`

Specifies that all scheduled tasks will be displayed, including system created scheduled tasks.

`-window <start_window_time> <end_window_time>`

Displays all predicted runtimes for all user scheduled tasks for a specified window of time. This displays up to 10,000 predicted runtimes. If schedname or pattern are specified, then the list of up to 10,000 is filtered to only those schedule names that match.

`-sortcol <col>[,<dir>][:<col>[,<dir>]...]`

Sorts command output based on column number (`<col>`). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (`<dir>`) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

`<schedname>|<pattern>...`

Specifies the schedules with the specified name (31 character maximum) or matching the glob-style pattern. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.)

This specifier can be repeated to display configuration information about multiple schedules. This specifier is not required. If not specified, configuration information for all non system scheduled tasks in the system is displayed.

`<start_window_time> <end_window_time>`

Specifies the time window for which to show the predicted runtimes. Start and end window times should each include a date and time. For example, they can be specified as "yyyy-mm-dd hh:mm:ss" where hh is specified as the hour interpreted on a 24 hour clock. Date and time pairs must be surrounded by double quotation marks.

Some flexibility exists around the format of date and time. Acceptable formats for time are hh:mm:ss or hhmm. Acceptable formats for date are mm/dd/yy, month_name dd, dd month_name yy, or yy-mm-dd. If the syntax yy-mm-dd is used, the year must be specified.

Authority

- Super, Service, Edit, Browse
- Any extended role in the system

Usage

- If the `-all` option is not specified then tasks that were created by the system will not be displayed.
- The `-window` option by itself will display only user created scheduled task runtimes. To see all scheduled task runtimes including system tasks, include the `-all` option in the command.

Examples

Use the following to show all the user schedules that have "createsv" in the name:

```
cli% showsched *createsv*
```

Use the following to show all the schedules including those created by the system:

```
cli% showsched -all
                                                    ------ Schedule ------
SchedName          File/Command                     Min Hour DOM Month DOW CreatedBy Status Alert NextRunTime
remove_expired_vvs removevv -f -expired              27    *   *     *    * 3parsvc   active     N 2015-06-19 04:27:00 MDT
move_back_chunklet moverelocpd -f -partial -p -tc_gt 0  17    2   *     *    0 3parsvc   active     N 2015-06-21 02:17:00 MDT
sample             sample 365                        37    3   *     *    * 3parsvc   active     N 2015-06-19 03:37:00 MDT
checkiperfserver   checkstopserver                   15    *   *     *    * 3parsvc   active     N 2015-06-19 04:15:00 MDT
coremanagement     coremanman                        47    4   *     *    * 3parsvc   active     N 2015-06-19 04:47:00 MDT
-------------------------------------------------------------------------------------------------------------------------
5                  total
```

Use the following to show all the runtimes, including schedules created by the system, for a 2 hour window:

```
cli% showsched -all -window "2013-12-19 12:00:00" "2013-12-19 14:00:00"
```

Use the following to show the runtimes, not including schedules created by the system, for a 2 week window and only show schedule names starting with "alpha":

```
cli% showsched -window "2014-02-14 00:00:00" "2014-02-28 00:00:00" alpha*
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showsched
```

showsnmpmgr

Syntax

`showsnmpmgr`

Description

The `showsnmpmgr` command displays a list of registered SNMP managers for receiving traps.

Options

None.

Specifiers

None.

Authority

Any role in the system

Usage

Managers are registered using the **`addsnmpmgr`** command.

Examples

The following example displays a list of registered managers:

```
cli% showsnmpmgr
```

```
HostIP                  Port SNMPVersion User Notify           AlertClear
82.185.98.76            162             2 --   standard    standard
82.162.9.7              8004            2 --   nodup                 all
5def:2008:abcd::161a    9162            3 joe  standard    nodup
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showsnmpmgr
```

showsnmppw

Syntax

```
showsnmppw [options]
```

Description

The `showsnmppw` command displays the SNMP community string passwords.

Options

`-rw|-r|-w`

Specifies that the read-write (-rw), read-only (-r), or write-only (-w) password is displayed. If not specified, the read-write community string password is displayed.

Specifiers

None.

Authority

Any role in the system

Usage

SNMP passwords are registered using the **`setsnmppw`** command.

Examples

The following example displays the read-write community string password:

```
cli% showsnmppw
[password]
```

showsnmpuser

Syntax

```
showsnmpuser [<user_name>]
```

```
showsnmpuser
```

Description

The `showsnmpuser` command displays information about one or all SNMP users, including the username, authentication and privacy protocols.

Options

`-listcols`

Lists the columns available to be shown in the -showcols option described below (see '`clihelp -col showuser`' for help on each column).

By default (if none of the information selection options below are specified) the following columns are shown:

`PrivProtocol AuthProtocol Username-showcols <column>[,<column>...]`

Explicitly select the columns to be shown using a comma-separated list of column names. For this option the full column names are shown in the header.

Run '`showsnmpuser -listcols`' to list the available columns.

Run '`clihelp -col showsnmpuser`' for a description of each column.

Specifiers

`<user_name>`

Specifies the SNMPv3 username, using up to 31 characters. This specifier is not required on the command line. If no specifier is used, information about all SNMP users is displayed.

Authority

Any role in the system

Usage

- The `showsnmpuser` command displays SNMPv3 user information for all local users when the <username> is not provided or the information for just the local user with the name <username> when it is provided. The output for users with Browse or Service role is limited to the user's own accounts.
- If the <username> is provided, it must be the user's name. The output shows the users' username, authentication protocol and privacy protocol.
- SNMP users are created using the **`createsnmpuser`** command.

Examples

```
cli% showsnmpuser joe
Username AuthProtocol PrivProtocol
snmpadm  HMAC-SHA-96  CFB128-AES-128
joe      HMAC-SHA-96  CFB128-AES-128
```

showspace

Syntax

`showspace [options <arg>]`

Description

The `showspace` command displays estimated free space for logical disk creation.

Options

The following options are used to select the logical disk creation parameters used for space calculation:

`-cpg <CPG_name>|<pattern>`

Specifies that logical disk creation parameters are taken from CPGs that match the specified CPG name or pattern, indicated with a glob-style pattern. Multiple CPG names or patterns can be specified using a comma separated list, for example cpg1,cpg2,cpg3. Only the -hist option can be specified with the -cpg option.

`-hist`

Specifies that free space history over time for CPGs specified with the -cpg option is displayed. This option can only be used if the -cpg option is specified.

-t <RAID_type>

Specifies the RAID type of the logical disk: r0 for RAID-0, r1 for RAID-1, r5 for RAID-5, or r6 for RAID-6. If no RAID type is specified, then the default is r6.

-ssz <size_number_chunklet>

Specifies the set size in terms of chunklets. The default depends on the RAID type specified: 2 for RAID-1, 4 for RAID-5, and 8 for RAID-6.

-rs <size>

Specifies the number of sets in a row. The <size> is a positive integer.

If not specified, no row limit is imposed.

-ss <size_KB>

Specifies the step size from 32 KB to 512 KB. The step size should be a power of 2 and a multiple of 32. The default value depends on raid type and device type used. If no value is entered and FC or NL drives are used, the step size defaults to 256 KB for RAID-0 and RAID-1, and 128 KB for RAID-5. If SSD drives are used, the step size defaults to 32 KB for RAID-0 and RAID-1, and 64 KB for RAID-5. For RAID-6, the default is a function of the set size.

-ha port|cage|mag

Specifies that the layout must support the failure of one port pair, one cage, or one drive magazine (mag). This option has no meaning for RAID-0. The default is cage availability.

-ch first|last

Specifies the chunklet location characteristics: either first (attempt to use the lowest numbered available chunklets) or last (attempt to use the highest numbered available chunklets). If no argument is specified, the default characteristic is first.

-p <pattern>

Specifies a pattern for candidate disks. Patterns are used to select disks that are used for creating logical disks. If no pattern is specified, the option defaults to Fast Class (FC) disks. If specified multiple times, each instance of the specified pattern adds additional candidate disks that match the pattern. The -devtype pattern cannot be used to mix Nearline (NL), FC, and Solid State Drive (SSD) drives. An item is specified as an integer, a comma-separated list of integers, or a range of integers specified from low to high.

The following arguments can be specified as patterns for this option:

An item is specified as an integer, a comma-separated list of integers, or a range of integers specified from low to high.

-nd <item>

Specifies one or more nodes. Nodes are identified by one or more integers (item). Multiple nodes are separated with a single comma (e.g. 1,2,3). A range of nodes is separated with a hyphen (e.g. 0-7). The primary path of the disks must be on the specified node(s).

-st <item>

Specifies one or more PCI slots. Slots are identified by one or more integers (item). Multiple slots are separated with a single comma (e.g. 1,2,3). A range of slots is separated with a hyphen (e.g. 0-7). The primary path of the disks must be on the specified PCI slot(s).

-pt <item>

Specifies one or more ports. Ports are identified by one or more integers (item). Multiple ports are separated with a single comma (e.g. 1,2,3). A range of ports is separated with a hyphen (e.g. 0-4). The primary path of the disks must be on the specified port(s).

-cg <item>

Specifies one or more drive cages. Drive cages are identified by one or more integers (item). Multiple drive cages are separated with a single comma (e.g. 1,2,3). A range of drive cages is separated with a hyphen (e.g. 0-3). The specified drive cage(s) must contain disks.

-mg <item>

Specifies one or more drive magazines. The "1." or "0." displayed in the CagePos column of showpd output indicating the side of the cage is omitted when using the -mg option. Drive magazines are identified by one or more integers (item). Multiple drive magazines are separated with a single comma (e.g. 1,2,3). A range of drive magazines is separated with a hyphen(e.g. 0-7). The specified drive magazine(s) must contain disks.

-pn <item>

Specifies one or more disk positions within a drive magazine. Disk positions are identified by one or more integers (item). Multiple disk positions are separated with a single comma (e.g. 1,2,3). A range of disk positions is separated with a hyphen(e.g. 0-3). The specified position(s) must contain disks.

-dk <item>

Specifies one or more physical disks. Disks are identified by one or more integers(item). Multiple disks are separated with a single comma (e.g. 1,2,3). A range of disks is separated with a hyphen(e.g. 0-3). Disks must match the specified ID(s).

-tc_gt <number>

Specifies that physical disks with total chunklets greater than the number specified be selected.

-tc_lt <number>

Specifies that physical disks with total chunklets less than the number specified be selected.

-fc_gt <number>

Specifies that physical disks with free chunklets greater than the number specified be selected.

-fc_lt <number>

Specifies that physical disks with free chunklets less than the number specified be selected.

-devid <model>

Specifies that physical disks identified by their models be selected. Models can be specified in a comma-separated list. Models can be displayed by issuing the "`showpd -i`" command.

-devtype <type>

Specifies that physical disks must have the specified device type (FC for Fast Class, NL for Nearline, SSD for Solid State Drive, SCM for Storage Class Memory) to be used. Device types can be displayed by issuing the `showpd` command.

If it is not specified, the default device type is FC.

-rpm <number>

Disks must be of the specified speed. Device speeds are shown in the RPM column of the `showpd` command. The number does not represent a rotational speed for the drives without spinning media (SSD). For FC and NL drives, the number corresponds to both a performance measure and actual rotational speed. For SSD & SCM drives, this field is "N/A".

Disks that satisfy all of the specified characteristics are used.

For example -p -fc_gt 60 -fc_lt 230 -nd 2 specifies all the disks that have greater than 60 and less than 230 free chunklets and that are connected to node 2 through their primary path.

Specifiers

None.

Authority

Any role in the system

Usage

- The -cpg and -hist options cannot be used with any other option.
- The `showspace` command includes chunklets in the process of being initialized as shown by the "Normal, Unused, Uninit" column of `showpd -c`. The space represented by these chunklets will not actually be available until the initialization process is complete.
- The space calculated is an estimate and not an exact figure.
- For this command MB = 1048576 bytes.

Examples

The following example displays the estimated free space for a RAID-1 logical disk:

```
cli% showspace -t r1
--Estimated(MB)---
RawFree    UsableFree
13824         6912
```

showspare

Syntax

```
showspare [-used]
```

Description

The `showspare` command displays information about chunklets in the system that are reserved for spares and previously free chunklets selected for spares by the system.

Options

`-used`

Shows only used spare chunklets.

By default all spare chunklets are shown.

Specifiers

None.

Authority

Any role in the system

Usage

The **`showpdch`** command is a more general and versatile command that can be used instead of `showspare`.

Examples

The following example displays information about spare chunklets on chunklets 53 through 56 on physical disk 4:

```
cli% showspare
PdId Chnk LdName  LdCh State Usage Media Sp Cl From To
   4   53 ronnie     0 normal   ld valid  N  N 2:37 ---
   4   54 james     28 normal   ld valid  N  N 0:29 ---
   4   55 dio       28 normal   ld valid  N  N 0:32 ---
   4   56 rocks      0 normal   ld valid  N  N 0:38 ---
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showspare
```

showsr

Syntax

```
showsr [options]
```

Description

The `showsr` command displays the amount of space consumed by the various System Reporter databases on the System Reporter volume.

Options

`-ldrg`

Displays which LD region statistic samples are available. This is used with the -btsecs and -etsecs options.

`-btsecs <secs>`

Selects the begin time in seconds for the report. The value can be specified as either-

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

If it is not specified then the time at which the report begins depends on the sample category (-hires, -hourly, -daily):

- For hires, the default begin time is 12 hours ago (-btsecs -12h).
- For hourly, the default begin time is 7 days ago (-btsecs -7d).
- For daily, the default begin time is 90 days ago (-btsecs -90d).

If begin time and sample category are not specified then the time the report begins is 12 hours ago and the default sample category is hires.

If -btsecs 0 is specified then the report begins at the earliest sample.

`-etsecs <secs>`

Selects the end time in seconds for the report. If -attime is specified, select the time for the report. The value can be specified as either of the following:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:

- Full time string including time zone: "2012-10-26 11:00:00 PDT"
- Full time string excluding time zone: "2012-10-26 11:00:00"
- Date string: "2012-10-26" or 2012-10-26
- Time string: "11:00:00" or 11:00:00

- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

If it is not specified then the report ends with the most recent sample.

Authority

Any role in the system

Usage

Each category is reserved a certain amount of space based on its type, so the percentage used is the amount used by the type relative to the amount of space reserved for that type, rather than relative to the entire System Reporter volume.

Examples

The following example displays the System Reporter status:

```
cli% showsr
```

The following example displays available LD region statistics over the last 2 hours:

```
cli% showsr -ldrg -btsecs -2h
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showsr
```

showsralertcrit

Syntax

```
showsralertcrit [options <arg>]
```

Description

The `showsralertcrit` command shows the criteria that System Reporter evaluates to determine if a performance alert should be generated.

Options

`-daily`

Displays only criteria evaluated on a daily basis.

`-hourly`

Displays only criteria evaluated on an hourly basis.

`-hires`

Displays only criteria evaluated on a high-resolution (5 minute) basis.

`-enabled`

Displays only criteria that are enabled.

-disabled

Displays only criteria that are disabled.

-critical

Displays only criteria that have critical severity.

-major

Displays only criteria that have major severity.

-minor

Displays only criteria that have minor severity.

-info

Displays only criteria that have informational severity.

Specifiers

None.

Authority

Any role in the system

Usage

- Both options and conditions are displayed in the Conditions column. The only exception is that frequency options (-daily, -hourly, or -hires) are only displayed under the Freq column.
- By default, all criteria are shown (all frequencies, enabled, disabled and all severities).

Examples

The following example displays all the criteria evaluated on an hourly basis:

```
cli% showsralertcrit -hourly
```

showsshkey

Syntax

`showsshkey`

Description

The `showsshkey` command will display all SSH public keys that have been set with `setsshkey`.

Options

None.

Specifiers

None.

Authority

Any role in the system

Usage

The public SSH keys that are displayed by this command are used to determine that the private keys that are paired with them are valid, and are not useful in determining the contents of that private key. As such, the contents of the public keys are not sensitive, and having access to them will not aid someone attempting to access an account to which they should not have access.

Examples

The following example shows the SSH public keys that have been set by the current user:

```
cli% showsshkey
ID Key
  0 ssh-dss AAAAB3NzaC1kc3MAAACBALDVpIVrNEGJn0tFDCSE1tLXdGVqBkb8Dm/
vziYKXShc6EIeU3kdD7pK9yhTEKx8q9urwgtKwyYH1KMWYp042kuwIM12UzIvkV6DqbA8Z4aOH81W
nBm7pEscDzHdq6IZJ47MszLj1xXO2RYyWTFwsbpKYLXoUrrj8vDqCtcVFGqvAAAAFQDtTx3ZLckWA
I9Pp/jgtJRIZ63wAAAIAtFIyscMhuyma3yMvnE8MyUyP1mJUzpKLHKwCNfvwL49bfXSoczitmY
+ENHTgtWduRaX0uIMmeDeeO5qUUKDfxp5t3KURhl47QcU4lNmLCWvR1S5ep1ff0+D7PVjJtRwlZxL
tfQMyGWzYCMzSRnk41pD6R392A8XLSzvqrLXat3wAAAIEAl+crC528VaZSiSbeW/FybPx3EhvvV/
Zyv6FTFIOE5+bsWbeQZP33M3yuwTvVuLvphSzpxTKQy1TOAK/
Q6XmsvE9ubUUdki5X1rIuFvOzU88KJ0X1c+XPw5+NaI8VXat74YR7dSBE5sbC3EKhZ142fD7IH
+nHfpHiHQIQSiIIFyw= user@example.com
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showsshkey
```

showsys

Syntax

```
showsys [options]
```

Description

The showsys command displays the HPE 3PAR storage system properties such as a system name, serial number, and system capacity information.

Options

-d

Specifies that more detailed information about the system is displayed.

-param

Specifies that the system parameters are displayed.

-fan

Displays the system fan information.

-space

Displays the system capacity information in MiB.

-vvspace

Displays the system capacity information in MiB with an emphasis on VVs.

-domainspace

Displays the system capacity information broken down by domain in MiB.

-desc

Displays the system descriptor properties.

-devtype FC|NL|SSD

Displays the system capacity information where the disks must have a device type string matching the specified device type: either Fast Class (FC), Nearline (NL), Solid State Drive (SSD). This option can only be issued with –space, –domainspace, or –vvspace.

Specifiers

None.

Authority

Any role in the system

Usage

- See the **setsys** command for information on setting the threshold parameters indicated by the Value column in the output for `showsys -param`.
- In the output for `showsys -param`, (from configured settings) indicates that the system parameters displayed have been successfully read from the Persistent Repository (PR). If the PR is not available (most likely because of problems with the admin volume), the output reads (from default settings) and the values displayed would indicate the system defaults. When (from default settings) is displayed, system parameters cannot be updated.
- If the VVRetentionTimeMax is 0, then the volume retention time in the system is disabled.
- For the system capacity information, there might be some overlaps among Volumes, System, and Failed Capacities.

Examples

The following example displays the system descriptor properties of a HPE 3PAR storage system:

```
cli% showsys -desc
------------System s36------------
System Name  :  s36
Location     :  Your Facility Address
Owner        :  Your Company Name
Contact      :  Joe Admin
Comment      :  Your Notes
```

The following example displays more detailed (`-d` option) information about the same storage server:

```
cli% showsys -d
------------General--------------
System Name       :          S424
System Model      : HPE_3PAR 7200
Serial Number     :       1600424
System ID         :           424
Number of Nodes   :             2
Master Node       :             0
Nodes Online      :           0,1
Nodes in Cluster  :           0,1

-----System Capacity (MiB)-----
Total Capacity      :  6277120
Allocated Capacity  :   687872
Free Capacity       :  5589248
Failed Capacity     :        0

---------System Fan---------
Primary Node ID    :       0
Secondary Node ID  :       1
State              :      OK
LED                :   Green
Speed              :  Normal
Primary Node ID    :       1
Secondary Node ID  :       0
State              :      OK
LED                :   Green
Speed              :  Normal

---------------Remote Syslog Status--------------
```

```
Active               :                              1
General Server       :          syslog.example.com
General Connection   :                            UDP
General Security     :   secure_syslog.example.com
Security Connection  :                            TLS

--------System Descriptors--------
Location  :
Owner     :
Contact   :
Comment   :
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showsys
```

showsysmgr

Syntax

```
showsysmgr [options]
```

Description

Displays startup state information about the system manager.

Options

-d

Shows additional detailed information if available.

-l

Shows field service diagnostics for System Manager specific Config Locks and MCALLs, and system-wide ioctl system calls.

Specifiers

None.

Authority

Any role in the system

Usage

- If the system has experienced a power failure, issue the **showvv**, **showld**, or **showpd** commands to determine if any physical disks, logical disks, or virtual volumes are unstarted. If the system stays in the powerfail state for longer than 15 minutes, the `setsysmgr` command can be issued with caution. See the **setsysmgr** command for more information.
- If the system has experienced a power failure, any cached data is saved to the IDE partition of each node. Upon restoration of power, the saved cached data is reapplied (powerfail recovery). During powerfail recovery, if kernel panics are encountered because of hardware or software errors, the powerfail recovery process fails. In such cases, the system attempts recovery three times. After three attempts, the system waits for user intervention. Issue the `setsysmgr force_idewipe` command to force the system to restart and check all logical disks and virtual volumes upon restarting.
- During powerfail recovery, a kernel panic because of hardware or software errors might indicate a damaged volume. In such cases, after the system restarts, it waits for user intervention. Issue the `setsysmgr force_iderecovery` command to force the system to start another powerfail recovery. Any saved data for virtual volumes with IDs displayed in the `showsysmgr` command output can be lost. Those virtual volumes are checked when the system restarts. See **setsysmgr** for more information.

Examples

The following example displays the startup state information about the system manager. In this example, the system manager reports that it is up and running.

```
cli% showsysmgr
System is up and running from 2016-07-21 06:29:06 PDT
```

If the system has experienced a power failure, issuing the `showsysmgr` command displays the following message:

```
cli% showsysmgr
System is recovering from a previous power failure. Please use the CLI
commands for 'showvv', 'showld', 'showpd' to see any unstarted vvs, lds, pds.
```

If the system has attempted powerfail recovery three times, and during the recovery attempts encountered kernel panics because of hardware or software errors, the following message is displayed.

```
cli% showsysmgr
System is recovering from a previous power failure.
Please use the 'showvv', 'showld', 'showpd' CLI commands to
check for not started vvs, lds, pds.
Use force_idewipe to wipe pfail partition and restart system with all
lds/vvs being checked. This can cause data inconsistency.
```

If the system has attempted powerfail recovery, and during the recovery encountered kernel panics because of hardware or software errors, the error could be associated with a specific volume. The following message is displayed.

```
cli% showsysmgr
System is recovering from a previous power failure.
Please use the 'showvv', 'showld', 'showpd' CLI commands to
check for not started vvs, lds, pds.
Use force_iderecovery to start pfail recovery.
VVs with the id(s) 1 will lose data.
```

Displays the following message when the system has online PDs with no cage position assigned yet.

```
cli% showsysmgr
       System has started, PD cage position refresh is pending.
```

showtarget

Syntax

```
showtarget [<Node_WWN>...]

showtarget -lun {all|<Node_WWN>...}

showtarget {-inq|-mode} [-page <num>] [-d] <LUN_WWN>...

showtarget [-force] -rescan
```

Description

The `showtarget` command displays information about unrecognized targets.

Options

`-lun`

Displays the exported Logical Unit Numbers (LUNs) from the unknown targets. Use the "all" specifier to display the exported LUNs from all of the unknown targets.

`-inq`

Displays SCSI inquiry page information.

-mode

Displays SCSI mode page information.

-page <num>

Specifies the SCSI page number for the inquiry and mode information.

<num> is a hex number. For SCSI inquiry information, the valid <num> is 0, 80, 83, and c0. For SCSI mode information, the valid <num> is 3 and 4. This option needs to be used together with -inq or -mode. If this option is not specified, the default <num> is 0.

-d

Displays the detail information of SCSI inquiry or mode page information.

-force

Specifies that the rescan is forced. If this option is not used, the rescan will be suppressed if the peer ports have already been rescanned within the last 10 seconds.

-rescan

Rescans the peer ports to find the unknown targets.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

<Node_WWN>

Indicates the World Wide Name (WWN) of the node.

<LUN_WWN>

Indicates the World Wide Name (WWN) of a LUN exported from an unknown target.

Authority

Any role in the system

Usage

Some hosts might appear as unknown targets.

Examples

The following example displays information about unrecognized targets:

```
cli% showtarget
```

The following example displays exported LUNs from unknown target 2FF70002AC00001F:

```
cli% showtarget -lun 2FF70002AC00001F
```

The following example displays exported LUNs from all unknown targets:

```
cli% showtarget -lun all
```

The following example displays the SCSI inquiry page 0 of target 50002AC00001001F:

```
cli% showtarget -inq -page 0 50002AC00001001F
```

The following example displays the detailed information for SCSI mode page 0x3 of target 50002AC00001001F:

```
cli% showtarget -mode -page 0x3 -d 50002AC00001001F
```

The following example rescans the peer ports to find unknown targets:

```
cli% showtarget -rescan
```

If unknown targets are found, information (Port, LUN_WWN, Port WWN, and State) about each is displayed.

Use the following to get a description of each column in the above examples:

```
clihelp -col showtarget
```

showtask

Syntax

```
showtask [options] [<task_name> | <pattern>]
```

Description

The `showtask` command displays information about tasks. By default, this command displays all non-system tasks on the system within the last 24 hours.

Options

`-all`

Displays all tasks. Unless the `-all` option is specified, system tasks are not displayed.

`-done`

Displays includes only tasks that are successfully completed.

`-failed`

Displays includes only tasks that are unsuccessfully completed.

`-active`

Displays includes only tasks that are currently in progress.

`-t <hours>`

Shows only tasks started within the past `<hours>`, where `<hours>` is an integer from 1 through 99999.

`-type <task_type | <pattern>`

Specifies that specified patterns are treated as glob-style patterns and that all tasks whose types match the specified pattern are displayed. To see the different task types, use the `showtask` column help.

`-d <task_ID>...`

Shows detailed task status for specified tasks. Tasks must be explicitly specified using their task IDs `<task_ID>`. Multiple task IDs can be specified. This option cannot be used in conjunction with other options.

`-sortcol <col>[,<dir>][:<col>[,<dir>]...]`

Sorts command output based on column number (`<col>`). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (`<dir>`) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

None.

Authority

Any role in the system

Usage

- See the *HPE 3PAR Command Line Interface Administrator Guide* for additional information and examples regarding task management and task management commands, as well as information about Dynamic Optimization.
- By default, `showtask` shows all non-system tasks that were active within the last 24 hours.
- The system stores information for the most recent 2,000 tasks. Task ID numbers roll at 29999.

Examples

The following example shows details about a specific task using the task ID (task ID is 3).

```
cli% showtask -d 3
Id Type                Name          Status Phase Step -------StartTime------- ------FinishTime----
 3 snapspace_accounting ss_accounting done     ---  --- 2015-02-27 11:37:46 PST 2015-02-27 11:37:48 PST

Detailed status:
2015-02-27 11:37:46 PST Created     task.
2015-02-27 11:37:46 PST Started     snapshot usage data collection process for VVs
2015-02-27 11:37:46 PST Updated     snapshot usage data for VV North
2015-02-27 11:37:47 PST Updated     snapshot usage data for VV VV_Tech_1.0020.0000.ro
2015-02-27 11:37:48 PST Updated     snapshot usage data for VV VV_Tech_1
2015-02-27 11:37:48 PST Finished     snapshot usage data collection process.
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showtask
```

showtemplate

Syntax

```
showtemplate [options] [<template_name_or_pattern> ...]
```

Description

The `showtemplate` command displays existing templates that can be used for Virtual Volume (VV), Logical Disk (LD) Common Provisioning Group (CPG) creation.

Options

`-t {vv|ld|cpg}`

Specifies that the template type displayed is a VV, LD, or CPG template.

`-fit`

Specifies that the properties of the template is displayed to fit within 80 character lines.

Specifiers

`<template_name_or_pattern>`

Specifies the name of a template, using up to 31 characters or glob-style pattern for matching multiple template names. If not specified, all templates are displayed.

Authority

Any role in the system

Usage

If no options are specified, all existing templates are displayed.

Examples

The following example displays the properties of all templates in CLI line format:

```
cli% showtemplate
Name         Type Other Options
CPGTemplate cpg  -nro -ro -t r6 -ha cage -ssz 2-ss 256 -sdgs 0
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showtemplate
```

showtoc

Syntax

```
showtoc
```

Description

The `showtoc` command displays the system table of contents summary that provides a summary of the system's resources.

Options

None.

Specifiers

None.

Authority

Any role in the system

Usage

None.

Examples

The following example displays the system table of contents:

```
cli% showtoc
gltab toc generation: 313
Toc header magic num: 42444854
Toc in use len: 70944
Toc not in use len: 0
Toc on disk len: 38912
Toc version: 69
Toc generation: 313
Toc pd_entries: 80
Toc ld_entries: 5
Toc vv_entries: 1
Toc cage_entries: 4
```

showtocgen

Syntax

`showtocgen`

Description

The `showtocgen` command displays the table of contents generation number.

Options

None.

Specifiers

None.

Authority

Any role in the system

Usage

The table of contents generation number increases each time there is a change in the system configuration.

Examples

The following example displays the table of contents generation number:

```
cli% showtocgen
Table of Contents generation number: 4292
```

showuser

Syntax

The syntax for the `showuser` command can be the following:

```
showuser [<user_name>]
showuser [-oneline] [<user_name>]
showuser -domain <domain_name> [<user_name>]
showuser -k
showuser -showcols <column>[,<column>...] [<user_name>]
showuser -listcols
```

Description

The `showuser` command displays information about one or all users, including the username, authority level, and system resources to which a given user has access. The `showuser` command shows account information for local users (those created with the **`createuser`** command) or shows the SSH key information for local and LDAP users (enabled by the **`setauthparam`** command).

Options

`-domain`

Shows the list of users for the specified domain.

`-k`

Shows users who have set a public SSH key.

`-oneline`

Shows all information about a user on one line.

-listcols

Lists the columns available to be shown in the -showcols option described below (see 'clihelp -col showuser' for help on each column).

By default (if none of the information selection options below are specified) the following columns are shown:

Default Domain Role Username-showcols <column>[,<column>...]

Explicitly select the columns to be shown using a comma-separated list of column names. For this option the full column names are shown in the header.

Run 'showuser -listcols' to list the available columns.

Run 'clihelp -col showuser' for a description of each column.

Specifiers

<user_name>

Specifies the login name, using up to 31 characters. This specifier is not required on the command line. If no specifier is used, information about all users is displayed.

Authority

Any role in the system

Usage

- For the first usage (without the -k option), the showuser command displays account information for all local users when the <username> is not provided or the account information for just the local user with the name <username> when it is provided. The output for users with Browse or Service role is limited to the user's own accounts. If the <username> is provided, it must be the user's name. The output shows the users' username, domain, role, and whether the domain is the default for the user.
- Because the output is only for local users and is further limited when the user is at the Browse or Service authority level, an LDAP user may not see any user account information. Alternatively, both local and LDAP users can use the showuserconn command to see the role of the current user.
- See the column help for **showuserconn** for the different definition of "local" used in the output of showuserconn.
- For the second usage (-k), users who have set a public SSH key with the **setsshkey** command are displayed, one per line. For users with Browse or Service role, the output is limited to the current user.
- Users are created using the **createuser** command.

Examples

The following example indicates user Joe being a member of all domains. Valid output includes all, or <domain_name>:

```
cli% showuser Joe
Username Domain Privilege Default
Joe      all    super     N
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showuser
```

showuseracl

Syntax

showuseracl

Description

The `showuseracl` command displays a user's access control list (ACL).

Options

None.

Specifiers

None.

Authority

Any role in the system

Usage

None.

Examples

The following example displays ACL information about all system users:

```
cli% showuseracl
User        Operation  Object_Names_or_Patterns
suser         updatevv avvro*
buser         updatevv avvr*,vv0,cpvv0,tpvv0 -f
ruser         updatevv vv0 cpvv0 tpvv0 avv*
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showuseracl
```

showuserconn

Syntax

```
showuserconn [options <arg>] [<pid> ...]
```

Description

The `showuserconn` command displays information about users who are currently connected (logged in) to the storage system.

Options

`-current`

Shows all information about the current connection only.

`-d`

Specifies the more detailed information about the user connection.

`-listcols`

Lists the columns available to be shown in the `-showcols` option described below (see '`clihelp -col showuserconn`' for help on each column).

By default (if none of the information selection options below are specified) the following columns are shown:

`Client ClientName Connected_since Current Id IP_Addr Name Role-showcols <column>[,<column>...]`

Explicitly select the columns to be shown using a comma-separated list of column names. For this option the full column names are shown in the header.

Run '`showuserconn -listcols`' to list the available columns.

Run 'clihelp -col showuserconn' for a description of each column.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

<pid>

Specifies the storage system process ID for the user connection.

Authority

Any role in the system

Usage

A Domain column may be included by using the **setclienv** command to set listdom to 1. The listdom option is also set to 1 if the CLI was started with the -listdom option, or if the CLI was started with the TPDLISTDOM environment variable set. Run "cli -h" and "setclienv -h" for details of the environment variables.

Examples

The following example displays information about user connections:

```
cli% showuserconn
   Id Name    -----IP_Addr------ -Role- ----Connected_since---- Current Client ClientName
28888 3paradm 192.168.17.7:20000 super  2017-07-19 10:38:23 MDT current local  SSH
-------------------------------------------------------------------------------------
    1 total
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showuserconn
```

showvasa

Syntax

```
showvasa [options]
```

Description

Displays configuration and status information of the VASA Provider service.

The VASA (VMware APIs for Storage Awareness) web service provider allows discovery and management of the HPE 3PAR storage system by a VMware vSphere environment. The VASA web service provider is based on SOAP and allows a vSphere environment to:

- Discover storage system configuration, including storage space (LUNs, vVols, Storage Containers and Protocol Endpoints), ports, processors and capabilities.
- Discover storage system statistics, such as I/O statistics.

- Monitor storage system events such as configuration changes, and alarms, such as space capacity constraints.
- Manage vVols, including creation, binding, snapshots, reverting, cloning (full copy or linked-clones for data-deduplication), migration, and removal.

The VASA provider must be registered (a storage system user credential is required) with the vCenter server, using the following URL: `https://<hostname>:9997/vasa`

If a VASA SSL certificate does not exist, showvasa will display the the proposed IP addresses and/or hostnames that should be included in the Subject Alternative Name (SAN) of the certificate. If a VASA SSL certificate exists, showvasa will display hostnames and/or IP addresses in the VASA URLs, based on the SAN entry/entries in the VASA certificate. If any of these SAN entries are invalid (do not match the proper address or host name for the array) their state will be Invalid while the VASA Provider is started. Only a valid URL should be used when registering the VASA Provider.

The `showvasa` command displays the following information:

- VASA Version
- VASA URL
- VASA SSL Server certificate subject and thumbprint
- VASA Provider state (Proposed/Disabled/Enabled/Invalid)
- VASA Provider memory consumption

Options

`-cert`

Shows VASA SSL certificate information including certificate subject and thumbprint as well as the certificate management mode.

`-config`

Show the current values for VASA configuration options.

Specifiers

None.

Authority

Any role in the system

Usage

The memory usage column may display a dash (–) if the VASA Provider server is disabled.

Examples

The following example shows the VASA Provider server information:

```
cli% showvasa
-State- --------VASA_API2_URL------- -MemUsage(MB)- -Version-
Enabled https://<hostname>:9997/vasa        55        0.1
Enabled   https://<althostname>:9997/vasa    -          -

cli% showvasa -cert
Mgmt_Mode -------------------Subject-------------------
-----------------------Fingerprint----------------------- -Locked_By_Uuid-
client    C=US/ST=CA/L=SanJose/O=HPE/OU=PSE/CN=VASATeam 5D:95:B9:1F:C8:7F:
43:07:F7:3B:66:73:6F:A0:38:93:52:DA:EE:1D -

cli% showvasa -config
Key          Value
```

```
TLSMethod    tlsv1_2
CopyPriority high
```

The following example shows VASA Provider server information if the VASA certificate contains invalid SAN entries:

```
cli% showvasa
    -Service- --------VASA_API2_URL-------------- -MemUsage(MiB)- -Version-
    Enabled   https://<validhostname>:9997/vasa                55 0.1
    Invalid   https://<invalidhostname>:9997/vasa               - -
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showvasa
```

showversion

Syntax

```
showversion [options]
```

Description

The `showversion` command displays information about the storage server software. Use this command to determine if the storage system is using the latest software version.

Options

`-a`

Shows all component versions.

`-b`

Shows build levels.

`-s`

Shows release version number only (useful for scripting).

Specifiers

None.

Authority

Any role in the system

Usage

- When displaying all versions, for certain components multiple versions might be installed. In such cases, multiple lines are displayed.
- If no options are specified, the overall version of the software is displayed.

Examples

The following example displays comprehensive information about the system software:

```
cli% showversion
Release version     3.2.2
Patches:            None
Component Name      Version
CLI Server          3.2.2
CLI Client          3.2.2
System Manager      3.2.2
```

```
Kernel               3.2.2
TPD Kernel Code      3.2.2
```

showvfs

Syntax

```
showvfs [options]
```

Description

The `showvfs` command displays information on Virtual File Servers.

Options

`-d`

Detailed output.

`-fpg <fpgname>`

Limits the display to VFSs contained within the File Provisioning Group.

`-vfs <vfs>`

Limits the display to the specified VFS name.

Specifiers

None.

Authority

Any role in the system.

Usage

- Access to all domains is required to run this command.
- VFS name is not globally unique, and the same VFS name may be in use in multiple File Provisioning Groups.
- If no filter options are provided the system will traverse all File Storage Pools and display all associated VFSs.

Examples

The following example shows a verbose listing of examplevfs:

```
cli% showvfs -d -vfs examplevfs
```

showvlun

Syntax

```
showvlun [options <arg>]
```

```
showvlun -listcols
```

Description

The `showvlun` command displays information about Virtual Logical Unit Numbers (VLUNs) in the system, such as all active and template VLUNs. The display is divided into two sections: the upper provides information about active VLUNs and the lower provides information about VLUN templates.

Options

-listcols

Lists the columns available to be shown in the -showcols option described below (see 'clihelp -col showvlun' for help on each column).

The [options] are generally of two kinds: those that select the type of information that is displayed, and those that filter the list of VLUNs that are displayed.

By default (if none of the information selection options below are specified) the following columns are shown:

```
Lun VVName HostName Host_WWN Port Type
```

Options that select the type of information shown include the following:

-showcols <column>[,<column>...]

Selects the columns to be shown using a comma-separated list of column names. For this option the full column names are shown in the header.

Run 'showvlun -listcols' to list the available columns.

Run 'clihelp -col showvlun' for a description of each column.

-lvw

Shows the WWN of the virtual volume associated with the VLUN.

-pathsum

Shows path summary information for active VLUNs.

-hostsum

Shows mount point, Bytes per cluster, capacity information from Host Explorer and user reserved space, VV size from showvv.

Options that filter the VLUNs that are displayed include the following:

-a

Shows only active VLUNs.

-t

Shows only VLUN templates.

-host {<hostname>|<pattern>|<hostset>}...

Displays only VLUNs exported to hosts that match <hostname> or glob-style patterns, or to the host sets that match <hostset> or glob-style patterns. (For more information, enter clihelp sub,globpat.) The host set name must start with set:. Multiple host names, host sets or patterns can be repeated using a comma-separated list.

-v {<VV_name>|<pattern>|<VV_set>}...

Displays only VLUNs of virtual volumes that match <VV_name> or glob-style patterns, or to the vv sets that match <VV-set> or glob-style patterns. (For more information, enter clihelp sub,globpat. The VV set name must start with "set:". Multiple volume names, vv sets or patterns can be repeated using a comma-separated list (for example -v <VV_name>, <VV_name>...).

-l {<LUN|pattern>}...

Specifies that only exports to the specified LUN are displayed. This specifier can be repeated to display information for multiple LUNs.

-nodes <nodelist>

Requests that only VLUNs for specific nodes are displayed. The node list is specified as a series of integers separated by commas (for example 0,1,2). The list can also consist of a single integer (for example 1).

-slots <slotlist>

Requests that only VLUNs for specific slots are displayed. The slot list is specified as a series of integers separated by commas (for example 0,1,2). The list can also consist of a single integer (for example 1).

-ports <portlist>

Requests that only VLUNs for specific ports are displayed. The port list is specified as a series of integers separated by commas ((for example 1,2). The list can also consist of a single integer (for example 1).

-domain {<domain_name|pattern>}...

Shows only the VLUNs whose virtual volumes are in domains with names that match one or more of the <domainname_or_pattern> options. This option does not allow listing objects within a domain of which the user is not a member. Multiple domain names or patterns can be repeated using a comma-separated list.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

None.

Authority

Any role in the system

Usage

- The use of VV sets for the -v option and host sets for the -host option affects the listing of VLUNs related to these sets, rather than each individual member. As a result only VLUN templates will be shown; active VLUNs always show each host and/or VV.
- A Domain column may be included by using the **setclienv** command to set listdom to 1. The listdom option is also set to 1 if the CLI was started with the -listdom option, or if the CLI was started with the TPDLISTDOM environment variable set. Run "`cli -h`" and "`setclienv -h`" for details of the environment variables.

Examples

The following example displays all active and template VLUNs using glob-style patterns:

```
cli% showvlun -v *.2,*.3
Active VLUNs
 Lun VVname     Host ----Host_WWN----  Port Type
   2  tpv.2 pe750-07 210000E08B056C21 0:2:1 host
   3  tpv.3 pe750-07 210000E08B056C21 0:2:1 host
   2  tpv.2 pe750-07 210100E08B256C21 1:2:1 host
   3  tpv.3 pe750-07 210100E08B256C21 1:2:1 host
   2 test.2      sun 210000E08B023F71 0:2:2 host
   3 test.3      sun 210000E08B023F71 0:2:2 host
   2 test.2      sun 210000E08B023C71 1:5:1 host
   3 test.3      sun 210000E08B023C71 1:5:1 host
 ----------------------------------------------
   8 total

VLUN Templates
 Lun VVname     Host ----Host_WWN----  Port Type
```

```
  2  tpv.2 pe750-07 ----------------   --- host
  3  tpv.3 pe750-07 ----------------   --- host
  2 test.2      sun ----------------   --- host
  3 test.3      sun ----------------   --- host
-----------------------------------------------
  4 total
```

The following will show all VLUNs exported to host set host_set1: s

```
cli% showvlun -host set:host_set1
```

The following will show all VLUNs exported from VV set vv_set:

```
cli% showvlun -v set:vv_set
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showvlun
```

showvv

Syntax

```
showvv [options] [<VV_name|pattern|VV_set> ...]
showvv -listcols
```

Description

The `showvv` command displays information about all Virtual Volumes (VVs) or a specific VV in a system.

Options

`-listcols`

Lists the columns available to be shown in the `-showcols` option described below (see '`clihelp -col showvv`' for help on each column).

The `[options]` are generally of two kinds -- those that select the type of information that is displayed, and those that filter the list of VVs that are displayed.

By default (if none of the information selection options below are specified) the following columns are shown:

`Id Name Prov Compr Dedup Type CopyOf BsId Rd State Adm_Rsvd_MB Snp_Rsvd_MB Usr_Rsvd_MB VSize_MB`

Options that select the type of information shown include the following:

`-showcols <column>[,<column>...]`

Explicitly select the columns to be shown using a comma-separated list of column names. For this option the full column names are shown in the header.

Run '`showvv -listcols`' to list the available columns.

Run '`clihelp -col showvv`' for a description of each column.

`-d`

Displays detailed information about the VVs. The following columns are shown:

`Id Name Rd Mstr Prnt Roch Rwch PPrnt PBlkRemain VV_WWN CreationTime Udid`

`-pol`

Displays policy information about the VVs. The following columns are shown:

```
Id Name Policies
```

-space (-s)

Displays Logical Disk (LD) space use by the VVs. The following columns are shown:

```
Id Name Prov Compr Dedup Type Adm_Rsvd_MB Adm_Used_MB Snp_Rsvd_MB
Snp_Used_MB Snp_Used_Perc Warn_Snp_Perc Limit_Snp_Perc Usr_Rsvd_MB
Usr_Used_MB Usr_Used_Perc Warn_Usr_Perc Limit_Usr_Perc Tot_Rsvd_MB
Tot_Used_MB VSize_MB Host_Wrt_MB Compaction Compression
```

NOTE: For snapshot (vcopy) VVs, the Adm_Used_MB, Snp_Used_MB, Usr_Used_MB and the corresponding _Perc columns have a '*' before the number for two reasons: to indicate that the number is an estimate that must be updated using the **updatesnapspace** command, and to indicate that the number is not included in the total for the column since the corresponding number for the snapshot's base VV already includes that number.

-r

Displays raw space use by the VVs. The following columns are shown:

```
Id Name Prov Compr Dedup Type Adm_RawRsvd_MB Adm_Rsvd_MB Snp_RawRsvd_MB
Snp_Rsvd_MB Usr_RawRsvd_MB Usr_Rsvd_MB Tot_RawRsvd_MB Tot_Rsvd_MB
```

-zone

Displays mapping zone information for VVs. The following columns are shown:

```
Id Name Prov Compr Dedup Type VSize_MB Adm_Zn Adm_Free_Zn Snp_Zn
Snp_Free_Zn Usr_Zn Usr_Free_Zn
```

-g

Displays the SCSI geometry settings for the VVs. The following columns are shown: `Id Name SPT HPC SctSz`

-alert

Indicates whether alerts are posted on behalf of the VVs. The following columns are shown:

```
Id Name Prov Compr Dedup Type VSize_MB Snp_Used_Perc Warn_Snp_Perc
Limit_Snp_Perc Usr_Used_Perc Warn_Usr_Perc Limit_Usr_Perc Alert_Adm_Fail_Y
Alert_Snp_Fail_Y Alert_Snp_Wrn_Y Alert_Snp_Lim_Y Alert_Usr_Fail_Y
Alert_Usr_Wrn_Y Alert_Usr_Lim_Y
```

-alerttime

Shows times when alerts were posted (when applicable). The following columns are shown:

```
Id Name Alert_Adm_Fail Alert_Snp_Fail Alert_Snp_Wrn Alert_Snp_Lim
Alert_Usr_Fail Alert_Usr_Wrn Alert_Usr_Lim
```

-cpprog

Shows the physical copy and promote progress. The following columns are shown:

```
Id Name Prov Compr Dedup Type CopyOf VSize_MB Copied_MB Copied_Perc
```

-cpgalloc

Shows CPGs associated with each VV. The following columns are shown:

```
Id Name Prov Compr Dedup Type UsrCPG SnpCPG
```

-state

Shows the detailed state information for the VVs. The following columns are shown:

```
Id Name Prov Compr Dedup Type State Detailed_State SedState
```

-hist

Shows the history information of the VVs. The following columns are shown:

```
Id Name Prov Compr Dedup Type CreationTime RetentionEndTime ExpirationTime
SpaceCalcTime Comment
```

-rcopy

This option appends two columns, RcopyStatus and RcopyGroup, to any of the display options above.

-notree

Do not display VV names in tree format.

Unless either the `-notree` or the `-sortcol` option described below are specified, the VVs are ordered and the names are indented in tree format to indicate the virtual copy snapshot hierarchy.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (`<col>`). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (`<dir>`) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

A Domain column may be included by using the **setclienv** command to set listdom to 1. The listdom option is also set to 1 if the CLI was started with the `-listdom` option, or if the CLI was started with the TPDLISTDOM environment variable set. Run "`cli -h`" and "`setclienv -h`" for details of the environment variables.

The VV filtering option include:

-p <pattern>

Pattern for matching VVs to show (see below for description of `<pattern>`) If the `-p` option is specified multiple times, each instance of `<pattern>` adds additional candidate VVs that match that pattern.

A `<pattern>` is one or more of the following:

-cpg {<cpgname>|<pattern>}[,{<cpgname>|<pattern>}...]

Shows only VVs whose UsrCPG or SnpCPG matches the one or more of the cpgname_or_patterns.

-prov {<prov>|<pattern>}[,{<prov>|<pattern>}...]

Shows only VVs with Prov (provisioning) values that match the prov_or_pattern.

-type {<type>|<pattern>}[,{<type>|<pattern>}...]

Shows only VVs of types that match the type_or_pattern.

-host {<hostname>|<pattern>}[,{<hostname>|<pattern>}...]

Shows only VVs that are exported as VLUNs to hosts with names that match one or more of the hostname_or_patterns.

-baseid {<baseid>|<pattern>}[,{<baseid>|<pattern>}...]

Shows only VVs whose BsId column matches one more of the baseid_or_patterns.

-copyof {<vvname>|<pattern>}[,{<vvname>|<pattern>}...]

Shows only VVs whose CopyOf column matches one more of the vvname_or_patterns.

-rcopygroup {<groupname>|<pattern>}[,{<groupname>|<pattern>}...]

Shows only VVs that are in Remote Copy groups that match one or more of the groupname_or_patterns.

-policy {<policy>|<pattern>}[,{<policy>|<pattern>}...]

Shows only VVs whose policy matches the one or more of the policy_or_pattern.

-vmname {<vmname>|<pattern>}[,{<vmname>|<pattern>}...]

Shows only VVs whose vmname matches one or more of the vvname_or_patterns.

-vmid {<vmid>|<pattern>}[,{<vmid>|<pattern>}...]

Shows only VVs whose vmid matches one or more of the vmids.

-vmhost {<vmhost>|<pattern>}[,{<vmhost>|<pattern>}...]

Shows only VVs whose vmhost matches one or more of the vmhost_or_patterns.

-vvolstate {<bound|unbound>}

Shows only VVs whose vvolstate matches the specified state - bound or unbound.

-vvolsc {<vvset_name>|<pattern>}[,{<vvset_name>|<pattern>}...]

Shows only VVs whose storage container (vvset) name matches one or more of the vvset_name_or_patterns.

-domain <domain_name|pattern>[,<domain_name|pattern>...]

Shows only VVs that are in domains with names matching one or more of the specified domain_name or patterns. This option does not allow listing objects within a domain of which the user is not a member.

Multiple domain names or patterns can be repeated using a comma-separated list (for example `-domain <domain_name>, <domain_name>...`).

-expired

Shows only VVs that have expired.

-exp <time>{d|D|h|H|m|M}

Specifies the relative time from the current time that volume will expire. <time> is a positive integer value and in the range of 1 minute - 1825 days. Time can be specified in days, hours, or minutes. Use "d" or "D" for days, "h" or "H" for hours, or "m" or "M" for minutes following the entered time value.

`<time>` is a positive integer value and in the range of 1 - 43,800 hours (1,825 days). Time can be optionally specified in days or hours providing either d or D for day and h or H for hours following the entered time value.

-retained

Shows only VVs that have a retention time.

-failed

Shows only failed VVs.

The following VV filtering options are deprecated and will be removed in a future release. Instead, use the appropriate `-p <pattern>` option described above.

-cpg {<cpgname>|<pattern>}[,{<cpgname>|<pattern>}...]

Shows only VVs whose UsrCPG or SnpCPG matches the one or more of the cpgname_or_patterns. This option is deprecated without the preceding `-p`. Use `-p -cpg` option described above instead.

-tpvv

Shows only thin provisioned VVs (TPVVs). This option is deprecated.

Use the `-p -prov tp*` option described above instead.

-host {<hostname>|<pattern>}[,{<hostname>|<pattern>}...]

Shows only VVs that are exported as VLUNs to hosts with names that match one or more of the hostname_or_patterns.

This option is deprecated without the preceding -p. Use -p -host option described above instead.

-rcopygroup {<groupname>|<pattern>}[,{<groupname>|<pattern>}...]

Shows only VVs that are in Remote Copy groups that match one or more of the groupname_or_patterns.

This option is deprecated without the preceding -p. Use -p -rcopygroup option described above instead.

Specifiers

<VV_name|pattern|VV_set>...

Specifies the VVs with the specified name (up to 31 characters in length), matching the glob-style pattern or that are members of the supplied VV set will have their information displayed. The VV set name must start with "set:". This specifier can be repeated to display configuration information about multiple VVs. This specifier is not required. If not specified, configuration information for all VVs in the system is displayed.

Authority

Any role in the system

Usage

- The per-snapshot space is approximate and must be calculated using the **updatesnapspace** command.
- For all the <name_or_pattern> fields, the patterns are glob-style (shell-style) patterns. (For more information, enter clihelp sub,globpat or see **Glob-style patterns**.)
- If the -showcols option is used the full column names are shown in the header otherwise the header contains abbreviated column names.
- Compression column is under Efficiency column which displays the compression ratio. Compr column denotes the current Compression state of the thinly compressed volume.
- The Compr column holds one of the following values:
  - Yes - Compression is currently enabled.
  - No - Compression is currently disabled.
  - NA - Compression is not applicable. (In case of non-compressed volumes.)
- The Dedup column indicates whether it is a thinly deduplicated volume or not.
  - Yes - Deduplicated volume.
  - No - Not a Deduplicated volume.
- A Domain column may be included by using the **setclienv** command to set listdom to 1. The listdom option is also set to 1 if the CLI was started with the -listdom option, or if the CLI was started with the TPDLISTDOM environment variable set. Run "cli -h" and "setclienv -h" for details of the environment variables.
- If there are greater than 8192 VVs and snapshots on the system then there may be issues displaying the Shared volume and TDVV tree. Some TDVVs may not appear under the Shared volume and instead appear later in the list.

  For this command, 1 MB = 1048576 bytes.

Examples

The following example displays space usage information for thin provisioned VVs that are exported to host hname:

```
cli% showvv -s -p -prov tp* -host hname
```

Note that the 'tp*' matches tpvv as well as tpsd.

The following example displays raw space usage information for thin provisioned VVs exported to host hname and all VVs in the tree with the base VV of Id 50:

```
cli% showvv -r -p -prov tp* -host hname -p -baseid 50
```

The following example lists the Id, Name and VSize_MB columns only for thin provisioned VVs exported to host hname:

```
cli% showvv -showcols Id,Name,VSize_MB -p -prov tp* -host hname
```

The following will show all VVs in VV set vv_set:

```
cli% showvv set:vv_set
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showvv
```

showvvcpg

Syntax

```
showvvcpg [options] [<vvname|pattern|VV_set> ...]
```

Description

The `showvvcpg` command displays the Virtual Volume (VV) space distribution among CPGs.

Options

`-domain <domain_name|pattern>[,<domain_name|pattern>...]`

Shows only VVs that are in domains with names matching one or more of the specified domain_name or patterns. Only VVs within domains that the user is a member of can be listed. Multiple domain names or patterns can be repeated using a comma-separated list (for example -domain <domain_name>,<domain_name>…).

`-sortcol <col>[,<dir>][:<col>[,<dir>]...]`

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

`<vvname|pattern>...`

Shows only virtual volumes with the specified names, names that match the glob-style patterns or belong to matching VV sets. The VV set specifier must start with 'set:' and be followed by a VV set name or glob-style pattern.

Authority

Any role in the system

Usage

- A Domain column may be included by using the **setclienv** command to set listdom to 1. The listdom option is also set to 1 if the CLI was started with the -listdom option, or if the CLI was started with the TPDLISTDOM environment variable set. Run "`cli -h`" and "`setclienv -h`" for details of the environment variables.
- The columns shown under the heading 'New (MB)' are empty (0) unless some regions are currently being moved from one CPG to another. For the duration of the move, the space being moved appears under both the normal '(MB)' heading for the old CPG and under the 'New (MB)' heading for the new CPG. After the move is completed, the space is removed from the old CPG and appears only under the normal '(MB)' heading for the new CPG.

Examples

The following shows VV space distribution among CPGs in the 'ESX' domain for VVs with names that end in 'sysrpt'.

```
cli% showvvcpg -domain ESX *sysrpt
                       -----(MB)------ -New (MB)--
Name       CPG         Adm  Snp    Usr Adm Snp Usr
esx-sysrpt esx-fc-r1   256 1920 114560   0   0   0
esx-sysrpt esx-ssd-r1  256 1152   7296   0   0   0
------------------------------------------------
1          -           512 3072 121856   0   0   0
```

The following shows VV space distribution for VVs belonging to VV sets with names that begin with 'tv': s

```
cli% showvvcpg set:tv*
               ----(MB)---- -New (MB)--
Name      CPG  Adm Snp  Usr Adm Snp Usr
testvv    cpg0   0   0 1024   0   0   0
newtestvv cpg0   0   0 1024   0   0   0
--------------------------------------
2         -      0   0 2048   0   0   0
```

The following example shows how Adaptive Optimization has distributed space for the Adm, Snp and Usr space for a VV among 3 CPGs:

```
cli% showvvcpg Win2k8-MySqlData
-----(MB)------ -New (MB)--
Name             CPG           Adm  Snp    Usr Adm Snp Usr
Win2k8-MySqlData sysrpt-fc-r1    -    -  41472   0   0   0
Win2k8-MySqlData sysrpt-nl-r6  352 8064 191872   0   0   0
Win2k8-MySqlData sysrpt-ssd-r1 160    -    384   0   0   0
-------------------------------------------------------
1                -             512 8064 233728   0   0   0
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showvvcpg
```

showvvmap

Syntax

`showvvmap <VV_name>`

Description

The `showvvmap` command displays information about how virtual volume regions are mapped to logical disks.

Options

None.

Specifiers

<VV_name>

The virtual volume name.

Authority

Any role in the system

Usage

None.

Examples

The following example shows the VV to LD mapping of a volume:

```
cli% showvvmap vv0
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showvvmap
```

showvvolsc

Syntax

```
showvvolsc [options] [<SC_name>...|<pattern>...]
showvvolsc -listcols
```

Description

Displays vVol storage containers that contain VMware Volumes for Virtual Machines (vVols). Storage containers include logical collections of vVols managed by the vSphere environment. In the HPE 3PAR OS, VV sets (see, **showvvset**) represent storage containers, and the **setvvolsc** and **showvvolsc** commands manage vSphere specific data associated with the VV set used as a storage container.

Options

-listcols

Lists the columns available for viewing with the `-showcols` option. Any column labeled `MB` is reported in `MiB`

-showcols <column>[,<column>...]

Selects the columns to show. Use a comma to separate multiple column names. This option lists the full column names in the header.

Run `showvvolsc -listcols` to list the available columns.

-d

Displays detailed information about the storage containers, including any vVols that have been auto-dismissed by remote copy DR operations.

Specifiers

None.

Authority

Any role in the system

Usage

None.

Examples

The following example displays all the Storage Containers available to the current user:

```
cli% showvvolsc
```

The following example displays the Storage Container with name SC1:

```
cli% showvvolsc SC1
```

showvvolvm

Syntax

```
showvvolvm [options] -sc <container_name> [<VM_name>]
showvvolvm -listcols
```

Description

Displays information about all virtual machines (vVol-based) or a specific virtual machine in a system. Use this command to determine the association between virtual machines and their associated virtual volumes. The command also shows the space usage information for a virtual machine.

Options

Use the following options to select the type of information you want to see:

-listcols

Lists all of the columns available to view with the -showcols option (enter clihelp -col showvvolvm for help on each column).

With mandatory -sc and no other options selected, displays the following columns:

```
VM_Name  GuestOS  VM_State  Num_vv  Physical Logical
```

-showcols *<column>*[,*<column>*...]

Selects the columns you want to see. Use a comma-separated list of column names. The command returns the full column names in the header.

To list available columns, enter showvvolvm -listcols.

To see a description of each column, enter clihelp -col showvvolvm for a description of each column.

-d

Displays detailed information about the VMs. Displays the following columns:

```
VM_Name UUID Num_vv Num_snap Physical Logical GuestOS VM_State UsrCPG SnpCPG Container CreationTime
```

When using other options, the columns displayed can change (see, the -vv option).

-sc *<container_name>*

Mandatory. Displays VMs from the specified *<container_name>* only.

Use sys:all to display all VMs.

-sp

Displays the storage profiles that have constraints associated with the VM or individual vVol. Often, all vVols associated with a VM use the same storage profile. However, if vSphere has provisioned different VMDK volumes with different storage profiles, only the storage profile for the first virtual disk (VMDK) vVol is displayed. In this case, use the -vv option to display storage profiles for individual volumes associated with the VM. Without the -vv option, the command displays the following columns:

```
VM_Name SP_Name SP_Constraint_List
```

-summary

Displays a summary of VMs in the system, including the total number of VMs, VVs, and total physical and exported space used. Displays the following columns:

```
Num_vm Num_vv Physical Logical
```

-binding

Displays the detailed binding information for the VMs. The binding state could be `PoweredOn`, `Bound (exported)`, or `Unbound`. When the VM is bound, `showvvolvm` displays the host names to which it is bound. When the VM is bound and you use the `-vv` option, `showvvolvm` displays the exported LUN templates for each volume, and the state for actively bound vVols. `PoweredOn` means the VM is powered on. `Bound` means the VM is not powered on, but is being created, modified, or queried, or the VM is changing a powered state (on to off or off to on). Due to bind caching, a VM may be reported as `Bound` if it was recently powered off. `Unbound` means the VM is powered off. Displays the following columns:

```
VM_Name VM_State Last_Host Last_State_Time Last_Pwr_Time
```

With the -vv option, displays the following columns:

```
VM_Name vVol_Name vVol_Type vVol_State vVol_LunId Bind_Host Last_State_Time
```

-vv

Displays all the VVs (Virtual Volumes) associated with the VM. Use the `-d` option to display the following columns:

```
VM_Name VV_ID vVol_Name vVol_Type Prov Compr Dedup Physical Logical
```

Use the -vv option to display the following columns:

```
VM_Name VV_ID vVol_Name vVol_Type Prov Compr Dedup Physical Logical VMW_Type
        Shared Container
```

The columns displayed can change when used with other options (see, `-binding` option).

-rcopy

Displays the remote copy group name, sync status, role, and last sync time of the volumes associated with a VM. If a VM does not report as synced, the last sync time for the VM **DOES NOT** represent a consistency point. True consistency points are only represented by the `showrcopy LastSyncTime`. You can combine this option with the `-vv`, `-binding`, `-d`, and `-sp` options.

-autodismissed

Shows only VMs containing automatically dismissed volumes. Combine with the `-vv` option to display automatically dismissed vVols only.

VM filtering options include:

-p *<pattern>*

Displays VMs that match the pattern. Use the `-p` option multiple times to add VMs that match the *<pattern>*.

A *<pattern>* is one or more of the following:

-vmhost *<vmhost_or_pattern>*[,*<vmhost_or_pattern>*...]

Displays only VMs where the vSphere host name matches one or more of the *<vmhost_or_pattern>*.

Host names are those shown by the **showhost** command.

-vmid *<uuid>*[,*uuid*>...]

Displays only VMs where the VM UUID matches the specified *<uuid>*.

-vmstate *<bindingstate>*

Displays only VMs where the binding state matches the requested *<bindingstate>*. Legal values for state are defined in `-showcols`.

-domain *<domain_name|pattern>*[,*<domain_name|pattern>*...]

Displays only VMs that are in domains with names matching one or more of the specified *<domain_name|pattern>*. You can list only those domains of which you are a member. Separate multiple domain names or repeated patterns using a comma.

Specifiers

<VM_name|pattern>

Specifies the VM name (up to 80 characters). Repeat this specifier to display information about multiple VMs. If not specified, `showvvolvm` displays information for all VMs.

<containerName>

The name of the virtual volume storage container. May be "sys:all" to display all VMs.

Authority

Any role in the system

Usage

- For all the *<pattern>* fields, the patterns are glob-style (shell-style). (For more information, enter the `clihelp sub,globpat` command or see, **Glob-style patterns**.)

  The `showvvolvm` command displays the original VM name. If you have renamed a VM in vSphere, the `showvvolvm` command does not reflect that change. This behavior is similar to renaming a VM in a traditional datastore, the VMFS directory name for the VM does not change. vSphere does not pass the new VM name to the array.

- The vVol-to-VM association data displayed is for informational purposes only. Because of limitations in how and when vSphere updates the vVols associated VM metadata, `showvvolvm` provides a best-effort of information when showing which vVols are associated with which VMs. This limitation is especially apparent when a vVol is shared or transferred between two VMs.

- Following are the allowed values for storage profile constraints:

  SpaceEfficiency

  The allowed values for VMware defined space efficiency settings. Legal values are Thin and Thick.

  CPG

  The allowed CPG(s) for base volumes. Note, although the VMware storage profile may allow more than one value, once provisioned only one of the allowed values will be used for the user CPG.

  SnapshotCPG

  The allowed CPG(s) for snapshots. Note, although the VMware storage profile may allow more than one value, once provisioned only one of the allowed values will be used for the snapshot CPG.

  ThinPersistence

  The allowed values for the Thin Persistence setting. Legal values are Enabled and Disabled. If not specified, zero detect will be enabled by default for thinly provisioned volumes.

DataReduction

The required setting for Data Reduction. The value can be True or False. If not specified, Data Reduction is enabled by default if the CPG supports SSD drives.

AdaptiveFlashCache

The required setting for adaptive flash cache. The value may be True or False. If not specified, flash cache will be enabled by default if flash cache has been configured for array-wide support.

SCMCache

The required setting for SCM cache. Legal values are Enabled and Disabled. If not specified, SCM cache will be enabled by default if SCM cache has been configured for array-wide support

Examples

Displaying information for all virtual machines in a storage container `scname`:

```
cli% showvvolvm -sc scname1
```

Displaying the volumes with a virtual machine name of `WebServer1` in a storage container `scname1`:

```
cli% showvvolvm -sc scname1 -vv WebServer1
```

Displaying space usage information for VMs that are exported to host hname:

```
cli% showvvolvm -sc scname1 -summary -p -vmhost hname
```

The following Lists the Name and Phys_MB columns only for VMs exported to host hname:

```
cli% showvvolvm -sc scname1 -showcols VM_Name,Physical_MB -p -vmhost hname
```

The following Provides a description of each column in the above examples:

```
clihelp -col showvvolvm
```

showvvpd

Syntax

```
showvvpd [option <arg>] {<VV_name|pattern>}...
```

Description

The `showvvpd` command displays virtual volume (VV) distribution across physical disks (PD).

Options

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

<VV_name>|<pattern>...

Specifies the virtual volume with the specified name (31 character maximum) or matches the glob-style pattern for which information is displayed. This specifier can be repeated to display configuration information about multiple virtual volumes. This specifier is not required. If not specified, configuration information for all virtual volumes in the system is displayed.

Authority

Any role in the system

Usage

- The command prints the number of chunklets in each physical disk that contain data for the Snapshot Admin (SA), Snapshot Data (SD), User space as well as the total for all spaces for all the virtual volumes that match the <VV_name> or <pattern> list. This can be useful in determining how evenly the virtual volume is striped across the disks.
- Not all the data in the chunklets is necessarily allocated exclusively to the selected virtual volumes. In some cases, only part of a chunklet could be data that is allocated for the virtual volumes.

Examples

The following example displays information about multiple virtual volumes. In this case multi.0, multi.1, multi.2 and multi.3 were created using the -cnt 4 option. Using multi.*, the aggregate chunklets for all the physical disks is displayed: s

```
cli% showvvpd multi.*
    Id Cage_Pos SA SD usr total
     0    0:0:0  0  0   0     0
     1    0:0:1  0  0   0     0
     2    0:0:2  0  0   0     0
     3    0:0:3  0  0   0     0
     4    0:1:0  0  0   0     0
     5    0:1:1  0  0   0     0
     6    0:1:2  0  0   0     0
     7    0:1:3  0  0   0     0
     8    0:2:0  0  0   0     0
     9    0:2:1  0  0   0     0
    10    0:2:2  0  0   0     0
    11    0:2:3  0  0   0     0
    12    0:3:0  0  0   0     0
    13    0:3:1  0  0   0     0
    14    0:3:2  0  0   0     0
    15    0:3:3  0  0   0     0
    16    1:0:0  0  0   4     4
    17    1:0:1  0  0   4     4
    18    1:0:2  0  0   4     4
    19    1:0:3  0  0   4     4
    20    1:1:0  0  0   4     4
    21    1:1:1  0  0   4     4
    [...]
    46    2:3:2  0  0   4     4
    47    2:3:3  0  0   4     4
    --------------------------
    48    total  0  0 128   128
```

For the example above, if only a single volume is specified you would get the same result. That is because the four virtual volumes are interleaved across the same logical disks and share the same chunklets.

The following example displays the distribution of space for a single virtual volume (multi.0):

```
cli% showvvpd multi.0
```

The following example displays information for a specific column:

```
cli% showvvpd -sortcol 5,dec junk
```

showvvset

Syntax

```
showvvset [options] [<setname_or_pattern>...]
showvvset -vv [options] [<vvname_or_pattern>...]
```

Description

The `showvvset` command lists the Virtual Volume(VV) sets defined on the HPE 3PAR storage system and their members.

Options

`-d`

Shows a more detailed listing of each set.

`-vv`

Shows VV sets that contain the supplied vvnames or patterns.

`-summary`

Shows VV sets with summarized output with VV sets names and number of VVs in those sets

Specifiers

`<setname_or_pattern>...`

An optional list of setnames or patterns. If no setnames or patterns are specified all sets are displayed, otherwise only sets with names matching one or more of the setnames or patterns are displayed.

The patterns are glob-style patterns. (For more information, issue "clihelp sub,globpat" or see **Glob-style patterns**.)

`<vvname_or_pattern>...`

Specifies that the sets containing virtual volumes with the specified names or matching the glob-style patterns should be displayed.

Authority

Any role in the system

Usage

- A Domain column may be included by using the **`setclienv`** command to set `listdom` to 1. The `listdom` option is also set to 1 if the CLI was started with the `-listdom` option, or if the CLI was started with the TPDLISTDOM environment variable set. Run `cli -h` and `setclienv -h` for details of the environment variables.
- Total line in the output specifies the total number of unique objects.

Examples

Use the following to show all VV sets defined to the system:

```
cli% showvvset
    Id Name   Members
     0 oravv  oravv.0
              oravv.1
              oravv.2
              oravv.3
              oravv.4
```

```
              oravv.5
              oravv.6
              oravv.7
              oravv.8
              oravv.9
    20 sia-1  test
              ttpvv.rw
              test-sv
              bar
    --------------------
     2 total  14
```

Use the following to show the sia-1 set only, with detail:

```
cli% showvvset -d sia-1
    20 sia-1 test     This set has a comment
             ttpvv.rw
             test-sv
             bar
    ----------------------------------------
     1 total 4
```

Use the following to show VV sets containing VVs matching test*:

```
cli% showvvset -vv test*
    Id Name  Members
    20 sia-1 test
             ttpvv.rw
             test-sv
             bar
    -----------------
     1 total 4
```

Use the following to show VV sets with summarized output:

```
cli% showvvset -summary
    Id Name  VV_Cnt
     0  oravv 10
    20  sia-1  4
    ---------------
     2  total 14
```

The following total line shows the unique number of objects:

```
cli% showvvset
    Id Name   Members
     0 oravv  oravv.0
              oravv.1
    20 sia-1  oravv.1
    --------------------
     2 total  2
```

Use the following to get a description of each column in the above examples:

```
clihelp -col showvvset
```

showwsapi

Syntax

```
showwsapi [option]
```

Description

Displays the WSAPI server service configuration state as either Enabled or Disabled. It displays the server current running status as Active, Inactive or Error. It also displays the current status of the HTTP and HTTPS ports and their port numbers. WSAPI server URL is also displayed.

Options

-d

Shows WSAPI information in table format.

Specifiers

None.

Authority

Any role in the system

Usage

If the version column displays `Unknown`, the Web Services API server might have restarted and was not not ready to accept queries.

Examples

The following shows the current Web Services API server status:

```
 cli% showwsapi
-Service-  -State- -HTTP_State- HTTP_Port -HTTPS_State- HTTPS_Port -Version- ---API_URL-----
Enabled     Active    Enabled    8008       Enabled       8080       1.5.2   https://name.example.com:8080/api/v1
```

The following lists a description of each column in the above examples:

```
clihelp -col showwsapi
```

showwsapisession

Syntax

```
showwsapisession [options] [{<id>|<pattern>}...]
```

Description

The `showwsapisession` command displays the WSAPI server connection information for the sessions. This includes the id, node, username, role, hostname, and IP Address of the connecting client. It also displays the session creation time and session type.

With no options specified, the display defaults to the following columns:

```
Id         Node     Name     Role     Client_IP_Addr     Connected_since
State      Session_Type
```

Options

-d

Shows WSAPI sessions information in table format. In addition to the default fields mentioned above, the display includes the node hostname. If session is of type Event, the event filter is also included.

Use this option only in conjunction with a single `<id>`.

-node <node_or_pattern>

Shows only records that match one or more of the `<node_or_pattern>`.

-role <role_or_pattern>

Shows only records that match one or more of the `<role_or_pattern>`.

-name <username_or_pattern>

Shows only records that match one or more of the <username_or_pattern>.

-ip <ip_or_pattern>

Shows only records that match one or more of the <ip_or_pattern>.

-t <t_or_pattern>

Shows only records that match one or more of the session connection time <t_or_pattern>.

-at <t_or_pattern>

Shows only records that match one or more of the session last access time <t_or_pattern>.

-type <Regular|Event>

Shows only records that match one or more of the session type. If type is Regular, the session can be used for regular WSAPI POST/PUT/GET/DELETE operations. If the type is Event, the session can be used only for server sent events operation.

-filter

Shows the filters associated with server sent event sessions. This option can be used only in conjunction with -type Event.

Specifiers

<id> <pattern>

Specifies the storage system process Id or pattern for the WSAPI session.

Authority

Any role in the system

Usage

- For all the <pattern> fields, the patterns are glob-style (shell-style) patterns. (For more information, enter clihelp sub,globpat or see **Glob-style patterns**.)
- The only exception is when -d option is provided. In this case, the session id is not a glob-style pattern and is checked against an exact match.
- The Web Services API server sessions information is only available if the WSAPI server is ready to accept queries.

Examples

The following shows the Web Services API server sessions:

```
cli% showwsapisession
      Id        Node -Name- -Role- -Client_IP_Addr- ----Connected_since---- -State- -Session_Type-
 75828955236970    0 root   super  192.168.28.191   2016-10-07 01:29:15 PDT Active  Regular
 75828782349741    0 root   super  192.168.28.191   2016-10-07 01:26:22 PDT Idle    Event
------------------------------------------------------------------------------
      2 total

cli% showwsapisession -type Event -filter
    Id          -Name- -Client_IP_Addr- -Session_Type- ---------------------Filters----------------------
1487738476616430   root 192.168.29.194   Event          category=lifecycle,alert:resource=volume,vlun,port
1487738461554474   root 192.168.29.194   Event          category=lifecycle
1487842498778244   root 192.168.29.195   Event          --
-------------------------------------------------------------------------------------------------------------
      3  total
```

Use the following to get descriptions of each column in the above examples:

```
clihelp -col showwsapisession
```

Shutdown Commands

shutdownnode

Syntax

```
shutdownnode halt|reboot|check <node_ID>
```

Description

The `shutdownnode` command shuts down a system node.

Specifiers

`<node_ID>`

Specifies the node, identified by its ID, to be shut down.

Subcommands

`halt`

Specifies halting the nodes after shutdown.

`reboot`

Specifies restarting the nodes after shutdown.

`restart`

Specifies restarting the storage services.

Authority

- Super, Service
- Any role granted the `node_shutdown` right

Usage

- Requires access to all domains.
- The system manager executes a set of validation checks before proceeding with the shutdown.
- Unless indicated otherwise, if any of the following conditions exist, the shutdown operation will not proceed:
  - The system checks for interrupting connectivity to various volumes and returns an error if multipathing is configured incorrectly.
  - System software upgrade is in progress.
  - Target node is not online.
  - If the system is processing tasks, the command returns a warning message to inform the user that tasks are running, and that the shutdown operation can cause some tasks to fail. If the user confirms the shutdown operation, the specified node reboots even if tasks are running.

    If no tasks are running when the initial checks are performed, and a new task starts afterward, the shutdown fails.
  - Any other node is online but not yet integrated into the cluster.
  - Another shutdown node operation is already in progress.
  - Shutdown node operation will result in the system shutting down due to loss of quorum.
  - One or more orphaned logical disks exist on the system that cannot be preserved.

- One or more admin logical disks cannot be reset, resulting in the kernel being unable to access meta data from those logical disks.
- One or more data (user or snap) logical disks cannot be reset, causing their associated VLUNs to become inaccessible to host applications.

Examples

Reboot node 0:

```
cli% shutdownnode reboot 0

-----System Information------
System Name   : ma1tp001
System Model  : HPE_3PAR 8400
Serial Number : 4UW0001212

------------------Partial List Of Hosts--------------------
Id Name        Persona        -WWN/iSCSI_Name- Port  IP_addr
0 sunx4150-01 Generic-legacy 2101001B32343495 0:5:1 n/a
1 sunx4150-02 Generic-legacy 2101001B32344695 0:5:1 n/a
2 sunx4150-03 Generic-legacy 2100001B32147595 0:5:1 n/a
3 sunx4150-04 Generic-legacy 2101001B32341C95 0:5:1 n/a
4 sunx4150-05 Generic-legacy 2101001B32342495 0:5:1 n/a
5 sunx4150-06 Generic-legacy 2100001B321A6E82 0:5:1 n/a
6 sunx4150-07 Generic-legacy 2101001B3231EE79 0:5:1 n/a
7 sunx4150-08 Generic-legacy 2101001B323BC44C 0:5:1 n/a
8 sunx4150-09 Generic-legacy 2101001B323BFF4C 0:5:1 n/a
9 sunx4150-10 Generic-legacy 2101001B323B0A4C 0:5:1 n/a

There are active tasks on this Storage System.
It is recommended that these tasks are either allowed to finish
or manually cancelled before continuing.
Continuing with this REBOOT operation may cause some of these tasks to fail.

Shutting down this node will impact applications running on the system displayed above

Do you REALLY want to REBOOT this node? yes or no: yes

The node will REBOOT in 15 seconds

Press the enter key to stop...

The node will REBOOT in 12 seconds

Press the enter key to stop...

The node will REBOOT in 9 seconds

Press the enter key to stop...

The node will REBOOT in 6 seconds

Press the enter key to stop...

The node will REBOOT in 3 seconds

Press the enter key to stop...
```

shutdownsys

Syntax

```
shutdownsys halt|reboot
```

Description

The `shutdownsys` command shuts down an entire system.

Options

None.

Specifiers

None.

Subcommands

`halt`

Specifies that the system should be halted after shutdown. If this subcommand is not specified, the reboot must be used.

`reboot`

Specifies that the system should be restarted after shutdown. If this subcommand is not given, the halt subcommand must be used.

Authority

- Super, Service
- Any role granted the `sys_shutdown` right

Usage

- Access to all domains is required to run this command.
- The execution of `shutdownsys` command can affect service.
- Hence, a confirmation is required before proceeding with this command.
- After the `shutdownsys` command is issued, there is no indication from the CLI that the shutdown is occurring. You can issue the **`showsys`** command to display the current status of the system during the initial stage of the shutdown process and after the system has fully restarted.
- If the node that was running on the system manager fails or if the system manager process exits while executing the `shutdownsys` command, the shutdown will not complete. The only safe action is to reissue the `shutdownsys` command.
- Do not issue any commands other than **`showsys`** while the system is shutting down.

Examples

The following example shows how to reboot the system:

```
cli% shutdownsys reboot

-----System Information------
System Name   : ma1tp001
System Model  : HPE_3PAR 8400
Serial Number : 4UW0001212

-------------------Partial List Of Hosts--------------------
Id Name         Persona        -WWN/iSCSI_Name- Port  IP_addr
0 sunx4150-01 Generic-legacy 2101001B32343495 6:5:1 n/a
1 sunx4150-02 Generic-legacy 2101001B32344695 6:5:1 n/a
2 sunx4150-03 Generic-legacy 2100001B32147595 6:5:1 n/a
3 sunx4150-04 Generic-legacy 2101001B32341C95 6:5:1 n/a
4 sunx4150-05 Generic-legacy 2101001B32342495 6:5:1 n/a
5 sunx4150-06 Generic-legacy 2100001B321A6E82 6:5:1 n/a
6 sunx4150-07 Generic-legacy 2101001B3231EE79 6:5:1 n/a
7 sunx4150-08 Generic-legacy 2101001B323BC44C 6:5:1 n/a
8 sunx4150-09 Generic-legacy 2101001B323BFF4C 6:5:1 n/a
9 sunx4150-10 Generic-legacy 2101001B323B0A4C 6:5:1 n/a

Shutting down this storage system will impact applications running on the hosts displayed above
```

```
Do you REALLY want to REBOOT this storage system? yes or no: yes
The system will REBOOT in 15 seconds
Press the enter key to stop...
The system will REBOOT in 12 seconds
Press the enter key to stop...
The system will REBOOT in 9 seconds
Press the enter key to stop...
The system will REBOOT in 6 seconds
Press the enter key to stop...
The system will REBOOT in 3 seconds
Press the enter key to stop...
The system will REBOOT NOW !!!
```

sraomoves

Syntax

```
sraomoves [options]
```

Description

The `sraomoves` command shows the space that AO has moved between tiers.

Options

`-btsecs <secs>`

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h, or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h), or days (e.g. -7d).

  If it is not specified, then the time at which the report begins is 12 hours ago.

  If -btsecs 0 is specified, then the report begins at the earliest sample.

`-etsecs <secs>`

Selects the end time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h, or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h), or days (e.g. -7d).

  If it is not specified, then the report ends with the most recent sample.

-oneline

Shows data in simplified format with one line per AOCFG.

-vv <VV_name|VV_set|pattern>[,<VV_name|VV_set|pattern>...]

Limits the analysis to VVs with names that match one or more of the specified names or glob-style patterns. VV set names must be prefixed by "set:". Note that snapshot VVs will not be considered since only base VVs have region space.

-withvv

Shows the data for each VV.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

None.

Authority

Any role in the system

Usage

- For this command, 1MB = 1048576 bytes.
- The space is reported is the space that AO has requested to be moved and the actual space moved may be less, for example if the move failed.

Examples

The following example shows space moved by AO beginning 2 hours (7200 seconds) ago:

```
cli% sraomoves -btsecs -7200
```

Use the following to get a description of each column in the above examples:

```
clihelp -col sraomoves
```

srcpgspace

Syntax

srcpgspace [options] [{*<CPG_name>*|*<pattern>*}...]

Description

Displays historical space data reports for common provisioning groups (CPGs).

Options

-attime

Performance is shown at a particular time interval, specified by the `-etsecs` option, with one row per object group described by the `-groupby` option. Without this option performance is shown versus time, with a row per time interval.

-btsecs *<secs>*

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins is 12 hours ago.

  If `-btsecs 0` is specified then the report begins at the earliest sample.

-etsecs <secs>

Selects the end time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.

- Specifying an etsecs value in the future, prefixed by "+", will provide a forecast of the space statistics. Either absolute or relative time may be specified, such as "+2017-03-24" or "+2d".

-hires

Selects high resolution samples (5 minute intervals) for the report. This is the default.

-hourly

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-summary <keyword>[,<keyword>]...

Summarize performance across requested objects and time range.

One of these summary keywords must be included:

- min—Display the minimum for each metric
- avg—Display the average for each metric
- max—Display the maximum for each metric
- <N>%—Display percentile for each metric. <N> may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

- detail

  Displays individual performance records in addition to one or more summaries. By default, -summary output excludes individual records and only displays the summary.

- per_time

  When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per time. By default, one summary is computed across all records.

- per_group

  When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per object grouping. By default, one summary is computed across all records.

- only_compareby

  When requesting data limited to certain object groupings with the -compareby option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all records and ignore the limitation of the -compareby option, though the "detail" output does conform to the -compareby object limitation.

-groupby <groupby>[,<groupby>...]

For -attime reports, generate a separate row for each combination of <groupby> items. Each <groupby> must be different and one of the following:

- DOM_NAME—Domain name
- CPGID—Common Provisioning Group ID
- CPG_NAME—Common Provisioning Group name
- DISK_TYPE—The disk type of the PDs used by the CPG
- RAID_TYPE—The RAID type of the CPG

-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>

The -compareby option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following:Shared_MB, Private_Base_MB, Private_Snap_MB, Free_MB, Total_Space_MB, Total_Used_MB, Growth_Avail_MB, Dedup_GC_KBPS, Compr_Ratio, Compact_Ratio, Dedup_Ratio, Data_Reduce

-disk_type <type>[,<type>...]

Limits the data to disks of the types specified. Allowed types are:

- FC—Fast Class
- NL—Nearline
- SSD—Solid State Drive
- SCM - Storage Class Memory Drive

-raid_type <type>[,<type>...]

Limits the data to RAID of the specified types. Allowed types are 0, 1, 5 and 6.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

<CPG_name>|<pattern>

CPGs matching either the specified CPG_name or glob-style pattern are included. This specifier can be repeated to display information for multiple CPGs. If not specified, all CPGs are included.

Authority

Any role in the system

Usage

- The Growth space available for each CPG is only sampled every hour and is shown as 0 for high resolution samples that are not hourly samples.
- For this command, 1MB = 1048576 bytes.

Examples

Displaying aggregate hourly CPG space information for CPGs with names that match the pattern "fc*" beginning 24 hours ago:

```
cli% srcpgspace -hourly -btsecs -24h fc*
```

Displaying CPG space information for the most recent daily sample aggregated by the disk type and raid type for CPGs with names that match the pattern "fc*".

```
cli% srcpgspace -daily -attime -groupby disk_type,raid_type fc*
```

srhistld

Syntax

```
srhistld [options] [{<LD_name>|<pattern>}...]
```

Description

The `srhistld` command displays historical histogram performance data reports for logical disks.

Options

`-attime`

Performance is shown at a particular time interval, specified by the -etsecs option, with one row per object group described by the -groupby option. Without this option performance is shown versus time, with a row per time interval.

`-btsecs <secs>`

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins is 12 hours ago.

  If -btsecs 0 is specified then the report begins at the earliest sample.

`-etsecs <secs>`

Selects the end time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h, or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.

`-hires`

Selects high resolution samples (5 minute intervals) for the report.

This is the default.

`-hourly`

Selects hourly samples for the report.

`-daily`

Selects daily samples for the report.

-metric both|time|size

Selects which metric to display. Metrics can be one of the following:

- both—(Default) Display both I/O time and I/O size histograms
- time—Display only the I/O time histogram
- size—Display only the I/O size histogram

-timecols <fcol> <lcol>

For the I/O time histogram, shows the columns from the first column <fcol> through last column <lcol>. The available columns range from 0 through 31.

The first column (<fcol>) must be a value greater than or equal to 0, but less than the value of the last column (<lcol>).

The last column (<lcol>) must be less than or equal to 31.

The first column includes all data accumulated for columns less than the first column and the last column includes accumulated data for all columns greater than the last column.

The default value of <fcol> is 6.

The default value of <lcol> is 15.

-sizecols <fcol> <lcol>

For the I/O size histogram, shows the columns from the first column (<fcol>) through the last column (<lcol>). Available columns range from 0 through 15.

The first column (<fcol>) must be a value greater than or equal to 0, but less than the value of the last column (<lcol>) (default value of 3).

The last column (<lcol>) must be less than or equal to 15 (default value of 11).

The default value of <fcol> is 3.

The default value of <lcol> is 11.

-pct

Shows the access count in each bucket as a percentage. If this option is not specified, the histogram shows the access counts.

-rw

Specifies that the display includes separate read and write data. If not specified, the total is displayed.

-groupby <groupby>[,<groupby>...]

For -attime reports, generate a separate row for each combination of <groupby> items. Each <groupby> must be different and one of the following:

- DOM_NAME—Domain name
- LDID—Logical disk ID
- LD_NAME—Logical disk name
- CPG_NAME—Common Provisioning Group name
- NODE—The node that owns the LD

-cpg <CPG_name|pattern>[,<CPG_name|pattern>...]

Limits the data to LDs in CPGs with names that match one or more of the specified names or glob-style patterns.

-node <node>[,<node>...]

Limits the data to that corresponding to one of the specified nodes.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

<LD_name>|<pattern>

LDs matching either the specified LD_name or glob-style pattern are included. This specifier can be repeated to display information for multiple LDs. If not specified, all LDs are included.

Authority

Any role in the system

Usage

For this command KB = 1000 bytes.

Examples

The following example displays aggregate hourly histogram performance statistics for all logical disks beginning 24 hours ago:

```
cli% srhistld -hourly -btsecs -24h
```

srhistpd

Syntax

```
srhistpd [options] [{<PDID>|<pattern>}...]
```

Description

Displays historical histogram performance data reports for physical disks.

Options

-attime

Performance is shown at a particular time interval, specified by the -etsecs option, with one row per object group described by the -groupby option. Without this option performance is shown versus time, with a row per time interval.

-btsecs *<secs>*

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"

    - Date string: "2012-10-26" or 2012-10-26
    - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified, then the time at which the report begins is 12 hours ago.

  If -btsecs 0 is specified, then the report begins at the earliest sample.

-etsecs *<secs>*

Selects the end time in seconds for the report. The value can be specified as either.

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
    - Full time string including time zone: "2012-10-26 11:00:00 PDT"
    - Full time string excluding time zone: "2012-10-26 11:00:00"
    - Date string: "2012-10-26" or 2012-10-26
    - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h, or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h), or days (e.g. -7d).

  If it is not specified, then the report ends with the most recent sample.

-hires

Selects high resolution samples (5 minute intervals) for the report.

This is the default.

-hourly

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-metric both|time|size

Selects which metric to display. Metrics can be one of the following:

- both—(Default) Display both I/O time and I/O size histograms
- time—Display only the I/O time histogram
- size—Display only the I/O size histogram

-timecols *<fcol> <lcol>*

For the I/O time histogram, shows the columns from the first column <fcol> through last column <lcol>. The available columns range from 0 through 31.

The first column (<fcol>) must be a value greater than or equal to 0, but less than the value of the last column (<lcol>).

The last column (<lcol>) must be less than or equal to 31.

The first column includes all data accumulated for columns less than the first column and the last column includes accumulated data for all columns greater than the last column.

The default value of *<fcol>* is 6.

The default value of *<lcol>* is 15.

-sizecols *<fcol>* *<lcol>*

For the I/O size histogram, shows the columns from the first column (*<fcol>*) through the last column (*<lcol>*). Available columns range from 0 through 15.

The first column (*<fcol>*) must be a value greater than or equal to 0, but less than the value of the last column (*<lcol>*) (default value of 3).

The last column (*<lcol>*) must be less than or equal to 15 (default value of 11).

The default value of *<fcol>* is 3.

The default value of *<lcol>* is 11.

-pct

Shows the access count in each bucket as a percentage. If this option is not specified, the histogram shows the access counts.

-rw

Specifies that the display includes separate read and write data. If not specified, the total is displayed.

-groupby *<groupby>*[,*<groupby>*...]

For -attime reports, generate a separate row for each combination of <groupby> items. Each <groupby> must be different and one of the following:

- PDID—Physical disk ID
- PORT_N—The node number for the primary port for the PD
- PORT_S—The PCI slot number for the primary port for the PD
- PORT_P—The port number for the primary port for the PD
- DISK_TYPE—The disktype of the PD
- SPEED—The speed of the PD

-disk_type *<type>*[,*<type>*...]

Limits the data to disks of the types specified. Allowed types are:

- FC—Fast Class
- NL—Nearline
- SSD—Solid State Drive
- SCM—Storage Class Memory Drive

-rpm <speed>[,<speed>...]

Limits the data to disks of the specified RPM. Allowed speeds are 7, 10, 15, 100, 150, and 1500.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

<PDID>|<pattern>

PDs with IDs that match either the specified PDID or glob-style pattern is included. This specifier can be repeated to include multiple PDIDs or patterns. If not specified, all PDs are included.

Authority

Any role in the system

Usage

For this command KB = 1000 bytes.

Examples

The following example displays aggregate hourly histogram performance statistics for all physical disks beginning 24 hours ago:

```
cli% srhistpd -hourly -btsecs -24h
```

srhistport

Syntax

```
srhistport [options] [{<npat>:<spat>:<ppat>} ...]
```

Description

The `srhistport` command displays historical histogram performance data reports for ports.

Options

-attime

Performance is shown at a particular time interval, specified by the -etsecs option, with one row per object group described by the -groupby option. Without this option performance is shown versus time, with a row per time interval.

-btsecs <secs>

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"

- Date string: "2012-10-26" or 2012-10-26
- Time string: "11:00:00" or 11:00:00

- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins is 12 hours ago.

  If -btsecs 0 is specified then the report begins at the earliest sample.

-etsecs <secs>

Selects the end time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.

-hires

Selects high resolution samples (5 minute intervals) for the report.

This is the default.

-hourly

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-metric both|time|size

Selects which metric to display. Metrics can be one of the following:

- both—(Default)Display both I/O time and I/O size histograms
- time—Display only the I/O time histogram
- size—Display only the I/O size histogram

-timecols <fcol> <lcol>

For the I/O time histogram, shows the columns from the first column <fcol> through last column <lcol>. The available columns range from 0 through 31.

The first column (<fcol>) must be a value greater than or equal to 0, but less than the value of the last column (<lcol>).

The last column (<lcol>) must be less than or equal to 31.

The first column includes all data accumulated for columns less than the first column and the last column includes accumulated data for all columns greater than the last column.

The default value of <fcol> is 6.

The default value of <lcol> is 15.

-sizecols <fcol> <lcol>

For the I/O size histogram, shows the columns from the first column (<fcol>) through the last column (<lcol>). Available columns range from 0 through 15.

The first column (<fcol>) must be a value greater than or equal to 0, but less than the value of the last column (<lcol>) (default value of 3).

The last column (<lcol>) must be less than or equal to 15 (default value of 11).

The default value of <fcol> is 3.

The default value of <lcol> is 11.

-pct

Shows the access count in each bucket as a percentage. If this option is not specified, the histogram shows the access counts.

-rw

Specifies that the display includes separate read and write data. If not specified, the total is displayed.

-groupby <groupby>[,<groupby>...]

For -attime reports, generate a separate row for each combination of <groupby> items. Each <groupby> must be different and one of the following:

- PORT_N—The node number for the port
- PORT_S—The PCI slot number for the port
- PORT_P—The port number for the port
- PORT_TYPE—The type of the port
- GBITPS—The speed of the port
- TRANS_TYPE—The transaction type - ctl or data

-both | -ctl | -data

Specifies that both control and data transfers are displayed(-both), only control transfers are displayed (-ctl), or only data transfers are displayed (-data). If this option is not specified, only data transfers are displayed.

-port_type <type>[,<type>...]

Limits the data to port of the types specified. Allowed types are:

- disk—Disk port
- host—Host Fibre channel port
- iscsi—Host ISCSI port
- free—Unused port
- fs—File Persona port
- peer—Data Migration FC port

- rcip—Remote copy IP port
- rcfc—Remote copy FC port

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

<npat>:<spat>:<ppat>

Ports with <port_n>:<port_s>:<port_p> that match any of the specified <npat>:<spat>:<ppat> patterns are included, where each of the patterns is a glob-style pattern. This specifier can be repeated to include multiple ports or patterns. If not specified, all ports are included.

Authority

Any role in the system

Usage

For this command KB = 1000 bytes.

Examples

The following example displays aggregate hourly histogram performance statistics for disk and host ports on nodes 0 and 1 beginning 24 hours ago:

```
cli% srhistport -hourly -btsecs -24h -port_type host,disk 0:*:* 1:*:*
```

srhistvlun

Syntax

```
srhistvlun [options]
```

Description

The `srhistvlun` command displays historical histogram performance data reports for VLUNs.

Options

-attime

Performance is shown at a particular time interval, specified by the -etsecs option, with one row per object group described by the -groupby option. Without this option performance is shown versus time, with a row per time interval.

-btsecs <secs>

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:

- Full time string including time zone: "2012-10-26 11:00:00 PDT"
- Full time string excluding time zone: "2012-10-26 11:00:00"
- Date string: "2012-10-26" or 2012-10-26
- Time string: "11:00:00" or 11:00:00

- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins is 12 hours ago.

  If -btsecs 0 is specified then the report begins at the earliest sample.

-etsecs <secs>

Selects the end time in seconds for the report. The value can be specified as either

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.

-hires

Selects high resolution samples (5 minute intervals) for the report.

This is the default.

-hourly

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-metric both|time|size

Selects which metric to display. Metrics can be one of the following:

- both—(Default)Display both I/O time and I/O size histograms
- time—Display only the I/O time histogram
- size—Display only the I/O size histogram

-timecols <fcol> <lcol>

For the I/O time histogram, shows the columns from the first column <fcol> through last column <lcol>. The available columns range from 0 through 31.

The first column (<fcol>) must be a value greater than or equal to 0, but less than the value of the last column (<lcol>).

The last column (<lcol>) must be less than or equal to 31.

The first column includes all data accumulated for columns less than the first column and the last column includes accumulated data for all columns greater than the last column.

The default value of <fcol> is 6.

The default value of <lcol> is 15.

-sizecols <fcol> <lcol>

For the I/O size histogram, shows the columns from the first column (<fcol>) through the last column (<lcol>). Available columns range from 0 through 15.

The first column (<fcol>) must be a value greater than or equal to 0, but less than the value of the last column (<lcol>) (default value of 3).

The last column (<lcol>) must be less than or equal to 15 (default value of 11).

The default value of <fcol> is 3.

The default value of <lcol> is 11.

-pct

Shows the access count in each bucket as a percentage. If this option is not specified, the histogram shows the access counts.

-rw

Specifies that the display includes separate read and write data. If not specified, the total is displayed.

-groupby <groupby>[,<groupby>...]

For -attime reports, generate a separate row for each combination of <groupby> items. Each <groupby> must be different and one of the following:

- DOM_NAME—Domain name
- VV_NAME—Virtual volume name
- HOST_NAME—Host name
- LUN—The LUN number for the VLUN
- HOST_WWN—The host WWN for the VLUN
- PORT_N—The node number for the VLUN port
- PORT_S—The PCI slot number for the VLUN port
- PORT_P—The port number for the VLUN port
- VVSET_NAME—Virtual volume set name
- HOSTSET_NAME—Host set name
- VM_NAME—Virtual Machine Name for vVol based VMs
- VM_ID—Virtual Machine Identification number for vVol based VMs
- VM_HOST—Virtual Machine host for vVol based VMs
- VVOLSC—Virtual Volume Storage Container for vVol based VMs

-host <host_name|host_set|host_wwn|pattern>[,<host_name|host_set|host_wwn|pattern>...]

Limits the data to hosts with names that match one or more of the specified names or glob-style patterns. Host set name must start with "set:" and can also include patterns. To specify the host by WWN, start with "wwn:". A WWN can also include glob-style patterns.

-vv <VV_name|VV_set|pattern>[,<VV_name|VV_set|pattern>...]

Limits the data to VVs with names that match one or more of the specified names or glob-style patterns. VV set name must be prefixed by "set:" and can also include patterns.

-l <LUN|pattern>[,<LUN|pattern>...]

Limits the data to LUNs that match one or more of the specified LUNs or glob-style patterns.

-port <npat>:<spat>:<ppat>[,<npat>:<spat>:<ppat>...]

Ports with <port_n>:<port_s>:<port_p> that match any of the specified <npat>:<spat>:<ppat> patterns are included, where each of the patterns is a glob-style pattern. If not specified, all ports are included.

-vlun <host>:<vv>[:<lun>:<port_n>:<port_s>:<port_p>]
[,<host>:<vv>[:<lun>:<port_n>:<port_s>:<port_p>]]...

Limits the data to VLUNs matching the specified combination of host, VV, lun, and port. Each of these components in this option may be a glob-style pattern. The host and VV components may specify a corresponding object set by prefixing "set:" to the component. The host component may specify a WWN by prefixing the component with "wwn:". The lun and port components are optional, and if not present, data will be filtered to any matching combination of host and VV. This option cannot be combined with -host, -vv, -l, or -port.

-vmname {<VM_name>|<pattern>}[,{<VM_name>|<pattern>}]

Limits the data to VMs that match one or more of the specified VM names or glob-styled patterns for vVol based VMs.

-vmid {<VM_ID>|<pattern>}[,{<VM_ID>|<pattern>}]

Limits the data to VMs that match one or more of the specified VM IDs or glob-styled patterns for vVol based VMs.

-vmhost {<VM_host_name>|<pattern>}[,{<VM_host_name>|<pattern>}]

Limits the data to VMs that match one or more of the specified VM host names or glob-styled patterns for vVol based VMs.

-vvolsc {<vVol_container_name>|<pattern>}[,{<VVol_container_name>|<pattern>}...]

Limits the data to vVol containers that match one or more of the specified vVol container names or glob-styled patterns.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

None.

Authority

Any role in the system

Usage

- For this command KB = 1000 bytes.
- Because VVs can belong to multiple VV sets, requesting data for more than one VV set can include duplicate data for any VVs with multiple membership.

Examples

The following example displays aggregate hourly histogram performance statistics for all VLUNs beginning 24 hours ago:

```
cli% srhistvlun -hourly -btsecs -24h
```

The following example displays VV or host sets can be specified with patterns:

```
cli% srhistvlun -btsecs -2h -host set:hostset -vv set:vvset*
```

srldspace

Syntax

`srldspace [options] [{<LD_name>|<pattern>}...]`

Description

The `srldspace` command displays historical space data reports for logical disks (LDs).

Options

`-attime`

Performance is shown at a particular time interval, specified by the `-etsecs` option, with one row per object group described by the `-groupby` option. Without this option performance is shown versus time, with a row per time interval.

`-btsecs <secs>`

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins is 12 hours ago.

  If `-btsecs 0` is specified then the report begins at the earliest sample.

`-etsecs <secs>`

Selects the end time in seconds for the report. The value can be specified as either

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.
- Specifying an etsecs value in the future, prefixed by "+", will provide a forecast of the space statistics. Either absolute or relative time may be specified, such as "+2017-03-24" or "+2d".

`-hires`

Selects high resolution samples (5 minute intervals) for the report.

This is the default.

`-hourly`

Selects hourly samples for the report.

`-daily`

Selects daily samples for the report.

`-summary <keyword>[,<keyword>]...`

Summarizes performance across requested objects and time range.

One of these summary keywords must be included:

- `min`—Display the minimum for each metric
- `avg`—Display the average for each metric
- `max`—Display the maximum for each metric
- `<N>%`—Display percentile for each metric. `<N>` may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

`detail`

Displays individual performance records in addition to one or more summaries. By default, `-summary` output excludes individual records and only displays the summary.

`per_time`

When requesting data across multiple points in time (vstime) and multiple object groupings (`-groupby`) compute summaries per time. By default, one summary is computed across all records.

`per_group`

When requesting data across multiple points in time (vstime) and multiple object groupings (`-groupby`) compute summaries per object grouping. By default, one summary is computed across all records.

only_compareby

When requesting data limited to certain object groupings with the -compareby option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all records and ignore the limitation of the -compareby option, though the "detail" output does conform to the -compareby object limitation.

-groupby <groupby>[,<groupby>...]

For -attime reports, generate a separate row for each combination of <groupby> items. Each <groupby> must be different and one of the following:

- DOM_NAME—Domain name
- CPG_NAME—Common Provisioning Group name
- LDID—Logical disk ID
- LD_NAME—Logical disk name
- DISK_TYPE—The disktype of the PDs used by the LD
- RAID_TYPE—The RAID type of the LD
- SET_SIZE—The RAID set size of the LD
- STEP_SIZE—The RAID step size of the LD
- ROW_SIZE—The RAID row size of the LD
- OWNER—The owner node for the LD

-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>

The -compareby option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following: Raw_MB, Used_MB, Free_MB, Total_MB

-cpg <CPG_name|pattern>[,<CPG_name|pattern>...]

Limits the data to LDs in CPGs with names that match one or more of the specified names or glob-style patterns.

-disk_type <type>[,<type>...]

Limits the data to disks of the types specified. Allowed types are:

- FC—Fast Class
- NL—Nearline
- SSD—Solid State Drive
- SCM—Storage Class Memory Drive

-raid_type <type>[,<type>...]

Limits the data to RAID of the specified types. Allowed types are 0, 1, 5 and 6.

-owner <node>[,<node>...]

Limits data to LDs owned by the specified nodes.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

<LD_name>|<pattern>

LDs matching either the specified LD_name or glob-style pattern are included. This specifier can be repeated to display information for multiple LDs. If not specified, all LDs are included.

Authority

Any role in the system

Usage

For this command, 1MB = 1048576 bytes.

Examples

The following example displays aggregate hourly LD space information for all RAID 5 LDs with names that match either "fc*" or "ssd*" patterns beginning 24 hours ago:

```
cli% srldspace -raid_type 5 -hourly -btsecs -24h fc* ssd*
```

The following example displays LD space information for the most recent daily sample aggregated by the LD name for LDs with names that match the pattern "tp*".

```
cli% srldspace -daily -attime -groupby ld_name tp*
```

srpdspace

Syntax

```
srpdspace [options] [{<PDID>|<pattern>}...]
```

Description

The srpdspace command displays historical space data reports for physical disks (PDs).

Options

-attime

Performance is shown at a particular time interval, specified by the -etsecs option, with one row per object group described by the -groupby option. Without this option performance is shown versus time, with a row per time interval.

-btsecs <secs>

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:

- Full time string including time zone: "2012-10-26 11:00:00 PDT"
- Full time string excluding time zone: "2012-10-26 11:00:00"
- Date string: "2012-10-26" or 2012-10-26
- Time string: "11:00:00" or 11:00:00

- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h, or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h), or days (e.g. -7d).

  If it is not specified, then the time at which the report begins is 12 hours ago.

  If `-btsecs` 0 is specified then the report begins at the earliest sample.

-etsecs <secs>

Selects the end time in seconds for the report. The value can be specified as either.

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h, or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h), or days (e.g. -7d).

  If it is not specified, then the report ends with the most recent sample.
- Specifying an etsecs value in the future, prefixed by "+", will provide a forecast of the space statistics. Either absolute or relative time may be specified, such as "+2017-03-24" or "+2d".

-hires

Selects high resolution samples (5 minute intervals) for the report.

This is the default.

-hourly

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-summary <keyword>[,<keyword>]...

Summarize performance across requested objects and time range.

One of these summary keywords must be included:

- `min`—Display the minimum for each metric
- `avg`—Display the average for each metric

- max—Display the maximum for each metric
- <N>%—Display percentile for each metric. <N> may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

detail

Displays individual performance records in addition to one or more summaries. By default, -summary output excludes individual records and only displays the summary.

per_time

When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per time. By default, one summary is computed across all records.

per_group

When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per object grouping. By default, one summary is computed across all records.

only_compareby

When requesting data limited to certain object groupings with the -compareby option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all records and ignore the limitation of the -compareby option, though the "detail" output does conform to the -compareby object limitation.

-groupby <groupby>[,<groupby>...]

For -attime reports, generate a separate row for each combination of <groupby> items. Each <groupby> must be different and one of the following:

- PDID—Physical disk ID
- CAGEID—Cage ID
- CAGESIDE—Cage Side
- MAG—Disk Magazine number within the cage
- DISK—Disk position within the magazine
- DISK_TYPE—The disk type of the PD
- SPEED—The disk speed

-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>

The compareby option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following: Nrm_Used_OK, Nrm_Used_Fail, Nrm_Avail_Clean, Nrm_Avail_Dirty, Nrm_Avail_Fail, Spr_Used_OK, Spr_Used_Fail, Spr_Avail_Clean, Spr_Avail_Dirty, Spr_Avail_Fail, LifeLeft_Pct, T_DegC

The following fields can only be used with the -capacity option: alloc_cap_MB, free_cap_MB, failed_cap_MB, total_cap_MB

-disk_type <type>[,<type>...]

Limits the data to disks of the types specified. Allowed types are:

- FC—Fast Class
- NL—Nearline
- SSD—Solid State Drive
- SCM—Storage Class Memory Drive

-capacity

Displays disk contributions to the system capacity categories: Allocated, Free, Failed, and Total

-rpm <speed>[,<speed>...]

Limits the data to disks of the specified RPM. Allowed speeds are 7, 10, 15, 100, 150, and 1500.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

<PDID>|<pattern>

PDs with IDs that match either the specified PDID or glob-style pattern is included. This specifier can be repeated to include multiple PDIDs or patterns. If not specified, all PDs are included.

Authority

Any role in the system

Usage

- The LifeLeft% and the T(C) metrics are expensive to sample and are only sampled for the hourly and daily resolutions.
- For this command, 1MB = 1048576 bytes.
- The temperature and life left values are calculated using cached values, so they are delayed by several minutes from the indicated time.

Examples

The following example displays aggregate hourly PD space information for all PDs beginning 24 hours ago:

```
cli% srpdspace -hourly -btsecs -24h
```

The following example displays current system capacity values of SSD PDs:

```
cli% srpdspace -capacity -attime -disk_type SSD
```

The following example displays system capacity values of all PDs over the last 2 hours. Since these capacity values are aggregated at each time sample and all PDs are included, this represents the entire system capacity equivalent to showsys:

```
cli% srpdspace -capacity -btsecs -2h
```

srrgiodensity

Syntax

```
srrgiodensity [options] -cpg {<cpg_name>|<pattern>}...
srrgiodensity [options] <aocfg_name>
```

Description

The `srrgiodensity` command shows the distribution of IOP/s intensity for Logical Disk (LD) regions for a common provisioning group (CPG) or Adaptive Optimization (AO) configuration. For a single CPG, this can be used to see whether AO can be effectively used. For an AO configuration the command shows how AO has moved regions between tiers.

Options

`-btsecs <secs>`

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins is 12 hours ago.

  If -btsecs 0 is specified then the report begins at the earliest sample.

`-etsecs <secs>`

Selects the end time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.

`-cmult 2|4|8`

Selects the step between histogram columns of the report. By default each column's IO density is 4 times the previous column, but a step of 2 or 8 can also be specified.

-cpg

Treats the specifiers as CPG names or glob-style patterns.

-vv <VV_name|VV_set|pattern>[,<VV_name|VV_set|pattern>...]

Limits the analysis to VVs with names that match one or more of the specified names or glob-style patterns. VV set names must be prefixed by "set:". Note that snapshot VVs will not be considered since only base VVs have region space.

-cumul

Shows data as cumulative including all the columns to the right.

-pct

Shows data as a percentage per row.

-totpct

Shows data as a totaled percentage across an AOCFG.

-withvv

Shows the data for each VV.

-rw

Specifies that the display includes separate read and write data. If not specified, the total is displayed.

Specifiers

<aocfg_name>

The name of the AO configuration to generate the report for.

<cpg_name>|<pattern>

If the -cpg option is specified, the report is generated for all the CPGs with names that match one or more CPG names or glob-style patterns.

Authority

Any role in the system

Usage

When the -rw option is combined with either the -pct or -totpct option, the read, write and total categories each compute separate percentages.

Examples

The following example shows region IO performance density report for a CPG beginning 2 hours (7200 seconds) ago:

```
cli% srrgiodensity -btsecs -7200 -cpg cpg_name
```

The following example will show region IO performance density report using the -rw, -withvv -btsecs and -etsecs options:

```
cli% srrgiodensity -btsecs "2015-05-15 10:30:00 PDT" -etsecs "2015-05-15 12:30:00 PDT"
Start Time: 2015-05-15 10:30:00 PDT (1431711000)
End Time: 2015-05-15 12:30:00 PDT (1431718200)
```

srstatcache

Syntax

```
srstatcache [options] [<node>...]
```

Description

The `srstatcache` command displays historical performance data reports for flash cache and data cache.

Options

`-attime`

Performance is shown at a particular time interval, specified by the `-etsecs` option, with one row per object group described by the `-groupby` option. Without this option performance is shown versus time, with a row per time interval.

`-btsecs <secs>`

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins is 12 hours ago.

  If `-btsecs 0` is specified then the report begins at the earliest sample.

`-etsecs <secs>`

Selects the end time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.

`-hires`

Selects high resolution samples (5 minute intervals) for the report.

This is the default.

`-hourly`

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-summary <keyword>[,<keyword>]...

Summarize performance across requested objects and time range.

One of these summary keywords must be included:

- min—Display the minimum for each metric
- avg—Display the average for each metric
- max—Display the maximum for each metric
- <N>%—Display percentile for each metric. <N> may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

detail

Displays individual performance records in addition to one or more summaries. By default, -summary output excludes individual records and only displays the summary.

per_time

When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per time. By default, one summary is computed across all records.

per_group

When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per object grouping. By default, one summary is computed across all records.

only_compareby

When requesting data limited to certain object groupings with the -compareby option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all records and ignore the limitation of the -compareby option, though the "detail" output does conform to the -compareby object limitation.

-groupby <groupby>[,<groupby>...]

For -attime reports, generate a separate row for each combination of <groupby> items. Each <groupby> must be different and one of the following:

NODE—The controller node

-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>

The compareby option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following:

reads_per_sec, writes_per_sec, cmp_read_hit_pct, cmp_write_hit_pct, fmp_read_hit_pct, fmp_write_hit_pct, read_back_iops, read_back_mbps, destage_write_iops, destage_write_mbps, fmp_used_pct, fmp_q_dormant, fmp_q_cold, fmp_q_norm, fmp_q_warm, fmp_q_hot, fmp_q_destage, fmp_q_read, fmp_q_flush, fmp_q_wrtback, cmp_q_free, cmp_q_clean, cmp_q_write1, cmp_q_writeN, cmp_q_wrtsched, cmp_q_writing, cmp_q_dcowpend, cmp_q_dcowproc

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

By default, the output of the command lists read and write access rates as well as cache (CMP) and flash cache (FMP) read and write hit rates. One of the following options can be used to increase what the command shows.

-internal_flashcache

Lists the IOPS and bandwidth each for both read back and destaged write flash cache activity. May be combined with -fmp_queue and -cmp_queue.

-fmp_queue

Lists the FMP queue statistics. May be combined with -cmp_queue and -internal_flashcache.

-cmp_queue

Lists the CMP queue statistics. May be combined with -fmp_queue and -internal_flashcache.

-full

Lists all the metrics for each row in a single line. The output for this option is very wide.

Specifiers

<node>

Only the specified node numbers are included, where each node is a number from 0 through 7. This specifier can be repeated to display information for multiple nodes. If not specified, all nodes are included.

Authority

Any role in the system

Usage

None.

Examples

The following example displays aggregate hourly performance statistics for flash cache and data cache beginning 24 hours ago:

```
cli% srstatcache -hourly -btsecs -24h
```

The following example displays daily flash cache and data cache performance aggregated by node for nodes 0 and 1:

```
cli% srstatcache -daily -attime -groupby node 0 1
```

srstatcmp

Syntax

```
srstatcmp [options] [<node>...]
```

Description

The srstatcmp command displays historical performance data reports for cache memory.

Options

-attime

Performance is shown at a particular time interval, specified by the `-etsecs` option, with one row per object group described by the `-groupby` option. Without this option performance is shown versus time, with a row per time interval.

-btsecs <secs>

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins is 12 hours ago.

  If `-btsecs 0` is specified then the report begins at the earliest sample.

-etsecs <secs>

Selects the end time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.

-hires

Selects high resolution samples (5 minute intervals) for the report.

This is the default.

-hourly

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-summary <keyword>[,<keyword>]...

Summarize performance across requested objects and time range.

One of these summary keywords must be included:

- min—Display the minimum for each metric
- avg—Display the average for each metric
- max—Display the maximum for each metric
- <N>%—Display percentile for each metric. <N> may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

detail

Displays individual performance records in addition to one or more summaries. By default, -summary output excludes individual records and only displays the summary.

per_time

When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per time. By default, one summary is computed across all records.

per_group

When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per object grouping. By default, one summary is computed across all records.

only_compareby

When requesting data limited to certain object groupings with the -compareby option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all records and ignore the limitation of the -compareby option, though the "detail" output does conform to the -compareby object limitation.

-groupby <groupby>[,<groupby>...]

For -attime reports, generate a separate row for each combination of <groupby> items. Each <groupby> must be different and one of the following:

NODE—The controller node

-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>

The -compareby option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following:

rd_hit_ps, rd_miss_ps, rd_ps, rd_hit_perc, wr_hit_ps, wr_miss_ps, wr_ps, wr_hit_perc, accesses_ps, lockblk_ps, delack_nl_7_ps, delack_fc_ps, delack_ssd_ps, pg_free, pg_clean, pg_write1, pg_writen, pg_wrtsched, pg_writing, pg_recov, cm_dirty_nl_7, cm_dirty_fc, cm_dirty_ssd, cm_max_nl_7, cm_max_fc, cm_max_ssd

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

By default, the output of the command lists cache read and write hit rates, access rates as well as rates for locked blocks and delayed write acknowledgments. One of the following options can be used to change what the command shows.

-full

Lists all the metrics for each row in a single line. The output for this option is very wide.

-page

Lists the page state information.

Specifiers

<node>

Only the specified node numbers are included, where each node is a number from 0 through 7. This specifier can be repeated to display information for multiple nodes. If not specified, all nodes are included.

Authority

Any role in the system

Usage

None.

Examples

The following example displays aggregate hourly performance statistics for all node caches beginning 24 hours ago:

```
cli% srstatcmp -hourly -btsecs -24h
```

The following example displays daily node cache performance aggregated by node for nodes 0 and 1:

```
cli% srstatcmp -daily -attime -groupby node 0 1
```

srstatcpu

Syntax

```
srstatcpu [options] [<node>...]
```

Description

The `srstatcpu` command displays historical performance data reports for CPUs.

Options

-attime

Performance is shown at a particular time interval, specified by the `-etsecs` option, with one row per object group described by the `-groupby` option. Without this option performance is shown versus time, with a row per time interval.

-btsecs <secs>

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins is 12 hours ago.

  If `-btsecs 0` is specified then the report begins at the earliest sample.

-etsecs <secs>

Selects the end time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.

-hires

Selects high resolution samples (5 minute intervals) for the report.

This is the default.

-hourly

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-summary <keyword>[,<keyword>]...

Summarize performance across requested objects and time range.

One of these summary keywords must be included:

- `min`—Display the minimum for each metric
- `avg`—Display the average for each metric

- max—Display the maximum for each metric
- <N>%—Display percentile for each metric. <N> may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

detail

Displays individual performance records in addition to one or more summaries. By default, –summary output excludes individual records and only displays the summary.

per_time

When requesting data across multiple points in time (vstime) and multiple object groupings (–groupby) compute summaries per time. By default, one summary is computed across all records.

per_group

When requesting data across multiple points in time (vstime) and multiple object groupings (–groupby) compute summaries per object grouping. By default, one summary is computed across all records.

only_compareby

When requesting data limited to certain object groupings with the –compareby option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all records and ignore the limitation of the –compareby option, though the "detail" output does conform to the –compareby object limitation.

-groupby <groupby>[,<groupby>...]

For –attime reports, generate a separate row for each combination of <groupby> items. Each <groupby> must be different and one of the following:

- NODE—The controller node
- CPU—The CPU within the controller node

-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>

The –compareby option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following: user_pct, sys_pct, idle_pct, intr_ps, ctxt_ps

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

<node>

Only the specified node numbers are included, where each node is a number from 0 through 7. This specifier can be repeated to display information for multiple nodes. If not specified, all nodes are included.

Authority

Any role in the system

Usage

None.

Examples

The following example displays aggregate hourly performance statistics for all CPUs beginning 24 hours ago:

```
cli% srstatcpu -hourly -btsecs -24h
```

The following example displays daily node cpu performance aggregated by node for nodes 0 and 1:

```
cli% srstatcpu -daily -attime -groupby node 0 1
```

srstatfsav

Syntax

```
srstatfsav [options] [{<fpg_name>|<pattern}]
```

Description

The `srstatfsav` command displays historical performance data reports for File Persona anti-virus activity.

Options

`-attime`

Performance is shown at a particular time interval, specified by the `-etsecs` option, with one row per object group described by the `-groupby` option. Without this option performance is shown versus time, with a row per time interval.

`-btsecs <secs>`

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins is 12 hours ago.

  If `-btsecs 0` is specified then the report begins at the earliest sample.

`-etsecs <secs>`

Selects the end time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:

- Full time string including time zone: "2012-10-26 11:00:00 PDT"
- Full time string excluding time zone: "2012-10-26 11:00:00"
- Date string: "2012-10-26" or 2012-10-26
- Time string: "11:00:00" or 11:00:00

- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.

-hires

Selects high resolution samples (5 minute intervals) for the report.

This is the default.

-hourly

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-summary <keyword>[,<keyword>]...

Summarizes performance across requested objects and time range.

One of these summary keywords must be included:

- min—Display the minimum for each metric
- avg—Display the average for each metric
- max—Display the maximum for each metric
- <N>%—Display percentile for each metric. <N> may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

detail

Displays individual performance records in addition to one or more summaries. By default, -summary output excludes individual records and only displays the summary.

per_time

When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per time. By default, one summary is computed across all records.

per_group

When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per object grouping. By default, one summary is computed across all records.

only_compareby

When requesting data limited to certain object groupings with the -compareby option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all records and ignore the limitation of the -compareby option, though the "detail" output does conform to the -compareby object limitation.

-groupby <groupby>[,<groupby>...]

For -attime reports, generate a separate row for each combination of <groupby> items. Each <groupby> must be different and one of the following:

NODE—The controller node

-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>

The compareby option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following: scanengine, maxscanengine, totalscanned, totalinfected, totalquarantined

-node <node>[,<node>...]

Limits the data to that corresponding to one of the specified nodes.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

None.

Authority

Any role in the system

Usage

None.

Examples

The following example displays aggregate hourly performance statistics for all File Persona anti-virus activity beginning 24 hours ago:

```
cli% srstatfsav -hourly -btsecs -24h
```

The following example displays daily File Persona anti-virus performance aggregated by node for nodes 0 and 1:

```
cli% srstatfsav -daily -attime -groupby NODE -node 0,1
```

srstatfsblock

Syntax

```
srstatfsblock [options] [{<blockdev_name>|<pattern>}...]
```

Description

The srstatfsblock command displays historical performance data reports for File Persona block devices.

Options

-attime

Performance is shown at a particular time interval, specified by the `-etsecs` option, with one row per object group described by the `-groupby` option. Without this option performance is shown versus time, with a row per time interval.

-btsecs <secs>

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins is 12 hours ago.

  If `-btsecs 0` is specified then the report begins at the earliest sample.

-etsecs <secs>

Selects the end time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.

-hires

Selects high resolution samples (5 minute intervals) for the report.

This is the default.

-hourly

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-summary <keyword>[,<keyword>]...

Summarize performance across requested objects and time range.

One of these summary keywords must be included:

- `min`—Display the minimum for each metric
- `avg`—Display the average for each metric
- `max`—Display the maximum for each metric
- `<N>%`—Display percentile for each metric. `<N>` may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

detail

Displays individual performance records in addition to one or more summaries. By default, `-summary` output excludes individual records and only displays the summary.

per_time

When requesting data across multiple points in time (vstime) and multiple object groupings (`-groupby`) compute summaries per time. By default, one summary is computed across all records.

per_group

When requesting data across multiple points in time (vstime) and multiple object groupings (`-groupby`) compute summaries per object grouping. By default, one summary is computed across all records.

only_compareby

When requesting data limited to certain object groupings with the `-compareby` option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all records and ignore the limitation of the `-compareby` option, though the "detail" output does conform to the `-compareby` object limitation.

-groupby <groupby>[,<groupby>...]

For `-attime` reports, generate a separate row for each combination of `<groupby>` items. Each `<groupby>` must be different and one of the following:

- `NODE`—The controller node
- `BLOCKDEV_NAME`—The block device name

-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>

The `-compareby` option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following: `reads, reads_merged, read_sectors, read_time_ms, writes, writes_merged, write_sectors, write_time_ms, ios_current, io_time_ms, io_time_weighted_ms`

-node <node>[,<node>...]

Limits the data to that corresponding to one of the specified nodes.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (`<col>`). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (`<dir>`) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

<blockdev_name>|<pattern>

Block Devices matching either the specified name or glob-style pattern are included. This specifier can be repeated to display information for multiple devices. If not specified, all block devices are included.

Authority

Any role in the system

Usage

For this command KB = 1000 bytes.

Examples

The following example displays aggregate hourly performance statistics for all File Persona block devices beginning 24 hours ago:

```
cli% srstatfsblock -hourly -btsecs -24h
```

The following example displays daily File Persona block device performance aggregated by node for nodes 0 and 1:

```
cli% srstatfsblock -daily -attime -groupby NODE -node 0,1
```

srstatfscpu

Syntax

```
srstatfscpu [options] [<cpu_id>...]
```

Description

The `srstatfscpu` command displays historical performance data reports for File Persona CPU utilization.

Options

-attime

Performance is shown at a particular time interval, specified by the `-etsecs` option, with one row per object group described by the `-groupby` option. Without this option performance is shown versus time, with a row per time interval.

-btsecs <secs>

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"

- Date string: "2012-10-26" or 2012-10-26
- Time string: "11:00:00" or 11:00:00

- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins is 12 hours ago.

  If -btsecs 0 is specified then the report begins at the earliest sample.

-etsecs <secs>

Selects the end time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.

-hires

Selects high resolution samples (5 minute intervals) for the report.

This is the default.

-hourly

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-summary <keyword>[,<keyword>]...

Summarize performance across requested objects and time range.

One of these summary keywords must be included:

- min—Display the minimum for each metric
- avg—Display the average for each metric
- max—Display the maximum for each metric
- <N>%—Display percentile for each metric. <N> may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

detail

Displays individual performance records in addition to one or more summaries. By default, `-summary` output excludes individual records and only displays the summary.

per_time

When requesting data across multiple points in time (vstime) and multiple object groupings (`-groupby`) compute summaries per time. By default, one summary is computed across all records.

per_group

When requesting data across multiple points in time (vstime) and multiple object groupings (`-groupby`) compute summaries per object grouping. By default, one summary is computed across all records.

only_compareby

When requesting data limited to certain object groupings with the `-compareby` option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all records and ignore the limitation of the `-compareby` option, though the "detail" output does conform to the `-compareby` object limitation.

-groupby <groupby>[,<groupby>...]

For `-attime` reports, generate a separate row for each combination of `<groupby>` items. Each `<groupby>` must be different and one of the following:

- `NODE`—The controller node
- `CPU`—The CPU within the controller node

-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>

The `-compareby` option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following: `usage_pct, iowait_pct, idle_pct`

-node <node>[,<node>...]

Limits the data to that corresponding to one of the specified nodes.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (`<col>`). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (`<dir>`) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

<cpu_id>...

Only the specified CPU ID numbers are included. This specifier can be repeated to display information for multiple CPUs. If not specified, all CPUs are included.

Authority

Any role in the system

Usage

For this command KB = 1000 bytes.

Examples

The following example displays aggregate hourly performance statistics for all File Persona CPU activity beginning 24 hours ago:

```
cli% srstatfscpu -hourly -btsecs -24h
```

The following example displays daily File Persona CPU performance aggregated by node for nodes 0 and 1:

```
cli% srstatfscpu -daily -attime -groupby NODE -node 0,1
```

srstatfsfpg

Syntax

```
srstatfsfpg [options] [{<fpg_name>|<pattern>}...]
```

Description

The `srstatfsfpg` command displays historical performance data reports for File Persona file provisioning groups.

Options

`-attime`

Performance is shown at a particular time interval, specified by the `-etsecs` option, with one row per object group described by the `-groupby` option. Without this option performance is shown versus time, with a row per time interval.

`-btsecs <secs>`

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins is 12 hours ago.

  If `-btsecs 0` is specified then the report begins at the earliest sample.

`-etsecs <secs>`

Selects the end time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:

- Full time string including time zone: "2012-10-26 11:00:00 PDT"
- Full time string excluding time zone: "2012-10-26 11:00:00"
- Date string: "2012-10-26" or 2012-10-26
- Time string: "11:00:00" or 11:00:00

- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.

-hires

Selects high resolution samples (5 minute intervals) for the report.

This is the default.

-hourly

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-summary <keyword>[,<keyword>]...

Summarize performance across requested objects and time range.

One of these summary keywords must be included:

- min—Display the minimum for each metric
- avg—Display the average for each metric
- max—Display the maximum for each metric
- <N>%—Display percentile for each metric. <N> may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

detail

Displays individual performance records in addition to one or more summaries. By default, -summary output excludes individual records and only displays the summary.

per_time

When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per time. By default, one summary is computed across all records.

per_group

When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per object grouping. By default, one summary is computed across all records.

only_compareby

When requesting data limited to certain object groupings with the -compareby option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all records and ignore the limitation of the -compareby option, though the "detail" output does conform to the -compareby object limitation.

-groupby <groupby>[,<groupby>...]

For -attime reports, generate a separate row for each combination of <groupby> items. Each <groupby> must be different and one of the following:

- FPG_NAME—File Provisioning Group name
- FPG_ID—File Provisioning Group ID
- NODE—The controller node

-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>

The -compareby option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following:

```
Totalblocks, Freeblocks, Numreads, Numbytesread, Numwrites,
NumBytesWritten, Creates, Removes, Errors, ReadLatency, WriteLatency
```

-node <node>[,<node>...]

Limits the data to that corresponding to one of the specified nodes.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

<fpg_name>|<pattern>

File provisioning groups matching either the specified name or glob-style pattern are included. This specifier can be repeated to display information for multiple FPGs. If not specified, all FPGs are included.

Authority

Any role in the system

Usage

For this command KB = 1000 bytes.

Examples

The following example displays aggregate hourly performance statistics for all File Persona FPGs beginning 24 hours ago:

```
cli% srstatfsfpg -hourly -btsecs -24h
```

The following example displays daily FPG performance aggregated by node for nodes 0 and 1:

```
cli% srstatfsfpg -daily -attime -groupby NODE -node 0,1
```

srstatfsmem

Syntax

```
srstatfsmem [options]
```

Description

The `srstatfsmem` command displays historical performance data reports for File Persona memory utilization.

Options

`-attime`

Performance is shown at a particular time interval, specified by the `-etsecs` option, with one row per object group described by the `-groupby` option. Without this option performance is shown versus time, with a row per time interval.

`-btsecs <secs>`

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins is 12 hours ago.

  If `-btsecs 0` is specified then the report begins at the earliest sample.

`-etsecs <secs>`

Selects the end time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.

`-hires`

Selects high resolution samples (5 minute intervals) for the report.

This is the default.

`-hourly`

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-summary <keyword>[,<keyword>]...

Summarize performance across requested objects and time range.

One of these summary keywords must be included:

- min—Display the minimum for each metric
- avg—Display the average for each metric
- max—Display the maximum for each metric
- <N>%—Display percentile for each metric. <N> may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

detail

Displays individual performance records in addition to one or more summaries. By default, -summary output excludes individual records and only displays the summary.

per_time

When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per time. By default, one summary is computed across all records.

per_group

When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per object grouping. By default, one summary is computed across all records.

only_compareby

When requesting data limited to certain object groupings with the -compareby option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all records and ignore the limitation of the -compareby option, though the "detail" output does conform to the -compareby object limitation.

-groupby <groupby>[,<groupby>...]

For -attime reports, generate a separate row for each combination of <groupby> items. Each <groupby> must be different and one of the following:

NODE—The controller node

-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>

The -compareby option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following: usage_pct, swap_pct, free_pct

-node <node>[,<node>...]

Limits the data to that corresponding to one of the specified nodes.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

None.

Authority

Any role in the system

Usage

For this command KB = 1000 bytes.

Examples

The following example displays aggregate hourly performance statistics for all File Persona memory utilization beginning 24 hours ago:

```
cli% srstatfsmem -hourly -btsecs -24h
```

The following example displays daily File Persona memory performance aggregated by node for nodes 0 and 1:

```
cli% srstatfsmem -daily -attime -groupby NODE -node 0,1
```

srstatfsnet

Syntax

```
srstatfsnet [options] [{<ethdev_name>|<pattern>}...]
```

Description

The `srstatfsnet` command displays historical performance data reports for File Persona networking devices.

Options

`-attime`

Performance is shown at a particular time interval, specified by the `-etsecs` option, with one row per object group described by the `-groupby` option. Without this option performance is shown versus time, with a row per time interval.

`-btsecs <secs>`

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"

- Date string: "2012-10-26" or 2012-10-26
- Time string: "11:00:00" or 11:00:00

- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins is 12 hours ago.

  If `-btsecs 0` is specified then the report begins at the earliest sample.

-etsecs <secs>

Selects the end time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.

-hires

Selects high resolution samples (5 minute intervals) for the report.

This is the default.

-hourly

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-summary <keyword>[,<keyword>]...

Summarize performance across requested objects and time range.

One of these summary keywords must be included:

- min—Display the minimum for each metric
- avg—Display the average for each metric
- max—Display the maximum for each metric
- <N>%—Display percentile for each metric. <N> may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

detail

Displays individual performance records in addition to one or more summaries. By default, `-summary` output excludes individual records and only displays the summary.

per_time

When requesting data across multiple points in time (vstime) and multiple object groupings (`-groupby`) compute summaries per time. By default, one summary is computed across all records.

per_group

When requesting data across multiple points in time (vstime) and multiple object groupings (`-groupby`) compute summaries per object grouping. By default, one summary is computed across all records.

only_compareby

When requesting data limited to certain object groupings with the `-compareby` option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all records and ignore the limitation of the `-compareby` option, though the "detail" output does conform to the `-compareby` object limitation.

-groupby <groupby>[,<groupby>...]

For `-attime` reports, generate a separate row for each combination of `<groupby>` items. Each `<groupby>` must be different and one of the following:

- `NODE`—The controller node
- `DEV_NAME`—Ethernet interface name

-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>

The `-compareby` option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following: `rx_bytes, rx_packets, tx_bytes, tx_packets`

-node <node>[,<node>...]

Limits the data to that corresponding to one of the specified nodes.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (`<col>`). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (`<dir>`) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

<ethdev_name><|pattern>

Ethernet interface devices matching either the specified name or glob-style pattern are included. This specifier can be repeated to display information for multiple devices. If not specified, all devices are included.

Authority

Any role in the system

Usage

For this command KB = 1000 bytes.

Examples

The following example displays aggregate hourly performance statistics for all File Persona ethernet devices beginning 24 hours ago:

```
cli% srstatfsnet -hourly -btsecs -24h
```

The following example displays daily File Persona ethernet performance aggregated by node for nodes 0 and 1:

```
cli% srstatfsnet -daily -attime -groupby NODE -node 0,1
```

srstatfsnfs

Syntax

```
srstatfsnfs [options]
```

Description

The `srstatfsnfs` command displays historical performance data reports for File Persona NFS shares.

Options

`-attime`

Performance is shown at a particular time interval, specified by the `-etsecs` option, with one row per object group described by the `-groupby` option. Without this option performance is shown versus time, with a row per time interval.

`-btsecs <secs>`

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins is 12 hours ago.

  If `-btsecs 0` is specified then the report begins at the earliest sample.

`-etsecs <secs>`

Selects the end time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:

- Full time string including time zone: "2012-10-26 11:00:00 PDT"
- Full time string excluding time zone: "2012-10-26 11:00:00"
- Date string: "2012-10-26" or 2012-10-26
- Time string: "11:00:00" or 11:00:00

- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.

-hires

Selects high resolution samples (5 minute intervals) for the report.

This is the default.

-hourly

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-summary <keyword>[,<keyword>]...

Summarize performance across requested objects and time range.

One of these summary keywords must be included:

- min—Display the minimum for each metric
- avg—Display the average for each metric
- max—Display the maximum for each metric
- <N>%—Display percentile for each metric. <N> may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

detail

Displays individual performance records in addition to one or more summaries. By default, -summary output excludes individual records and only displays the summary.

per_time

When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per time. By default, one summary is computed across all records.

per_group

When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per object grouping. By default, one summary is computed across all records.

only_compareby

When requesting data limited to certain object groupings with the -compareby option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all records and ignore the limitation of the -compareby option, though the "detail" output does conform to the -compareby object limitation.

-groupby <groupby>[,<groupby>...]

For -attime reports, generate a separate row for each combination of <groupby> items. Each <groupby> must be different and one of the following:

NODE—The controller node

-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>

The -compareby option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following:

```
Client_RPC_calls, Client_RPC_retrans, Server_RPC_calls,
Server_RPC_badcalls, V3_Null, V3_GetAttr, V3_SetAttr, V3_lookup, V3_access,
V3_ReadLink, V3_Read, V3_Write, V3_Create, V3_MkDir, V3_Symlink, V3_Mknod,
V3_Remove, V3_RmDir, V3_Rename, V3_Link, V3_ReadDir, V3_ReadDirPlus,
V3_FsStat, V3_FsInfo, V3_PathConf, V3_Commit, V4_op0_unused, V4_op1_unused,
V4_op2_future, V4_access, V4_close, V4_commit, V4_create, V4_delegpurge,
V4_delegreturn, V4_getattr, V4_getfh, V4_link, V4_lock, V4_lockt, V4_locku,
V4_lookup, V4_lookup_root, V4_nverify, V4_open, V4_openattr, V4_open_conf,
V4_open_dgrd, V4_putfh, V4_putpubfh, V4_putrootfh, V4_Read, V4_reddir,
V4_readlink, V4_remove, V4_rename, V4_renew, V4_restorefh, V4_savefh,
V4_secinfo, V4_setattr, V4_setcltid, V4_setcltidconf, V4_verify, V4_Write,
V4_rellockowner, V4_bc_ctl, V4_bind_conn, V4_exchange_id, V4_create_ses,
V4_destroy_ses, V4_free_stateid, V4_getdirdeleg, V4_getdevinfo,
V4_getdevlist, V4_layoutcommit, V4_layoutget, V4_layoutreturn,
V4_secinfononam, V4_sequence, V4_set_ssv, V4_test_stateid, V4_want_deleg,
V4_destroy_clid, V4_reclaim_comp
```

-node <node>[,<node>...]

Limits the data to that corresponding to one of the specified nodes.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

None.

Authority

Any role in the system

Usage

For this command KB = 1000 bytes.

Examples

The following example displays aggregate hourly performance statistics for all File Persona NFS share activity beginning 24 hours ago:

```
cli% srstatfsnfs -hourly -btsecs -24h
```

The following example displays daily File Persona NFS share performance aggregated by node for nodes 0 and 1:

```
cli% srstatfsnfs -daily -attime -groupby NODE -node 0,1
```

srstatfssmb

Syntax

```
srstatfssmb [options]
```

Description

The `srstatfssmb` command displays historical performance data reports for File Persona SMB shares.

Options

`-attime`

Performance is shown at a particular time interval, specified by the `-etsecs` option, with one row per object group described by the `-groupby` option. Without this option performance is shown versus time, with a row per time interval.

`-btsecs <secs>`

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins is 12 hours ago.

  If `-btsecs 0` is specified then the report begins at the earliest sample.

`-etsecs <secs>`

Selects the end time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

If it is not specified then the report ends with the most recent sample.

-hires

Selects high resolution samples (5 minute intervals) for the report.

This is the default.

-hourly

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-summary <keyword>[,<keyword>]...

Summarize performance across requested objects and time range.

One of these summary keywords must be included:

- `min`—Display the minimum for each metric
- `avg`—Display the average for each metric
- `max`—Display the maximum for each metric
- `<N>%`—Display percentile for each metric. `<N>` may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

detail

Displays individual performance records in addition to one or more summaries. By default, `-summary` output excludes individual records and only displays the summary.

per_time

When requesting data across multiple points in time (vstime) and multiple object groupings (`-groupby`) compute summaries per time. By default, one summary is computed across all records.

per_group

When requesting data across multiple points in time (vstime) and multiple object groupings (`-groupby`) compute summaries per object grouping. By default, one summary is computed across all records.

only_compareby

When requesting data limited to certain object groupings with the `-compareby` option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all records and ignore the limitation of the `-compareby` option, though the "detail" output does conform to the `-compareby` object limitation.

-groupby <groupby>[,<groupby>...]

For -attime reports, generate a separate row for each combination of <groupby> items. Each <groupby> must be different and one of the following:

NODE—Statistics per node

-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>

The -compareby option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following:

connections, maxConnections, sessions, maxSessions, treeConnects, maxTreeConnects, openFiles, maxOpenFiles, ReadSumRecorded, ReadSampleRecorded, WriteSumRecorded, WriteSampleRecorded

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

None.

Authority

Any role in the system

Usage

None.

Examples

The following example displays aggregate hourly performance statistics for all File Persona SMB share activity beginning 24 hours ago:

```
cli% srstatfssmb -hourly -btsecs -24h
```

The following example displays daily File Persona SMB share performance aggregated by node for nodes 0 and 1:

```
cli% srstatfssmb -daily -attime -groupby NODE -node 0,1
```

srstatfssnapshot

Syntax

```
srstatfssnapshot [options]
```

Description

The srstatfssnapshot command displays historical performance data reports for File Persona snapshots.

Options

-attime

Performance is shown at a particular time interval, specified by the `-etsecs` option, with one row per object group described by the `-groupby` option. Without this option performance is shown versus time, with a row per time interval.

-btsecs <secs>

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins is 12 hours ago.

  If `-btsecs 0` is specified then the report begins at the earliest sample.

-etsecs <secs>

Selects the end time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.

-hires

Selects high resolution samples (5 minute intervals) for the report.

This is the default.

-hourly

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-summary <keyword>[,<keyword>]...

Summarize performance across requested objects and time range.

One of these summary keywords must be included:

- min—Display the minimum for each metric
- avg—Display the average for each metric
- max—Display the maximum for each metric
- <N>%—Display percentile for each metric. <N> may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

detail

Displays individual performance records in addition to one or more summaries. By default, -summary output excludes individual records and only displays the summary.

per_time

When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per time. By default, one summary is computed across all records.

per_group

When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per object grouping. By default, one summary is computed across all records.

only_compareby

When requesting data limited to certain object groupings with the -compareby option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all records and ignore the limitation of the -compareby option, though the "detail" output does conform to the -compareby object limitation.

-groupby <groupby>[,<groupby>...]

For -attime reports, generate a separate row for each combination of <groupby> items. Each <groupby> must be different and one of the following:

NODE—The controller node

-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>

The -compareby option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following: numredirectonwrite

-node <node>[,<node>...]

Limits the data to that corresponding to one of the specified nodes.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

None.

Authority

Any role in the system

Usage

None.

Examples

The following example displays aggregate hourly performance statistics for all File Persona snapshot activity beginning 24 hours ago:

```
cli% srstatfssnapshot -hourly -btsecs -24h
```

The following example displays daily File Persona snapshot performance aggregated by node for nodes 0 and 1:

```
cli% srstatfssnapshot -daily -attime -groupby NODE -node 0,1
```

srstatiscsi

Syntax

```
srstatiscsi [options] [{<npat>:<spat>:<ppat>} ...]
```

Description

The `srstatiscsi` command displays historical performance data reports for iSCSI ports.

Options

`-attime`

Performance is shown at a particular time interval, specified by the `-etsecs` option, with one row per object group described by the `-groupby` option. Without this option performance is shown versus time, with a row per time interval.

`-btsecs <secs>`

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

If it is not specified then the time at which the report begins depends on the sample category (-hires, -hourly, -daily):

- For hires, the default begin time is 12 hours ago (`-btsecs -12h`).
- For hourly, the default begin time is 7 days ago (`-btsecs -7d`).
- For daily, the default begin time is 90 days ago (`-btsecs -90d`).

If begin time and sample category are not specified then the time the report begins is 12 hours ago and the default sample category is hires. If `-btsecs 0` is specified then the report begins at the earliest sample.

`-etsecs <secs>`

Selects the end time in seconds for the report. If `-attime` is specified, select the time for the report.

The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.

`-hires`

Selects high resolution samples (5 minute intervals) for the report. This is the default.

`-hourly`

Selects hourly samples for the report.

`-daily`

Selects daily samples for the report.

`-summary <keyword>[,<keyword>]...`

Summarize performance across requested objects and time range.

One of these summary keywords must be included:

- `min`—Display the minimum for each metric
- `avg`—Display the average for each metric
- `max`—Display the maximum for each metric
- `<N>%`—Display percentile for each metric. `<N>` may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

detail

Displays individual performance records in addition to one or more summaries. By default, –summary output excludes individual records and only displays the summary.

per_time

When requesting data across multiple points in time (vstime) and multiple object groupings (–groupby) compute summaries per time. By default, one summary is computed across all records.

per_group

When requesting data across multiple points in time (vstime) and multiple object groupings (–groupby) compute summaries per object grouping. By default, one summary is computed across all records.

only_compareby

When requesting data limited to certain object groupings with the –compareby option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all records and ignore the limitation of the –compareby option, though the "detail" output does conform to the –compareby object limitation.

-groupby <groupby>[,<groupby>...]

For –attime reports, generate a separate row for each combination of <groupby> items. Each <groupby> must be different and one of the following:

- PORT_N—The node number for the port
- PORT_S—The PCI slot number for the port
- PORT_P—The port number for the port
- PROTOCOL—The protocol type for the port

-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>

The –compareby option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following:

rx_pkts, rx_kbps, tx_pkts, tx_kbps, tot_kbps, tot_pkts, errs

-protocol <protocol>[,<protocol>...]

Limits the data to protocol of the types specified. Allowed types are: Eth, IP, TCP and iSCSI.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

None.

Authority

Any role in the system

Usage

None.

Examples

The following example displays aggregate daily performance statistics for iSCSI ports on nodes 0 and 1 beginning 2 days ago:

```
cli% srstatiscsi -daily -btsecs -2d 0:*:* 1:*:*
```

srstatiscsisession

Syntax

```
srstatiscsisession [options] [{<npat>:<spat>:<ppat>} ...]
```

Description

The `srstatiscsisession` command displays historical performance data reports for iSCSI sessions.

Options

`-attime`

Performance is shown at a particular time interval, specified by the `-etsecs` option, with one row per object group described by the `-groupby` option. Without this option performance is shown versus time, with a row per time interval.

`-btsecs <secs>`

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins depends on the sample category (`-hires`, `-hourly`, `-daily`):
  - For hires, the default begin time is 12 hours ago (`-btsecs -12h`).
  - For hourly, the default begin time is 7 days ago (`-btsecs -7d`).
  - For daily, the default begin time is 90 days ago (`-btsecs -90d`).

  If begin time and sample category are not specified then the time the report begins is 12 hours ago and the default sample category is hires. If `-btsecs 0` is specified then the report begins at the earliest sample.

`-etsecs <secs>`

Selects the end time in seconds for the report. If `-attime` is specified, select the time for the report.

The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.

-hires

Selects high resolution samples (5 minute intervals) for the report. This is the default.

-hourly

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-summary <keyword>[,<keyword>]...

Summarize performance across requested objects and time range.

One of these summary keywords must be included:

- min—Display the minimum for each metric
- avg—Display the average for each metric
- max—Display the maximum for each metric
- <N>%—Display percentile for each metric. <N> may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

detail

Displays individual performance records in addition to one or more summaries. By default, -summary output excludes individual records and only displays the summary.

per_time

When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per time. By default, one summary is computed across all records.

per_group

When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per object grouping. By default, one summary is computed across all records.

only_compareby

When requesting data limited to certain object groupings with the -compareby option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all

records and ignore the limitation of the -compareby option, though the "detail" output does conform to the -compareby object limitation.

-groupby <groupby>[,<groupby>...]

For -attime reports, generate a separate row for each combination of <groupby> items. Each <groupby> must be different and one of the following:

- PORT_N—The node number for the port
- PORT_S—The PCI slot number for the port
- PORT_P—The port number for the port
- ISCSI_NAME—The iSCSI name for the session
- TPGT—The TPGT ID for the session

-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>

The -compareby option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following:

```
command_pdus_transferred, response_pdus_transferred, tot_pdus,
kbps_transmitted, kbps_received, tot_kbps, digest_errors,
connection_timeout_errors
```

-iscsi_name {<iSCSI_name>|<pattern>}[,{<iSCSI_name>|<pattern>...}]

Limits the data to iSCSI sessions that match one or more of the specified names or glob-styled patterns for iSCSI sessions.

-tpgt <tpgt_id>[,tpgt_id>...]

Limits the data to that corresponding to one of the specified Target Portal Group Tag IDs.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

None.

Authority

Any role in the system

Usage

None.

Examples

The following example displays aggregate daily performance statistics for iSCSI sessions on nodes 0 and 1 beginning 2 days ago:

```
cli% srstatiscsisession -daily -btsecs -2d 0:*:* 1:*:*
```

srstatld

Syntax

```
srstatld [options] [{<LD_name>|<pattern>}...]
```

Description

The `srstatld` command displays historical performance data reports for logical disks.

Options

`-attime`

Performance is shown at a particular time interval, specified by the `-etsecs` option, with one row per object group described by the `-groupby` option. Without this option performance is shown versus time, with a row per time interval.

`-btsecs <secs>`

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins is 12 hours ago.

  If `-btsecs 0` is specified then the report begins at the earliest sample.

`-etsecs <secs>`

Selects the end time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.

-hires

Selects high resolution samples (5 minute intervals) for the report.

This is the default.

-hourly

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-summary <keyword>[,<keyword>]...

Summarize performance across requested objects and time range.

One of these summary keywords must be included:

- min—Display the minimum for each metric
- avg—Display the average for each metric
- max—Display the maximum for each metric
- <N>%—Display percentile for each metric. <N> may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

detail

Displays individual performance records in addition to one or more summaries. By default, -summary output excludes individual records and only displays the summary.

per_time

When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per time. By default, one summary is computed across all records.

per_group

When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per object grouping. By default, one summary is computed across all records.

only_compareby

When requesting data limited to certain object groupings with the -compareby option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all records and ignore the limitation of the -compareby option, though the "detail" output does conform to the -compareby object limitation.

-groupby <groupby>[,<groupby>...]

For -attime reports, generate a separate row for each combination of <groupby> items. Each <groupby> must be different and one of the following:

- DOM_NAME—Domain name
- LDID—Logical disk ID
- LD_NAME—Logical disk name
- CPG_NAME—Common Provisioning Group name
- NODE—The node that owns the LD

-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>

The -compareby option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following:

read_iops, write_iops, total_iops, read_kbps, write_kbps, total_kbps, read_svctms, write_svctms, total_svctms, read_ioszkb, write_ioszkb, total_ioszkb, total_qlen, busy_pct

-cpg <CPG_name|pattern>[,<CPG_name|pattern>...]

Limits the data to LDs in CPGs with names that match one or more of the specified names or glob-style patterns.

-node <node>[,<node>...]

Limits the data to that corresponding to one of the specified nodes.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

<LD_name>|<pattern>

LDs matching either the specified LD_name or glob-style pattern are included. This specifier can be repeated to display information for multiple LDs. If not specified, all LDs are included.

Authority

Any role in the system

Usage

For this command KB = 1000 bytes.

Examples

The following example displays aggregate hourly performance statistics for all logical disks beginning 24 hours ago:

```
cli% srstatld -hourly -btsecs -24h
```

srstatlink

Syntax

```
srstatlink [options] [<node>...]
```

Description

The srstatlink command displays historical performance data reports for links (internode, PCI and cache memory).

Options

-attime

Performance is shown at a particular time interval, specified by the `-etsecs` option, with one row per object group described by the `-groupby` option. Without this option performance is shown versus time, with a row per time interval.

-btsecs <secs>

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins is 12 hours ago.

  If `-btsecs 0` is specified then the report begins at the earliest sample.

-etsecs <secs>

Selects the end time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.

-hires

Selects high resolution samples (5 minute intervals) for the report.

This is the default.

-hourly

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-summary <keyword>[,<keyword>]...

Summarize performance across requested objects and time range.

One of these summary keywords must be included:

- min—Display the minimum for each metric
- avg—Display the average for each metric
- max—Display the maximum for each metric
- <N>%—Display percentile for each metric. <N> may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

detail

Displays individual performance records in addition to one or more summaries. By default, -summary output excludes individual records and only displays the summary.

per_time

When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per time. By default, one summary is computed across all records.

per_group

When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per object grouping. By default, one summary is computed across all records.

only_compareby

When requesting data limited to certain object groupings with the -compareby option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all records and ignore the limitation of the -compareby option, though the "detail" output does conform to the -compareby object limitation.

-groupby <groupby>[,<groupby>...]

For -attime reports, generate a separate row for each combination of <groupby> items. Each <groupby> must be different and one of the following:

- NODE—The source controller node for the link
- QUEUE—The XCB queue
- NODE_TO—The destination controller node for the link
- ASIC_FROM—The source ASIC for the link
- ASIC_TO—The destination ASIC for the link

-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>

The -compareby option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following:

xfers_ps, kbps, szkb

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

<node>

Only the specified node numbers are included, where each node is a number from 0 through 7. This specifier can be repeated to display information for multiple nodes. If not specified, all nodes are included.

Authority

Any role in the system

Usage

None.

Examples

The following example displays aggregate hourly performance statistics for all nodes and links beginning 24 hours ago:

```
cli% srstatlink -hourly -btsecs -24h
```

The following example displays daily node link performance aggregated by node for nodes 0 and 1:

```
cli% srstatlink -daily -attime -groupby node 0 1
```

srstatpd

Syntax

```
srstatpd [options] [{<PDID>|<pattern>}...]
```

Description

The `srstatpd` command displays historical performance data reports for physical disks.

Options

-attime

Performance is shown at a particular time interval, specified by the `-etsecs` option, with one row per object group described by the `-groupby` option. Without this option performance is shown versus time, with a row per time interval.

-btsecs <secs>

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"

- Date string: "2012-10-26" or 2012-10-26
- Time string: "11:00:00" or 11:00:00

- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins is 12 hours ago.

  If `-btsecs 0` is specified then the report begins at the earliest sample.

-etsecs <secs>

Selects the end time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.

-hires

Selects high resolution samples (5 minute intervals) for the report.

This is the default.

-hourly

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-summary <keyword>[,<keyword>]...

Summarize performance across requested objects and time range.

One of these summary keywords must be included:

- `min`—Display the minimum for each metric
- `avg`—Display the average for each metric
- `max`—Display the maximum for each metric
- `<N>%`—Display percentile for each metric. `<N>` may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

detail

Displays individual performance records in addition to one or more summaries. By default, `-summary` output excludes individual records and only displays the summary.

per_time

When requesting data across multiple points in time (vstime) and multiple object groupings (`-groupby`) compute summaries per time. By default, one summary is computed across all records.

per_group

When requesting data across multiple points in time (vstime) and multiple object groupings (`-groupby`) compute summaries per object grouping. By default, one summary is computed across all records.

only_compareby

When requesting data limited to certain object groupings with the `-compareby` option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all records and ignore the limitation of the `-compareby` option, though the "detail" output does conform to the `-compareby` object limitation.

-groupby <groupby>[,<groupby>...]

For `-attime` reports, generate a separate row for each combination of `<groupby>` items. Each `<groupby>` must be different and one of the following:

- `PDID`—Physical disk ID
- `PORT_N`—The node number for the primary port for the PD
- `PORT_S`—The PCI slot number for the primary port for the PD
- `PORT_P`—The port number for the primary port for the PD
- `DISK_TYPE`—The disktype of the PD
- `SPEED`—The speed of the PD

-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>

The `-compareby` option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following:

```
read_iops, write_iops, total_iops, read_kbps, write_kbps, total_kbps,
read_svctms, write_svctms, total_svctms, read_ioszkb, write_ioszkb,
total_ioszkb, total_qlen, busy_pct
```

-disk_type <type>[,<type>...]

Limits the data to disks of the types specified. Allowed types are:

- FC—Fast Class
- NL—Nearline
- SSD—Solid State Drive
- SCM—Storage Class Memory Drive

-rpm <speed>[,<speed>...]

Limits the data to disks of the specified RPM. Allowed speeds are 7, 10, 15, 100, 150, and 1500.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

<PDID>|<pattern>

PDs with IDs that match either the specified PDID or glob-style pattern are included. This specifier can be repeated to include multiple PDIDs or patterns. If not specified, all PDs are included.

Authority

Any role in the system

Usage

For this command KB = 1000 bytes.

Examples

The following example displays aggregate hourly performance statistics for all physical disks beginning 24 hours ago:

```
cli% srstatpd -hourly -btsecs -24h
```

srstatport

Syntax

```
srstatport [options] [{<npat>:<spat>:<ppat>} ...]
```

Description

The srstatport command displays historical performance data reports for ports.

Options

-attime

Performance is shown at a particular time interval, specified by the -etsecs option, with one row per object group described by the -groupby option. Without this option performance is shown versus time, with a row per time interval.

-btsecs <secs>

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"

- Date string: "2012-10-26" or 2012-10-26
- Time string: "11:00:00" or 11:00:00

- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins is 12 hours ago.

  If `-btsecs 0` is specified then the report begins at the earliest sample.

-etsecs <secs>

Selects the end time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.

-hires

Selects high resolution samples (5 minute intervals) for the report.

This is the default.

-hourly

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-summary <keyword>[,<keyword>]...

Summarize performance across requested objects and time range.

One of these summary keywords must be included:

- min—Display the minimum for each metric
- avg—Display the average for each metric
- max—Display the maximum for each metric
- <N>%—Display percentile for each metric. <N> may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

detail

Displays individual performance records in addition to one or more summaries. By default, `-summary` output excludes individual records and only displays the summary.

per_time

When requesting data across multiple points in time (vstime) and multiple object groupings (`-groupby`) compute summaries per time. By default, one summary is computed across all records.

per_group

When requesting data across multiple points in time (vstime) and multiple object groupings (`-groupby`) compute summaries per object grouping. By default, one summary is computed across all records.

only_compareby

When requesting data limited to certain object groupings with the `-compareby` option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all records and ignore the limitation of the `-compareby` option, though the "detail" output does conform to the `-compareby` object limitation.

-groupby <groupby>[,<groupby>...]

For `-attime` reports, generate a separate row for each combination of `<groupby>` items. Each `<groupby>` must be different and one of the following:

- `PORT_N`—The node number for the port
- `PORT_S`—The PCI slot number for the port
- `PORT_P`—The port number for the port
- `PORT_TYPE`—The type of the port
- `GBITPS`—The speed of the port
- `TRANS_TYPE`—The transaction type - ctl or data

-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>

The `-compareby` option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following:

`read_iops, write_iops, total_iops, read_kbps, write_kbps, total_kbps, read_svctms, write_svctms, total_svctms, read_ioszkb, write_ioszkb, total_ioszkb, total_qlen, busy_pct`

-port_type <type>[,<type>...]

Limits the data to port of the types specified. Allowed types are:

- `disk`—Disk port
- `host`—Host Fibre channel port
- `iscsi`—Host ISCSI port
- `free`—Unused port
- `fs`—File Persona port
- `peer`—Data Migration FC port

- `rcip`—Remote copy IP port
- `rcfc`—Remote copy FC port

`-both | -ctl | -data`

Specifies that both control and data transfers are displayed (`-both`), only control transfers are displayed (`-ctl`), or only data transfers are displayed (`-data`). If this option is not specified, only data transfers are displayed.

`-sortcol <col>[,<dir>][:<col>[,<dir>]...]`

Sorts command output based on column number (`<col>`). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (`<dir>`) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

`<npat>:<spat>:<ppat>`

Ports with `<port_n>:<port_s>:<port_p>` that match any of the specified `<npat>:<spat>:<ppat>` patterns are included, where each of the patterns is a glob-style pattern. This specifier can be repeated to include multiple ports or patterns. If not specified, all ports are included.

Authority

Any role in the system

Usage

For this command KB = 1000 bytes.

Examples

The following example displays aggregate hourly performance statistics for disk and host ports on nodes 0 and 1 beginning 24 hours ago:

```
cli% srstatport -port_type disk,host -hourly -btsecs -24h 0:*:* 1:*:*
```

srstatqos

Syntax

```
srstatqos [options]
```

Description

The `srstatqos` command displays historical performance data reports for QoS rules.

Options

`-attime`

Performance is shown at a particular time interval, specified by the `-etsecs` option, with one row per object group described by the `-groupby` option. Without this option performance is shown versus time, with a row per time interval.

`-btsecs <secs>`

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins is 12 hours ago.

  If `-btsecs 0` is specified then the report begins at the earliest sample.

-etsecs <secs>

Selects the end time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.

-hires

Selects high resolution samples (5 minute intervals) for the report.

This is the default.

-hourly

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-summary <keyword>[,<keyword>]...

Summarize performance across requested objects and time range.

One of these summary keywords must be included:

- `min`—Display the minimum for each metric
- `avg`—Display the average for each metric

- `max`—Display the maximum for each metric
- `<N>%`—Display percentile for each metric. `<N>` may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

`detail`

Displays individual performance records in addition to one or more summaries. By default, `-summary` output excludes individual records and only displays the summary.

`per_time`

When requesting data across multiple points in time (vstime) and multiple object groupings (`-groupby`) compute summaries per time. By default, one summary is computed across all records.

`per_group`

When requesting data across multiple points in time (vstime) and multiple object groupings (`-groupby`) compute summaries per object grouping. By default, one summary is computed across all records.

`only_compareby`

When requesting data limited to certain object groupings with the `-compareby` option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all records and ignore the limitation of the `-compareby` option, though the "detail" output does conform to the `-compareby` object limitation.

`-vvset <VVSet_name|pattern>[,<VVSet_name|pattern>...]`

Limits the data to VVSets with names that match one or more of the specified names or glob-style patterns.

This option is deprecated and will be removed in a subsequent release.

`-all_others`

Displays statistics for all other I/O not regulated by a QoS rule.

This option is deprecated and will be removed in a subsequent release.

`-target {{vvset|domain}:{<name>|<pattern>}|sys:all_others}[,...]`

Limits the data to the specified QoS target rule(s).

Include a target type either `{vvset|domain}`, and a name or glob-style pattern.

The sys:all_others rule can be selected to display statistics for all other host I/O not regulated by any "on" QoS rule.

Multiple targets types can be specified as a comma separated list.

`-groupby <groupby>[,<groupby>...]`

For `-attime` reports, generate a separate row for each combination of `<groupby>` items. Each `<groupby>` must be different and one of the following:

- `DOM_NAME`—Domain name
- `TARGET_TYPE`—Type of QoS rule target, i.e. vvset
- `TARGET_NAME`—Name of QoS rule target
- `IOPS_LIMIT`—The I/O per second limit
- `BW_LIMIT_KBPS`—The KB per second bandwidth limit

-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>

The -compareby option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following:

```
read_iops, write_iops, total_iops, read_kbps, write_kbps, total_kbps,
read_svctms, write_svctms, total_svctms, read_ioszkb, write_ioszkb,
total_ioszkb, total_qlen, busy_pct read_wait_ms, write_wait_ms,
total_wait_ms, total_wqlen, total_io_rej, io_limit, bw_limit, priority,
io_guarantee, bw_guarantee, latency_target_ms, latency_ms
```

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

None.

Authority

Any role in the system

Usage

For this command KB = 1000 bytes.

Examples

The following example displays aggregate hourly performance statistics for all QoS rules beginning 24 hours ago:

```
cli% srstatqos -hourly -btsecs -24h
```

The following example displays the most recent daily (midnight) aggregate performance statistics for QoS domain rules starting with dom_2, OR for the special QoS rule which covers any host IO which is not regulated by any other rule:

```
cli% srstatqos -daily -attime -target domain:dom_2*,sys:all_others
```

srstatrcopy

Syntax

```
srstatrcopy [options]
```

Description

The srstatrcopy command displays historical performance data reports for Remote Copy links.

Options

-attime

Performance is shown at a particular time interval, specified by the -etsecs option, with one row per object group described by the -groupby option. Without this option performance is shown versus time, with a row per time interval.

-btsecs <secs>

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins is 12 hours ago.

  If `-btsecs 0` is specified then the report begins at the earliest sample.

-etsecs <secs>

Selects the end time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.

-hires

Selects high resolution samples (5 minute intervals) for the report.

This is the default.

-hourly

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-summary <keyword>[,<keyword>]...

Summarize performance across requested objects and time range.

One of these summary keywords must be included:

- `min`—Display the minimum for each metric
- `avg`—Display the average for each metric
- `max`—Display the maximum for each metric
- `<N>%`—Display percentile for each metric. `<N>` may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

`detail`

Displays individual performance records in addition to one or more summaries. By default, `-summary` output excludes individual records and only displays the summary.

`per_time`

When requesting data across multiple points in time (vstime) and multiple object groupings (`-groupby`) compute summaries per time. By default, one summary is computed across all records.

`per_group`

When requesting data across multiple points in time (vstime) and multiple object groupings (`-groupby`) compute summaries per object grouping. By default, one summary is computed across all records.

`only_compareby`

When requesting data limited to certain object groupings with the `-compareby` option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all records and ignore the limitation of the `-compareby` option, though the "detail" output does conform to the `-compareby` object limitation.

`-groupby <groupby>[,<groupby>...]`

For `-attime` reports, generate a separate row for each combination of `<groupby>` items. Each `<groupby>` must be different and one of the following:

- `TARGET_NAME`—The name of the Remote Copy target created with `creatercopytarget`
- `LINK_ID`—The ID of the Remote Copy target created with `creatercopytarget`
- `LINK_ADDR`—The address (IP or FC) of the Remote Copy target created with `creatercopytarget`
- `PORT_N`—The node number for the port used by a Remote Copy link
- `PORT_S`—The PCI slot number for the port used by a Remote Copy link
- `PORT_P`—The port number for the port used by a Remote Copy link

`-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>`

The `-compareby` option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following:

`kbs, kbps, hbrttms`

`-sortcol <col>[,<dir>][:<col>[,<dir>]...]`

Sorts command output based on column number (`<col>`). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (`<dir>`) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

-target <TARGET_NAME|pattern>[,<TARGET_NAME|pattern>...]

Limits the data to target names that match one or more of the specified TARGET_NAME or glob-style patterns.

-port <npat>:<spat>:<ppat>[,<npat>:<spat>:<ppat>...]

Ports with <port_n>:<port_s>:<port_p> that match any of the specified <npat>:<spat>:<ppat> patterns are included, where each of the patterns is a glob-style pattern. If not specified, all ports are included.

Specifiers

None.

Authority

Any role in the system

Usage

For this command KB = 1000 bytes.

Examples

The following example displays aggregate hourly performance statistics for all Remote Copy links beginning 24 hours ago:

```
cli% srstatrcopy -hourly -btsecs -24h
```

The following example displays most recent daily Remote Copy link performance for all targets grouped by target name:

```
cli% srstatrcopy -daily -attime -groupby target_name
```

The following example displays most recent daily Remote Copy link performance for all targets grouped by port_n, port_s, port_p:

```
cli% srstatrcopy -daily -attime -groupby port_n,port_s,port_p
```

srstatrcvv

Syntax

```
srstatrcvv [options]
```

Description

The srstatrcvv command displays historical performance data reports for Remote Copy volumes.

Options

-attime

Performance is shown at a particular time interval, specified by the -etsecs option, with one row per object group described by the -groupby option. Without this option performance is shown versus time, with a row per time interval.

-btsecs <secs>

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins is 12 hours ago.

  If `-btsecs 0` is specified then the report begins at the earliest sample.

-etsecs <secs>

Selects the end time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.

-hires

Selects high resolution samples (5 minute intervals) for the report.

This is the default.

-hourly

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-summary <keyword>[,<keyword>]...

Summarize performance across requested objects and time range.

One of these summary keywords must be included:

- `min`—Display the minimum for each metric
- `avg`—Display the average for each metric

- max—Display the maximum for each metric
- <N>%—Display percentile for each metric. <N> may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

detail

Displays individual performance records in addition to one or more summaries. By default, -summary output excludes individual records and only displays the summary.

per_time

When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per time. By default, one summary is computed across all records.

per_group

When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per object grouping. By default, one summary is computed across all records.

only_compareby

When requesting data limited to certain object groupings with the -compareby option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all records and ignore the limitation of the -compareby option, though the "detail" output does conform to the -compareby object limitation.

-groupby <groupby>[,<groupby>...]

For -attime reports, generate a separate row for each combination of <groupby> items. Each <groupby> must be different and one of the following:

- VV_NAME—The name of a volume admitted to a Remote Copy volume group with admitrcopyvv
- DOM_NAME—The domain name for a Remote Copy group when group was created with creatercopygroup
- TARGET_NAME—The target name of the Remote Copy target created with creatercopytarget
- TARGET_MODE—The target mode - Per: Periodic, Sync: Synchronous or Async: Asynchronous
- GROUP_NAME—The name of the Remote Copy group created with creatercopygroup
- GROUP_ROLE—The role (primary=1 or secondary=0) of the Remote Copy group
- PORT_TYPE—The port type (IP or FC) of the Remote Copy link(s) created with creatercopytarget
- PORT_N—The node number for the port used by a Remote Copy link
- PORT_S—The PCI slot number for the port used by a Remote Copy link
- PORT_P—The port number for the port used by a Remote Copy link
- VVSET_NAME—The virtual volume set name

-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>

The -compareby option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following:

```
lcl_read_iops, lcl_write_iops, lcl_total_iops, lcl_read_kbps,
lcl_write_kbps, lcl_total_kbps, lcl_read_svctms, lcl_write_svctms,
lcl_total_svctms, lcl_read_ioszkb, lcl_write_ioszkb, lcl_total_ioszkb,
lcl_busy_pct, lcl_total_qlen, rmt_read_iops, rmt_write_iops,
rmt_total_iops, rmt_read_kbps, rmt_write_kbps, rmt_total_kbps,
rmt_read_ioszkb, rmt_write_ioszkb, rmt_total_ioszkb, rmt_busy_pct,
rmt_total_qlen, rpo_timeInt
```

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

-vv <VV_name|VV_set|pattern>[,<VV_name|VV_set|pattern>...]

Limits the data to VVs with names that match one or more of the specified names or glob-style patterns. VV set name must be prefixed by "set:" and can also include patterns.

-target <TARGET_NAME|pattern>[,<TARGET_NAME|pattern>...]

Limits the data to target names that match one or more of the specified TARGET_NAME or glob-style patterns.

-mode <target_mode>[,<target_mode>]

Limits the data to TARGET_MODEs of the specified mode. Allowed modes are:

- Per—Periodic
- Sync—Synchronous
- Async—Asynchronous

-group <GROUP_NAME|pattern>[,<GROUP_NAME|pattern>...]

Limits the data to group names that match one or more of the specified GROUP_NAME or glob-style patterns.

Specifiers

None.

Authority

Any role in the system

Usage

For this command KB = 1000 bytes.

Examples

The following example displays aggregate hourly performance statistics for all Remote Copy links beginning 24 hours ago:

```
cli% srstatrcvv -hourly -btsecs -24h
```

The following example displays most recent daily Remote Copy volumes performance for volumes grouped by volume name:

```
cli% srstatrcvv -daily -attime -groupby vv_name
```

The following example displays most recent daily Remote Copy volumes performance for volumes grouped by target name:

```
cli% srstatrcvv -daily -attime -groupby target_name
```

The following example displays most recent daily Remote Copy volumes performance for volumes grouped by group name:

```
cli% srstatrcvv -daily -attime -groupby group_name
```

The following example displays most recent daily Remote Copy volumes performance for volumes grouped by port_n, port_s, port_p:

```
cli% srstatrcvv -daily -attime -groupby port_n,port_s,port_p
```

srstatvlun

Syntax

```
srstatvlun [options]
```

Description

The `srstatvlun` command displays historical performance data reports for VLUNs.

Options

`-attime`

Performance is shown at a particular time interval, specified by the `-etsecs` option, with one row per object group described by the `-groupby` option. Without this option performance is shown versus time, with a row per time interval.

`-btsecs <secs>`

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins is 12 hours ago.

  If `-btsecs 0` is specified then the report begins at the earliest sample.

`-etsecs <secs>`

Selects the end time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.

-hires

Selects high resolution samples (5 minute intervals) for the report.

This is the default.

-hourly

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-summary <keyword>[,<keyword>]...

Summarize performance across requested objects and time range.

One of these summary keywords must be included:

- `min`—Display the minimum for each metric
- `avg`—Display the average for each metric
- `max`—Display the maximum for each metric
- `<N>%`—Display percentile for each metric. `<N>` may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

detail

Displays individual performance records in addition to one or more summaries. By default, `-summary` output excludes individual records and only displays the summary.

per_time

When requesting data across multiple points in time (vstime) and multiple object groupings (`-groupby`) compute summaries per time. By default, one summary is computed across all records.

per_group

When requesting data across multiple points in time (vstime) and multiple object groupings (`-groupby`) compute summaries per object grouping. By default, one summary is computed across all records.

only_compareby

When requesting data limited to certain object groupings with the -compareby option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all records and ignore the limitation of the -compareby option, though the "detail" output does conform to the -compareby object limitation.

-groupby <groupby>[,<groupby>...]

For -attime reports, generate a separate row for each combination of <groupby> items. Each <groupby> must be different and one of the following:

- DOM_NAME—Domain name
- VV_NAME—Virtual Volume name
- HOST_NAME—Host name
- LUN—The LUN number for the VLUN
- HOST_WWN—The host WWN for the VLUN
- PORT_N—The node number for the VLUN port
- PORT_S—The PCI slot number for the VLUN port
- PORT_P—The port number for the VLUN port
- VVSET_NAME—Virtual volume set name
- HOSTSET_NAME—Host set name
- VM_NAME—Virtual Machine Name for vVol based VMs
- VM_ID—Virtual Machine Identification number for vVol based VMs
- VM_HOST—Virtual Machine host for vVol based VMs
- VVOLSC—Virtual Volume Storage Container for vVol based VMs

-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>

The -compareby option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following:

read_iops, write_iops, total_iops, read_kbps, write_kbps, total_kbps, read_svctms, write_svctms, total_svctms, read_ioszkb, write_ioszkb, total_ioszkb, total_qlen, busy_pct

-host <host_name|host_set|host_wwn|pattern>[,<host_name|host_set|host_wwn|pattern>...]

Limits the data to hosts with names that match one or more of the specified names or glob-style patterns. Host set name must start with "set:" and can also include patterns. To specify the host by WWN, start with "wwn:". A WWN can also include glob-style patterns.

-vv <VV_name|VV_set|pattern>[,<VV_name|VV_set|pattern>...]

Limits the data to VVs with names that match one or more of the specified names or glob-style patterns. VV set name must be prefixed by "set:" and can also include patterns.

-l <LUN|pattern>[,<LUN|pattern>...]

Limits the data to LUNs that match one or more of the specified LUNs or glob-style patterns.

-port <npat>:<spat>:<ppat>[,<npat>:<spat>:<ppat>...]

Ports with <port_n>:<port_s>:<port_p> that match any of the specified <npat>:<spat>:<ppat> patterns are included, where each of the patterns is a glob-style pattern. If not specified, all ports are included.

-vlun <host>:<vv>[:<lun>:<port_n>:<port_s>:<port_p>]
[,<host>:<vv>[:<lun>:<port_n>:<port_s>:<port_p>]]...

Limits the data to VLUNs matching the specified combination of host, VV, lun, and port. Each of these components in this option may be a glob-style pattern. The host and VV components may specify a corresponding object set by prefixing "set:" to the component. The host component may specify a WWN by prefixing the component with "wwn:". The lun and port components are optional, and if not present, data will be filtered to any matching combination of host and VV. This option cannot be combined with -host, -vv, -l, or -port.

-vmname {<VM_name>|<pattern>}[,{<VM_name>|<pattern>}]

Limits the data to VMs that match one or more of the specified VM names or glob-styled patterns for vVol based VMs.

-vmid {<VM_ID>|<pattern>}[,{<VM_ID>|<pattern>}]

Limits the data to VMs that match one or more of the specified VM IDs or glob-styled patterns for vVol based VMs.

-vmhost {<VM_host_name>|<pattern>}[,{<VM_host_name>|<pattern>}]

Limits the data to VMs that match one or more of the specified VM host names or glob-styled patterns for vVol based VMs.

-vvolsc {<VVol_container_name>|<pattern>}[,{<VVol_container_name>|
<pattern>}...]

Limits the data to vVol containers that match one or more of the specified vVol container names or glob-styled patterns.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

None.

Authority

Any role in the system

Usage

- For this command KB = 1000 bytes.
- Because VVs can belong to multiple VV sets, requesting data for more than one VV set can include duplicate data for any VVs with multiple membership.

Examples

The following example displays aggregate hourly performance statistics for all VLUNs beginning 24 hours ago:

```
cli% srstatvlun -hourly -btsecs -24h
```

VV or host sets can be specified with patterns:

```
cli% srstatvlun -btsecs -2h -host set:hostset -vv set:vvset*
```

srstatvv

Syntax

```
srstatvv [options] [{<VV_name>|<VV_set>|<pattern>}...]
```

Description

Displays historical performance data reports for virtual volumes (VVs).

Options

-attime

Performance is shown at a particular time interval, specified by the -etsecs option, with one row per object group described by the -groupby option. Without this option performance is shown versus time, with a row per time interval.

-btsecs <secs>

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h, or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h), or days (e.g. -7d).

  If it is not specified then the time at which the report begins depends on the sample category (-hires, -hourly, -daily):
  - For hires, the default begin time is 12 hours ago (-btsecs -12h).
  - For hourly, the default begin time is 7 days ago (-btsecs -7d).
  - For daily, the default begin time is 90 days ago (-btsecs -90d).

  If begin time and sample category are not specified then the time the report begins is 12 hours ago and the default sample category is hires. If -btsecs 0 is specified then the report begins at the earliest sample.

-etsecs <secs>

Selects the end time in seconds for the report. If -attime is specified, select the time for the report.

The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:

- Full time string including time zone: "2012-10-26 11:00:00 PDT"
- Full time string excluding time zone: "2012-10-26 11:00:00"
- Date string: "2012-10-26" or 2012-10-26
- Time string: "11:00:00" or 11:00:00

- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified, then the report ends with the most recent sample.

-hires

Selects high resolution samples (5 minute intervals) for the report. This is the default.

-hourly

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-summary <keyword>[,<keyword>]...

Summarize performance across requested objects and time range. One of these 4 summary keywords must be included:

- `min`—Display the minimum for each metric.
- `avg`—Display the average for each metric.
- `max`—Display the maximum for each metric.
- `<N>%`—Display percentile for each metric. <N> may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

detail

Displays individual performance records in addition to one or more summaries. By default, `-summary` output excludes individual records and only displays the summary.

per_time

When requesting data across multiple points in time (vstime) and multiple object groupings (`-groupby`) compute summaries per time. By default, one summary is computed across all records.

per_group

When requesting data across multiple points in time (vstime) and multiple object groupings (`-groupby`) compute summaries per object grouping. By default, one summary is computed across all records.

only_compareby

When requesting data limited to certain object groupings with the `-compareby` option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all records and ignore the limitation of the `-compareby` option, though the "detail" output does conform to the `-compareby` object limitation.

-groupby <groupby>[,<groupby>...]

For -attime reports, generate a separate row for each combination of <groupby> items. Each <groupby> must be different and one of the following:

- DOM_NAME—Domain name
- VVID—Virtual volume ID
- VV_NAME—Virtual volume name
- WWN—Virtual volume world wide name (WWN)
- SNP_CPG_NAME—Snap space Common Provisioning Group name
- USR_CPG_NAME—User space Common Provisioning Group name
- PROV_TYPE—The virtual volume provisioning type
- VV_TYPE—The type of the virtual volume
- VVSET_NAME—Virtual volume set name

-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>

The -compareby option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following:

read_iops, write_iops, total_iops, read_kbps, write_kbps, total_kbps, read_svctms, write_svctms, total_svctms, read_ioszkb, write_ioszkb, total_ioszkb, total_qlen, busy_pct

-usr_cpg <CPG_name|pattern>[,<CPG_name|pattern>...]

Only include VVs whose usr space is mapped to a CPG whose name matches one of the specified CPG_name or glob-style patterns.

-snp_cpg <CPG_name|pattern>[,<CPG_name|pattern>...]

Only include VVs whose snp space is mapped to a CPG whose name matches one of the specified CPG_name or glob-style patterns.

-prov <prov_type>[,<prov_type>...]

Only include VVs of the specified provisioning type(s). The possible values are: cpvv dds full peer snp tdvv tpsd tpvv

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

<VV_name>|<VV_set>|<pattern>

VVs with names that match either the specified VV_name or glob-style pattern are included. Data for all VV members of a specified VV set may also be requested. VV set names must be prefixed by "set:" and can also include patterns. This specifier can be repeated to include multiple VV names, sets, or patterns. If not specified, all VVs are included.

Authority

Any role in the system

Usage

- For this command KB = 1000 bytes.
- Because VVs can belong to multiple VV sets, requesting data for more than one VV set can include duplicate data for any VVs with multiple membership.

Examples

The following example displays aggregate hourly performance statistics for all VVs beginning 24 hours ago:

```
cli% srstatvv -hourly -btsecs -24h
```

VV or host sets can be specified with patterns:

```
cli% srstatvv -btsecs -2h set:vvset_ESX_*
```

srsysspace

Syntax

```
srsysspace [options]
```

Description

The `srsysspace` command displays historical space data reports for the system.

Options

-attime

Performance is shown at a particular time interval, specified by the `-etsecs` option, with one row per object group described by the `-groupby` option. Without this option performance is shown versus time, with a row per time interval.

-btsecs <secs>

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

If it is not specified then the time at which the report begins depends on the sample category (`-hires`, `-hourly`, `-daily`):

- For hires, the default begin time is 12 hours ago (`-btsecs -12h`).
- For hourly, the default begin time is 7 days ago (`-btsecs -7d`).
- For daily, the default begin time is 90 days ago (`-btsecs -90d`).

If begin time and sample category are not specified then the time the report begins is 12 hours ago and the default sample category is hires. If `-btsecs 0` is specified then the report begins at the earliest sample.

`-etsecs <secs>`

Selects the end time in seconds for the report. If `-attime` is specified, select the time for the report.

The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.
- Specifying an etsecs value in the future, prefixed by "+", will provide a forecast of the space statistics. Either absolute or relative time may be specified, such as "+2017-03-24" or "+2d".

`-hires`

Selects high resolution samples (5 minute intervals) for the report. This is the default.

`-hourly`

Selects hourly samples for the report.

`-daily`

Selects daily samples for the report.

`-summary <keyword>[,<keyword>]...`

Summarize performance across requested objects and time range.

One of these summary keywords must be included:

- `min`—Display the minimum for each metric
- `avg`—Display the average for each metric
- `max`—Display the maximum for each metric
- `<N>%`—Display percentile for each metric. `<N>` may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

detail

Displays individual performance records in addition to one or more summaries. By default, `-summary` output excludes individual records and only displays the summary.

per_time

When requesting data across multiple points in time (vstime) and multiple object groupings (`-groupby`) compute summaries per time. By default, one summary is computed across all records.

per_group

When requesting data across multiple points in time (vstime) and multiple object groupings (`-groupby`) compute summaries per object grouping. By default, one summary is computed across all records.

only_compareby

When requesting data limited to certain object groupings with the `-compareby` option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all records and ignore the limitation of the `-compareby` option, though the "detail" output does conform to the `-compareby` object limitation.

-groupby <groupby>[,<groupby>...]

For `-attime` reports, generate a separate row for each combination of `<groupby>` items. Each `<groupby>` must be different and one of the following:

- `DOM_NAME`—Domain name
- `DISK_TYPE`—The disktype of the PDs

-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>

The `-compareby` option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following:

```
Total_MB, Alloc_MB, Used_MB, System_MB, Free_MB, Failed_MB, Compact_Ratio,
Dedup_Ratio, OverProv_Ratio
```

```
-disk_type <type>[,<type>...]
```

Limits the data to disks of the types specified. Allowed types are:

- FC—Fast Class
- NL—Nearline
- SSD—Solid State Drive
- SCM—Storage Class Memory Drive

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (`<col>`). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (`<dir>`) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

None.

Authority

Any role in the system

Usage

For this command, 1MB = 1048576 bytes.

Examples

The following example displays aggregate hourly space information for the system, organized by PD disk type, beginning 24 hours ago:

```
cli% srsysspace -hourly -btsecs -24h -groupby disk_type
```

srvvspace

Syntax

```
srvvspace [options] [{<VV_name>|<VV_set>|<pattern>}...]
```

Description

The `srvvspace` command displays historical space data reports for virtual volumes (VVs).

Options

`-attime`

Performance is shown at a particular time interval, specified by the `-etsecs` option, with one row per object group described by the `-groupby` option. Without this option performance is shown versus time, with a row per time interval.

`-btsecs <secs>`

Selects the begin time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, `<secs>` can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the time at which the report begins is 12 hours ago.

  If `-btsecs 0` is specified then the report begins at the earliest sample.

`-etsecs <secs>`

Selects the end time in seconds for the report. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:

- Full time string including time zone: "2012-10-26 11:00:00 PDT"
- Full time string excluding time zone: "2012-10-26 11:00:00"
- Date string: "2012-10-26" or 2012-10-26
- Time string: "11:00:00" or 11:00:00

- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the report ends with the most recent sample.
- Specifying an etsecs value in the future, prefixed by "+", will provide a forecast of the space statistics. Either absolute or relative time may be specified, such as "+2017-03-24" or "+2d".

-hires

Selects high resolution samples (5 minute intervals) for the report.

This is the default.

-hourly

Selects hourly samples for the report.

-daily

Selects daily samples for the report.

-summary <keyword>[,<keyword>]...

Summarize performance across requested objects and time range.

One of these summary keywords must be included:

- min—Display the minimum for each metric
- avg—Display the average for each metric
- max—Display the maximum for each metric
- <N>%—Display percentile for each metric. <N> may be any number from 0 to 100. Multiple percentiles may be specified.

Other keywords which modify the summary display or computation:

detail

Displays individual performance records in addition to one or more summaries. By default, -summary output excludes individual records and only displays the summary.

per_time

When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per time. By default, one summary is computed across all records.

per_group

When requesting data across multiple points in time (vstime) and multiple object groupings (-groupby) compute summaries per object grouping. By default, one summary is computed across all records.

only_compareby

When requesting data limited to certain object groupings with the -compareby option, use this keyword to compute summaries using only that reduced set of object groupings. By default, summaries are computed from all

records and ignore the limitation of the –compareby option, though the "detail" output does conform to the –compareby object limitation.

-groupby <groupby>[,<groupby>...]

For –attime reports, generate a separate row for each combination of <groupby> items. Each <groupby> must be different and one of the following:

- DOM_NAME—Domain name
- VVID—Virtual volume ID
- VV_NAME—Virtual volume name
- BSID—Virtual volume ID of the base virtual volume
- WWN—Virtual volume world wide name (WWN)
- SNP_CPG_NAME—Snap space Common Provisioning Group name
- USR_CPG_NAME—User space Common Provisioning Group name
- PROV_TYPE—The virtual volume provisioning type
- VV_TYPE—The type of the virtual volume
- VVSET_NAME—Virtual volume set name
- VM_NAME—Virtual Machine name for vVol based VMs
- VM_ID—Virtual Machine Identification number for vVol based VMs
- VM_HOST—Virtual Machine host for vVol based VMs
- VVOLSC—Virtual Volume Storage Container for vVol based VMs
- VVOL_STATE—Virtual Volume state, either bound or unbound
- COMPR—Whether Compression is enabled, disabled, or NA

-compareby <top|bottom>,<number_of_objects>,<field_used_for_comparison>

The –compareby option limits output records to only certain objects, compared by a specified field. Either the top or bottom X objects can be displayed, up to 32 objects for vstime reports or 128 objects for attime reports. The field used for comparison can be any of the groupby fields or one of the following:

userrawrsvd_mb, snaprawrsvd_mb, totalrawrsvd_mb, userused_mb, userfree_mb, userrsvd_mb, snapused_mb, snapfree_mb, snaprsvd_mb, snapvcopy_mb, totalvcopy_mb, totalused_mb, totalrsvd_mb, virtualsize_mb, ce_virtual_blks, ce_used_blks, ce_tdvv_wrt_blks, ce_dd_used_blks, compact_ratio, hostwrt_mb, compr_ratio, compr_gc_kbps

-usr_cpg <CPG_name|pattern>[,<CPG_name|pattern>...]

Only include VVs whose usr space is mapped to a CPG whose name matches one of the specified CPG_name or glob-style patterns.

-snp_cpg <CPG_name|pattern>[,<CPG_name|pattern>...]

Only include VVs whose snp space is mapped to a CPG whose name matches one of the specified CPG_name or glob-style patterns.

-prov <prov_type>[,<prov_type>...]

Only include VVs of the specified provisioning type(s). The possible values are: dds peer snp tdvv tpvv.

cpvv dds full peer snp tdvv tpsd tpvv

-compr <Yes|No|Off|NA>

Only include VVs of the specified compression setting.

-vmname {<VM_name>|<pattern>}[,{<VM_name>|<pattern>}]

Limits the data to VMs that match one or more of the specified VM names or glob-styled patterns for vVol based VMs.

-vmid {<VM_ID>|<pattern>}[,{<VM_ID>|<pattern>}]

Limits the data to VMs that match one or more of the specified VM IDs or glob-styled patterns for vVol based VMs.

-vmhost {<VM_host_name>|<pattern>}[,{<VM_host_name>|<pattern>}]

Limits the data to VMs that match one or more of the specified VM host names or glob-styled patterns for vVol based VMs.

-vvolstate {bound|unbound}

Limits the data to vVOLs that have states in either the Bound or Unbound state.

-vvolsc {<vVol_container_name>|<pattern>}[,{<vVol_container_name>|<pattern>}...]

Limits the data to vVol containers that match one or more of the specified vVol container names or glob-styled patterns.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (`<col>`). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (`<dir>`) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

Specifiers

<VV_name>|<VV_set>|<pattern>

VVs with names that match either the specified `VV_name` or glob-style pattern are included. Data for all VV members of a specified VV set may also be requested. VV set name must be prefixed by "set:" and can also include patterns. This specifier can be repeated to include multiple VV names, sets, or patterns. If not specified, all VVs are included.

Authority

Any role in the system

Usage

- For this command, 1MB = 1048576 bytes.
- Because VVs can belong to multiple VV sets, requesting data for more than one VV set can include duplicate data for any VVs with multiple membership.

Examples

The following example displays aggregate hourly VV space information for VVs with names matching either "dbvv*" or "testvv*" patterns beginning 24 hours ago:

```
cli% srvvspace -hourly -btsecs -24h dbvv* testvv*
```

The following example displays VV space information for the most recent daily sample aggregated by the VV name for VVs with names that match the pattern "tp*":

```
cli% srvvspace -daily -attime -groupby vv_name tp*
```

The following example displays all VVs with compression enabled:

```
cli% srvvspace -compr Yes -attime
```

Start Commands

startao

Syntax

```
startao [option] <aocfg_name>
```

Description

The `startao` command starts execution of an Adaptive Optimization (AO) configuration using data region level performance data collected for the specified number of hours.

Options

`-btsecs <secs>`

Selects the begin time in seconds for the analysis period. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00

If it is not specified then the analysis begins 12 hours ago.

`-etsecs <secs>`

Selects the end time in seconds for the analysis period. The value can be specified as either:

- The absolute epoch time (for example 1351263600).
- The absolute time as a text string in one of the following formats:
  - Full time string including time zone: "2012-10-26 11:00:00 PDT"
  - Full time string excluding time zone: "2012-10-26 11:00:00"
  - Date string: "2012-10-26" or 2012-10-26
  - Time string: "11:00:00" or 11:00:00
- A negative number indicating the number of seconds before the current time. Instead of a number representing seconds, <secs> can be specified with a suffix of m, h or d to represent time in minutes (e.g. -30m), hours (e.g. -1.5h) or days (e.g. -7d).

  If it is not specified then the analysis ends with the most recent sample.

`-compact <mode>`

Specifies if and how CPGs should be compacted. Choices for <mode> are:

- auto—Automatically select the compactcpg mode (default). This will free up the most space but can potentially take longer because it may cause additional region moves to increase consolidation. This is the default mode. In

this mode, trimonly is run first to trim LDs without performing region moves, and then the full compactcpg is run if space is still needed.

- trimonly—Only run compactcpg with the -trimonly option. This will not perform any region moves during compactcpg.
- no—Do not run compactcpg. This option may be used if compactcpg is run or scheduled separately.

-dryrun

Do not execute the region moves, only show which moves would be done.

-maxrunh <hours>

Selects the approximate maximum run time in hours (default is 6 hours).

The number should be between 1 and 24 hours.

The command will attempt to limit the amount of data to be moved so the command can complete by the specified number of hours. If the time runs beyond the specified hours, the command will abort at an appropriate time.

-min_iops <min_iops>

Do not execute the region moves if the average IOPS during the measurement interval is less than <min_iops>. If the -vv option is not specified, the IOPS are for all the LDs in the AOCFG.

If the -vv option is specified, the IOPS are for all the VLUNs that include matching VVs.

If min_iops is not specified, the default value is 50.

-mode <mode>

Overrides the optimization bias of the AO config, for instance to control AO differently for different VVs. Can be one of the following:

- Performance—Move more regions towards higher performance tier.
- Balanced—Balanced between higher performance and lower cost.
- Cost—Move more regions towards lower cost tier.

-vv <VV_name|VV_set|pattern>[,<VV_name|VV_set|pattern>...]

Limits the analysis and data movement to VVs with names that match one or more of the specified names or glob-style patterns. VV set names must be prefixed by "set:". Note that snapshot VVs will not be considered since only base VVs have region space. Each VV's user CPG must be part of the specified AOCFG in order to be optimized. Snapshots in a VV's tree will not be optimized.

`-t0min <size>[g|G|t|T]`

`-t1min <size>[g|G|t|T]`

`-t2min <size>[g|G|t|T]`

Specifies the minimum space utilization of the tier CPG for AO to maintain when optimizing regions between tiers. The size can be specified in MB (default) or GB (using g or G) or TB (using t or T).

Setting a minimum to 0 (default) indicates that no minimum space utilization will be enforced.

`-t0max <size>[g|G|t|T]`

`-t1max <size>[g|G|t|T]`

`-t2max <size>[g|G|t|T]`

Specifies the maximum space utilization of the tier CPG. AO will move regions into and out of the CPG based on their relative access rate history, but will not exceed this maximum size in the CPG.

The size can be specified in MB (default) or GB (using g or G) or TB (using t or T). Setting a max to 0 (default) indicates that AO will use other indicators to decide the maximum CPG space utilization: either the CPG sdgl, sdgw, or maximum possible growth size.

Specifiers

<aocfg_name>

The AO configuration name, using up to 31 characters.

Authority

- Super, Edit
- Any role granted the `aocfg_start` right

Usage

- The `startao` command can be run manually, but it can also be scheduled to run automatically at specific times using the **createsched** command.
- The behavior of -compact option is to run only when there is a little space available to grow. To judge if space is needed, one or more of the following conditions must be true (otherwise, **compactcpg** is NOT run):
  - There is unused space in the CPG and the current allocated space is above the warning limit.
  - The unused space in the CPG is more than a certain fraction (25%) of the CPG space. This is the total unused space across user, snapshot, and admin space.
  - The space available to grow the CPG (i.e. free chunklets for the CPG) is less than 4 times the CPG growth increment. This can be examined by comparing the LDFree output of `showspace -cpg` with `showcpg -sdg`.
- If desired, **compactcpg** can be scheduled separately from `startao` to always run `compactcpg`.
- The -min_iops option can be used to prevent movement of normally busy regions to slower tiers if the application associated with an AOCFG (or the matching VVs if -vv is specified) is down or inactive (i.e. total LD IOP/s for the AOCFG (or the VLUN IOP/s for the matching VVs) is less than <min_iops>) during the measurement interval. This ensures that when the application resumes normal operation, its busy regions are still in the faster tiers.

Examples

Starting execution of AO config prodaocfg using data for the past 3 hours:

```
cli% startao -btsecs -3h prodaocfg
```

Starting execution of AO config prodaocfg using data from 12 hours ago until 3 hours ago, allowing up to 6 hours to complete:

```
cli% startao -btsecs -12h -etsecs -3h -maxrunh 6 prodaocfg
```

Starting execution of AO for the vvset dbvvset in AOCFG prodaocfg using data for the past 3 hours:

```
cli% startao -btsecs -3h -vv set:dbvvset prodaocfg
```

startcim

Syntax

`startcim`

Description

The `startcim` command starts the CIM server to service CIM requests. By default, the CIM server is not started until this command is issued.

Options

None.

Specifiers

None.

Authority

- Super, Service
- Any role granted the `cim_start` right

Usage

- Access to all domains is required to run this command.
- By default, the CIM server is not started until this command is issued.
- Use **`stopcim`** to stop the CIM server.

Examples

The following example starts the CIM server:

```
cli% startcim
CIM server will start shortly.
```

startfs

Syntax

```
startfs [options] <n:s:p> <n:s:p> [{<n:s:p> <n:s:p>} ...]
startfs -enable [<nodeid>,...]
```

Description

The `startfs` command initializes and starts File Persona on the system.

Options

`-enable`

Previously stopped File Persona will be started on the specified nodes. If no nodes are specified File Persona will be started on all nodes.

`-nowait`

Will not wait for the tasks to complete.

This option must not be used with -enable option.

`-cpgname <name>`

Specifies the name of the CPG to be used. If no CPG of that name is present, a CPG is created with "-p -devtype FC". If that fails, then it is retried with NL then SSD.

Specifiers

<n:s:p>

Specifies the node, slot and port on which File Persona will be configured. Nodes must be specified in matched pairs e.g. 0 and 1, 2 and 3 etc. Only one valid port on the slot needs to be specified as the whole NIC will be configured.

<nodeid>

Specifies the node on which File Persona is to be started. This is valid only with the -enable option.

Authority

- Super, Service
- Any role granted the `fs_start` right.

Usage

- Access to all domains is required to run this command.
- The `startfs` command without the -enable option will install and start File Persona and configure networking and must be run before any File Persona commands can be used.
- If -enable is provided, File Persona is started on all or specified nodes. File Persona must have been previously installed on those nodes.

Examples

The following example starts File Persona on a node pair 0 and 1:

```
cli% startfs 0:2:1 1:2:1
```

The following example enables File Persona on nodes 0 and 1:

```
cli% startfs -enable 0,1
```

startfsarchive

Syntax

```
startfsarchive scan [-resume <jobid>] [-fpg <fpgname>] -fstore <fstorename>
[-path <path>] <vfs>
```

Description

The `startfsarchive` command is used to start and resume the data validation scan on the retained/WORM'ed file.

Options

-fpg <fpgname>

Specifies the FPG that contains the VFS. This is required if VFS exists under multiple FPGs.

-fstore <fstorename>

Name of the file store under VFS.

-path <path>

This is full path starting from FPG mount point, VFS, File Store, share and subfolders.

-resume <jobid>

Job Id of the paused retention validation scan process.

Specifiers

<vfs>

Specifies the name of the VFS.

Subcommands

scan

Runs the retention validation scan at the FSTORE level on the given path for specific VFS. If path is not provided, scan will be initiated at the root of the FSTORE. Irrespective of VFS, a user can queue a maximum of 50 validation jobs per file serving node, other than one in the running state.

This subcommand can also resume a paused job if -resume is used with the job id.

Authority

- Super, Service
- Any role granted the `fsarchive_start` right.

Usage

Access to all domains is required to run this command.

Examples

The following example starts data validation scan on retained VFS:

```
cli% startfsarchive scan -fpg fpg1 -fstore fstore1
-path /fpg1/vfs1/fstore1/ vfs1
```

The following example resumes paused data validation scan process:

```
cli% startfsarchive scan -resume 10 -fpg fpg1 -fstore fstore1 vfs1
```

startfsav

Syntax

The syntax for the `startfsav` command can be one of the following:

```
startfsav svc
startfsav update
startfsav scan [-resume <scan_id>] [-fpg <fpgname>] [-fstore <fstore>] <vfs>
startfsav scan [-fpg <fpgname>] [-fstore <fstore> [-path <path_name>]] <vfs>
```

Description

The `startfsav` command is used to both start the antivirus service as well as to initiate a scan on a Virtual File Server (VFS) or File Store.

Options

-fpg <fpgname>

The name of the File Provisioning Group in which the VFS was created.

-fstore <fstore>

Specifies the File Store name, using up to 31 characters.

-path <path_name>

The path, on the containing File Store, which the scan must be initiated.

-resume <scan_id>

When the specified scan_id corresponds to an ongoing scan in the paused state, will cause the scan to resume.

Specifiers

<vfs>

The Virtual File Server (VFS) name, on which scan to Start or Resume.

Subcommands

svc

Specifies antivirus service.

update

Specifies to update virus scan engine definitions.

scan

Starts and resumes antivirus scan.

Authority

- Super, Edit
- Any role granted the `fsav_start` right

Usage

- Access to all domains is required to run this command.
- The "`stopfsav scan <vfs> <scan_id>`" command can be used to stop the running scan.
- The "`stopfsav scan -pause <vfs> <scan_id>`" command can be used to pause the running VFS scan. See **stopfsav**.

Examples

Starting antivirus service:

```
cli% startfsav svc
cli% startfsav update
```

Initiating an antivirus scan on path "/foo/bar" in the engineering File Store, unityVFS VFS:

```
cli% startfsav scan -fstore engineering -path /foo/bar unityVFS
Initiated scan id: AvScanTask_45
```

Resuming a previous paused antivirus scan with VFS "unityVFS" and scan_id "AvScanTask_20":

```
cli% startfsav scan -resume AvScanTask_20 unityVFS
```

startfsnapclean

Syntax

The syntax for the `startfsnapclean` command is as follows:

```
startfsnapclean [-resume] <fpgname>
```

```
startfsnapclean [-reclaimStrategy {<maxspeed>|<maxspace>}] <fpgname>
```

```
startfsnapclean <fpgname>
```

Description

The `startfsnapclean` command starts an on-demand reclamation task on a file provisioning group to reclaim space from snapshots which have been marked as deleted.

Options

`-resume`

Specifies a paused reclamation task must be resumed.

`-reclaimStrategy <maxspeed>|<maxspace>`

Specifies the strategy to be used while reclaiming snap space.

- <maxspeed>—Suggests optimize for speedy reclamation.
- <maxspace>—Suggests optimize to reclaim maximum space.

Specifiers

`<fpgname>`

Specifies the name of the file provisioning group.

Authority

- Super, Service
- Any role granted the `fsnapclean_start` right

Usage

Access to all domains is required to run this command.

Examples

The following example starts a reclamation task on samplepool file provisioning group with a maximum speed reclamation strategy:

```
cli% startfsnapclean -reclaimStrategy maxspeed samplepool
```

The following example starts a reclamation task on samplepool file provisioning group with a maximum space reclamation strategy:

```
cli% startfsnapclean -reclaimStrategy maxspace samplepool
```

The following example resumes a paused reclamation task on samplepool file provisioning group:

```
cli% startfsnapclean -resume samplepool
```

startfsndmp

Syntax

`startfsndmp`

Description

The `startfsndmp` command is used to start both NDMP service and ISCSI service.

Options

None.

Specifiers

None.

Authority

- Super, Service.
- Any role granted the `fsndmp_start` right

Usage

- Access to all domains is required to run this command.
- NDMP and ISCSI services can be stopped by using the **stopfsndmp** command.
- The status of NDMP and ISCSI service can be checked by using the "`showfsndmp -conf`" command. See **showfsndmp**.

Examples

Start NDMP and ISCSI services:

```
cli% startfsndmp
```

startld

Syntax

```
startld [option] <LD_name>
```

Description

The `startld` command starts data services on a LD that has not yet been started.

Options

-ovrd

Specifies that the LD is forced to start, even if some underlying data is missing.

Specifiers

<LD_name>

Specifies the LD name, using up to 31 characters.

Authority

- Super, Service
- Any role granted the `ld_start` right

Usage

Access to all domains is required to run this command.

Examples

The following example starts data services on LD ld5:

```
cli% startld ld5
```

startnoderescue

Syntax

```
startnoderescue -node <node>
```

Description

Initiates a node rescue that initializes the internal node disk of the specified node to match the contents of the other node disks. Reports progress as a task. Completes the copy over an Ethernet path internal to the system and does not require an external Ethernet connection to the network.

On legacy systems (7000 series, 8000 series, and 10000 series), the copy is done over the network, so the node to be rescued must have an ethernet connection. It will automatically select a valid unused link local address.

Specifiers

<node>

Specifies the node to rescue. This node must be physically present in the system and powered on, but not part of the cluster.

Authority

Super, Service

Usage

On systems other than T and F class, node rescue will automatically be started when a blank node disk is inserted into a node. The startnoderescue command only needs to be manually issued if the node rescue must be redone on a disk that is not blank. For T and F class systems, startnoderescue must always be issued to perform a node rescue.

Examples

The following example shows starting a node rescue of node 2:

```
cli% startnoderescue -node 2
Node rescue from node 0 to node 2 started.

cli% showtask
Id Type        Name          Status Phase Step -------StartTime------- -FinishTime- -Priority- ---User----
96 node_rescue node_2_rescue active  1/1   0/1  2018-06-15 18:19:38 PDT -            n/a        sys:3parsys
```

startrcopy

Syntax

```
startrcopy
```

Description

The `startrcopy` command starts the Remote Copy Service.

Options

None.

Specifiers

None.

Authority

- Super, Edit
- Any role granted the `rcopy_start` right

Usage

- Access to all domains is required to run this command.
- This command requires the HPE 3PAR Remote Copy license. Contact your local service provider for more information.
- This command must be executed before any other Remote Copy commands.

Examples

The following example starts Remote Copy on a system:

```
cli% startrcopy
```

startrcopygroup

Syntax

```
startrcopygroup [options] <group_name> [<vv:sv> ...]
startrcopygroup [options] <pattern>
```

Description

The `startrcopygroup` command enables Remote Copy for the specified Remote Copy volume groups.

Options

`-nosync`

Prevents the initial synchronization and sets the virtual volumes to a synchronized state.

`-wait`

Specifies that the command waits until the initial synchronization is completed. The system generates an event when the synchronization is completed.

If a group's target has the mirror_config policy set and the group is a secondary group, -wait option is not supported. If `startrcopygroup` command is run on secondary group with -wait specified, then the command will fail.

`-t <target_name>`

Only start the group on the specified target.

`-pat`

Specifies that specified patterns are treated as glob-style patterns and all Remote Copy groups matching the specified pattern will be started. -pat can specify a list of patterns. This option must be used if pattern specifier is used.

Specifiers

`<group_name>`

The name of the Remote Copy volume group. The group name can be obtained using the **`showrcopy`** command.

`<pattern>`

Specifies a glob-style pattern. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.) This specifier can be repeated to start multiple groups. If this specifier is not used, the <group_name> specifier must be used.

`<vv:sv>`

Member volumes and snapshots can be specified by vv:sv syntax, where vv is the base volume name and sv is the snapshot volume name. To indicate a full resync, specify the starting, read-only snapshot with "-".

When starting snapshots are applied, all vvs must be specified.

If a group's target has the mirror_config policy set and the group is a secondary group, `startrcopygroup` command will be forwarded to the corresponding primary group but starting snapshots (vv:sv) are not supported in this case.

Authority

- Super, Edit
- Any role granted the `rcopygroup_start` right

Usage

- This command requires the HPE 3PAR Remote Copy license. Contact your local service provider for more information.
- Synchronous groups are automatically synchronized when started. Asynchronous periodic volume groups are synchronized only on the first time they are started.
- You must enter this command on the secondary server before entering it on the primary if the mirror_config policy is not set.
- If a group's target has the mirror_config policy set and the group is a primary group, then this command will be mirrored to that target and the corresponding secondary group will be started. If the policy is set and the group is a secondary, then this command will be automatically forwarded to the corresponding primary group on remote side.
- If the mirror_config policy is not set, then the corresponding secondary group must already be started or this command will fail.
- When a secondary group says that it is in the started state, it does not necessarily mean that the group is receiving data from the primary system.
- The primary group might be stopped, or its system might even be down. The fact that the group is started only indicates that the group is ready to receive writes from the primary group. To determine if writes to the volumes of the primary group are being mirrored to the secondary group, the status of the primary group must be examined.

Examples

The following example starts Remote Copy for Group1:

```
cli% startrcopygroup Group1
```

To start Group2, which contains 4 vvs and specify starting snapshots, with vv4 starting from a full resync:

```
cli% startrcopygroup Group2 vv1:sv1 vv2:sv2 vv3:sv3 vv4:-
```

starttask

Syntax

```
starttask <command>
```

Description

Provides users with the ability to execute commands with long running times. The commands run in the background.

Options

None.

Specifiers

<command>

Specifies the command to run using `starttask`. Available commands are:

- `admithw`
- `creategroupsv`

- `createsv`
- `moverelocpd`
- `removevv`
- `updatevv`
- `upgradecage`
- `upgradepd`
- `startao`

If the command you choose to run requires confirmation, you must use the `-f` specifier to bypass user input. Otherwise, the command will fail.

Authority

- Super, Service, Edit
- Any role granted the `task_start` right

Examples

The following example shows the `removevv` command to be started with the forced option:

```
starttask removevv -f vv1
```

startvasa

Syntax

`startvasa`

Description

The `startvasa` command starts the VASA Provider server to service HTTPS requests. By default, the VASA Provider server is not started until this command is issued.

Options

None.

Specifiers

None.

Authority

- Super, Service
- Any role granted the `vasa_start` right

Usage

- Access to all domains is required to run this command.
- Use the **`stopvasa`** command to stop the VASA Provider server.
- The VASA Provider server only listens for HTTPS requests.

Examples

The following example starts the VASA Provider server:

```
cli% startvasa
```

The VASA Provider server will start shortly.

startvv

Syntax

```
startvv [options] <VV_name>
```

Description

The `startvv` command starts data services on a Virtual Volume (VV) that has not yet been started.

Options

`-ovrd`

Specifies that the logical disk is forced to start, even if some underlying data is missing.

Specifiers

`<VV_name>`

Specifies the VV name, using up to 31 characters.

Authority

- Super, Service
- Any role granted the `vv_start` right

Usage

Access to all domains is required to run this command.

Examples

None.

startwsapi

Syntax

```
startwsapi
```

Description

Starts the Web Services API server to support HTTP and HTTPS requests and to start the 3PAR GUI.. Both WSAPI and the UI default to `started`.

Options

None.

Specifiers

None.

Authority

- Super, Service
- Any role granted the `wsapi_start` right

Usage

- This command requires access to all domains.
- Use the `stopwsapi` command to stop the Web Services API server and the 3PAR GUI.
- The Web Services API server only listens for HTTPS requests.
- Use the `setwsapi` command to enable HTTP requests.

Examples

The following example starts the Web Services API server and the 3PAR GUI :

```
cli% startwsapi
The Web Services API server and the 3PAR GUI
 will start shortly.
```

statcache

Syntax

```
statcache [options <arg>]
```

Description

The `statcache` command displays Flash cache Memory Page (FMP) and Cache Memory Page (CMP) statistics by node or by virtual volume.

Options

`-v`

Specifies that CMP and FMP statistics by virtual volume instead of by node are displayed. By default, all virtual volumes are displayed unless the `-n` option is specified.

`-n <name|pattern>`

Specifies that statistics are displayed for virtual volumes matching the specified name or pattern. This option is valid only if `-v` is also specified.

`-domain {<domainname|pattern>}...`

Shows virtual volumes that are in domains with names that match one or more of the specified domains or patterns. If `-domain` is not specified, the virtual volumes that are in the current domain are shown. See the `currentdomain` parameter in the **`setclienv`** and **`showclienv`** commands. This option is only valid if the `-v` option is also specified.

`-metadata`

Specifies that statistics for metadata volumes (SA volumes) are displayed along with the statistics for user volumes. This option is only valid if `-v` is also specified.

`-d <seconds>`

Specifies the interval, in seconds, that statistics are sampled using an integer from 1 through 2147483. If no interval is specified, the option defaults to an interval of two seconds.

`-iter <number>`

Specifies that statistics are displayed a specified number of times as indicated by the number argument using an integer from 1 through 2147483647.

Specifiers

None.

Authority

Any role in the system

Usage

Patterns are glob-style (shell-style) patterns. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.)

Examples

The following example displays one iteration of cache statistics for all nodes:

```
cli% statcache -iter 1
```

The following example displays CMP (DRAM) and FMP (Flashcache) statistics on a node basis:

```
cli% statcache

15:12:35 04/22/2014 ------- Current --------    -------- Total ---------
                              CMP  FMP Total              CMP FMP Total
Node Type           Accesses Hit% Hit% Hit%      Accesses Hit% Hit% Hit%
    0 Read             0       0   0   0               0   0   0      0
    0 Write            0       0   0   0               0   0   0      0
    0 Read             0       0   0   0               0   0   0      0
    0 Write            0       0   0   0               0   0   0      0

                         Internal Flashcache Activity
                    ----- Current ------    ------- Total --------
Node Type           Accesses IO/s MB/s      Accesses IO/s MB/s
    0 Read Back            0    0    0             0    0    0
    0 Destaged Write       0    0    0             0    0    0
    1 Read Back            0    0    0             0    0    0
    1 Destaged Write       0    0    0             0    0    0

          --------------- FMP Queue Statistics ----------------
Node Dormant Cold Norm Warm Hot Destage Read Flush WrtBack
   0 8388608    0   0   0   0       0   0       0       0
   1 8388608    0   0   0   0       0   0       0       0
           -------------------- CMP Queue Statistics --------------------
Node    Free Clean Write1 WriteN WrtSched Writing DcowPend DcowProc
    0 1796802 58859     0   0       0           0       0       0
    1 1796521 58200     0   0       0           0       0       0
```

The following example displays CMP (DRAM) and FMP (Flashcache) statistics on a VV basis:

```
cli% statcache -v

15:14:29 04/22/2014 ------- Current -------- -------- Total ---------
                                 CMP FMP Total           CMP FMP Total
VVid VVname         Type Accesses Hit% Hit% Hit% Accesses Hit% Hit% Hit%
    0 admin         Read        0   0   0      0        0    0   0   0
    0 admin         Write       4   25  0      25       4    25  0   25
    1 .srdata       Read        0   0   0      0        0    0   0   0
    1 .srdata       Write       0   0   0      0        0    0   0   0
    5 VV25150_0001  Read        0   0   0      0        0    0   0   0
    5 VV25150_0001  Write       0   0   0      0        0    0   0   0
    7 VV25202_0001  Read        0   0   0      0        0    0   0   0
    7 VV25202_0001  Write       4   25  0      25       4    25  0   25
    8 VV25205_0001  Read        0   0   0      0        0    0   0   0
    8 VV25205_0001  Write       0   0   0      0        0    0   0   0
    9 VV25214_0000  Read        0   0   0      0        0    0   0   0
    9 VV25214_0000  Write       0   0   0      0        0    0   0   0
                    - Internal Flashcache Activity -
                                 - Current -- -- Total ---
VVid     VVname Type                Accesses Accesses
0         admin Read Back                  0        0
0         admin Destaged Write             0        0
1       .srdata Read Back                  0        0
1       .srdata Destaged Write             0        0
5       VV25150_0001 Read Back             0        0
5       VV25150_0001 Destaged Write        0        0
7       VV25202_0001 Read Back             0        0
7       VV25202_0001 Destaged Write        0        0
8       VV25205_0001 Read Back             0        0
8       VV25205_0001 Destaged Write        0        0
9       VV25214_0000 Read Back             0        0
9       VV25214_0000 Destaged Write        0        0
```

The following example displays CMP (DRAM) and FMP (Flashcache) statistics for both user volumes and metadata volumes:

```
cli% statcache -v -metadata

15:16:01 04/22/2014         ------- Current -------- -------- Total ---------
                                     CMP FMP Total                CMP FMP Total
VVid VVname Type         Accesses   Hit% Hit% Hit%            Accesses Hit% Hit% Hit%
   0 admin Read                 0      0    0    0                   0    0    0    0
   0 admin Write                0      0    0    0                   0    0    0    0
   0 admin MRead                0      0    0    0                   0    0    0    0
   0 admin MWrite               0      0    0    0                   0    0    0    0
   1 .srdata Read               0      0    0    0                   0    0    0    0
   1 .srdata Write              0      0    0    0                   0    0    0    0
   1 .srdata MRead              0      0    0    0                   0    0    0    0
   1 .srdata MWrite             0      0    0    0                   0    0    0    0
   5 VV25150_0001 Read          0      0    0    0                   0    0    0    0
   5 VV25150_0001 Write         0      0    0    0                   0    0    0    0
   5 VV25150_0001 MRead         0      0    0    0                   0    0    0    0
   5 VV25150_0001 MWrite        0      0    0    0                   0    0    0    0
   7 VV25202_0001 Read          0      0    0    0                   0    0    0    0
   7 VV25202_0001 Write         0      0    0    0                   0    0    0    0
   7 VV25202_0001 MRead         0      0    0    0                   0    0    0    0
   7 VV25202_0001 MWrite        0      0    0    0                   0    0    0    0
   8 VV25205_0001 Read          0      0    0    0                   0    0    0    0
   8 VV25205_0001 Write         0      0    0    0                   0    0    0    0
   8 VV25205_0001 MRead         0      0    0    0                   0    0    0    0
   8 VV25205_0001 MWrite        0      0    0    0                   0    0    0    0
   9 VV25214_0000 Read          0      0    0    0                   0    0    0    0
   9 VV25214_0000 Write         0      0    0    0                   0    0    0    0
   9 VV25214_0000 MRead         0      0    0    0                   0    0    0    0
   9 VV25214_0000 MWrite        0      0    0    0                   0    0    0    0
                                    - Internal Flashcache Activity -
                                       - Current -- -- Total ---
VVid VVname          Type (user+meta) Accesses Accesses
   0 admin           Read Back               0        0
   0 admin           Destaged Write          0        0
   1 .srdata         Read Back               0        0
   1 .srdata         Destaged Write          0        0
   5 VV25150_0001    Read Back               0        0
   5 VV25150_0001    Destaged Write          0        0
   7 VV25202_0001    Read Back               0        0
   7 VV25202_0001    Destaged Write          0        0
   8 VV25205_0001    Read Back               0        0
   8 VV25205_0001    Destaged Write          0        0
   9 VV25214_0000    Read Back               0        0
   9 VV25214_0000    Destaged Write          0        0
```

Use the following to get a description of each column in the above examples:

```
clihelp -col statcache
```

statch

Syntax

```
statch [options <arg>]
```

Description

The `statch` command displays chunklet statistics in a timed loop.

Options

-rw

Specifies that the display includes separate read and write data. If not specified, the total is displayed.

-d <secs>

Specifies the interval in seconds that statistics are sampled from using an integer from 1 through 2147483. If no count is specified, the command defaults to 2 seconds.

-iter <number>

Specifies that the statistics are to stop after the indicated number of iterations using an integer from 1 through 2147483647.

-begin

Specifies that I/O averages are computed from the system start time. If not specified, the average is computed since the first iteration of the command.

-idlep

Specifies the percent of idle columns in the output.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

-filt <fspec>

Specifies that statistics that fall below the threshold as specified by the <fspec> arguments, where <fspec> is <type>, <op>, <meas>, <val>, are filtered out and not displayed. The <type>, <op>, <meas>, and <val> arguments are separated with one comma.

- <type>

  The type argument can be specified with one of the following arguments:

  curs

  Specifies that only devices with current statistics above the threshold are displayed.

  avgs

  Specifies that only devices with average statistics above the threshold are displayed.

  maxs

  Specifies that only devices with maximum values above the threshold are displayed.

- <op>

  The operation argument can be specified with one of the following arguments:

  r

  Specifies that read-only statistics are displayed.

  w

  Specifies that write-only statistics are displayed.

t|rw

Specifies that statistics for read and write totals are displayed.

- `<meas>`

  The meas argument can be specified with one of the following arguments:

  iops

  Specifies that I/O operations per second are displayed. If this argument is used, the minimum threshold value must be specified using the `<val>` argument.

  bw

  Specifies that statistics for bandwidth in bytes per second are displayed. If this argument is used, the minimum threshold value must be specified using the `<val>` argument.

  svct

  Specifies that statistics for service time in milliseconds are displayed.

  size

  Specifies that statistics for I/O operations in bytes are displayed.

- `<val> [k|K]|[m|M]|[g|G]`

  Specifies the minimum threshold using any integer. The integer can optionally be followed with k or K to indicate a multiple of 1000, m or M to indicate a multiple of 1,000,000, or g or G to indicate a multiple of 1,000,000,000.

An example of this option is `-filt curs,r,iops,10k`. This indicates that statistics for a device that has more than 10,000 current read-only I/O operations are displayed.

-ni

Specifies that statistics for only non-idle devices are displayed. This option is shorthand for the option `-filt curs,t,iops,0`.

`-ld <LD_name>|-ch <chunk_num>`

-ld <LD_name>

Specifies that statistics are restricted to chunklets from a particular logical disk.

-ch <chunk_num>

Specifies that statistics are restricted to a particular chunklet number.

Specifiers

None.

Authority

Any role in the system

Usage

- If no options are used, the command defaults to show reads, writes, and totals separately for all chunklets in intervals of 2 seconds.
- The `statch` command can only be used after the **setstatch** command has been issued to enable chunklet statistics.
- For this command KB = 1000 bytes.

Examples

The following example displays the collection of statistics:

```
cli% statch -iter 1

                  04:08:35 07/15/2015 r/w I/O per second KBytes per sec      Svt ms IOSz KB
Ldid            Ldname LdCh Pdid PdCh      Cur  Avg  Max  Cur  Avg  Max   Cur    Avg Cur Avg Qlen
   0       admin.usr.0    0    1  407   t    0    0    0    0    0    0  0.00   0.00 0.0 0.0    0
   1       admin.usr.1    0    0  407   t    0    0    0    2    2    2 10.10  10.10 4.1 4.1    0
   4            log0.0    0   15  404   t    0    0    0    0    0    0  0.00   0.00 0.0 0.0    0
   5            log1.0    0   14  404   t    0    0    0    0    0    0  0.00   0.00 0.0 0.0    0
  11          pdsld0.0    0    1  406   t    0    0    0    0    0    0  0.00   0.00 0.0 0.0    0
  12          pdsld0.1    0   19  407   t    0    0    0    0    0    0  0.00   0.00 0.0 0.0    0
  13          pdsld0.2    0   12  406   t    0    0    0    0    0    0  0.00   0.00 0.0 0.0    0
  14      .srdata.usr.0    0   19  406   t    0    0    0    0    0    0  0.00   0.00 0.0 0.0    0
  15      .srdata.usr.1    0   12  405   t    0    0    0    0    0    0  0.00   0.00 0.0 0.0    0
  16           slog0.0    0    5    0   t    0    0    0    0    0    0  0.00   0.00 0.0 0.0    0
  17           slog1.0    0    8    0   t    0    0    0    0    0    0  0.00   0.00 0.0 0.0    0
  18        tp-0-sa-0.0    0   33    3   t    0    0    0    0    0    0  0.00   0.00 0.0 0.0    0
  19        tp-0-sa-0.1    0   18    3   t    0    0    0    0    0    0  0.00   0.00 0.0 0.0    0
  20        tp-0-sd-0.0    0   19    3   t    0    0    0    0    0    0  0.00   0.00 0.0 0.0    0
  21        tp-0-sd-0.1    0   12    3   t    0    0    0    0    0    0  0.00   0.00 0.0 0.0    0
  22 VV44326_0006.usr.0    0   19    4   t    0    0    0    0    0    0  0.00   0.00 0.0 0.0    0
  23 VV44326_0006.usr.1    0   10    4   t    0    0    0    0    0    0  0.00   0.00 0.0 0.0    0
  24 VV44326_0007.usr.0    0   21    5   t    0    0    0    0    0    0  0.00   0.00 0.0 0.0    0
  25 VV44326_0007.usr.1    0   12    4   t    0    0    0    0    0    0  0.00   0.00 0.0 0.0    0
------------------------------------------------------------------------------------------------
                    19                  t    0    0         2    2      10.10  10.10 4.1 4.1    0
```

For the previous example, before the `statch` command was issued, the `setstatch start` command was issued for chunklets 0 and 2 on logical disk vv0.usr.1 and for chunklets 0 and 1 on logical disk vv0.usr.2.

Use the following to get a description of each column in the above examples:

```
clihelp -col statch
```

statcmp

Syntax

```
statcmp [options <arg>]
```

Description

The `statcmp` command displays Cache Memory Page (CMP) statistics by node or by virtual volume.

Options

-v

Specifies that CMP statistics by virtual volume instead of by node are displayed. By default, all virtual volumes are displayed unless the -n option is specified.

-ni

Specifies that statistics for only non-idle VVs are displayed.

This option is valid only if -v is also specified.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified.

In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns. This option is valid only if -v is also specified.

-n <name|pattern>

Specifies that statistics are displayed for virtual volumes matching the specified name or pattern. This option is valid only if -v is also specified.

-domain {<domainname|pattern>}...

Shows virtual volumes that are in domains with names that match one or more of the specified domains or patterns. If -domain is not specified, the virtual volumes that are in the current domain are shown. This option is only valid if the -v option is also specified.

-d <seconds>

Specifies the interval, in seconds, that statistics are sampled using an integer from 1 through 2147483. If no interval is specified, the option defaults to an interval of two seconds.

-iter <number>

Specifies that CMP statistics are displayed a specified number of times as indicated by the number argument using an integer from 1 through 2147483647.

Specifiers

None.

Authority

Any role in the system

Usage

- Patterns are specified as regular expressions. (For more information, enter the `clihelp sub,regexpat` command.)
- Temporary and Page Credits refer to the number of credits being given by each node to other nodes in the system. The credits plus pages (free, clean, scheduled and writing) should add up to the total amount of memory in the node.

Examples

The following example displays one iteration of CMP statistics for all nodes:

```
cli% statcmp -iter 1

17:40:53 10/16/2019 ---- Current ----- --------- Total ----------
Node Type       Accesses Hits Hit% Accesses Hits Hit% LockBlk
0 Read           5800    0    0     5800    0    0       0
0 Write         11196   46    0    11196   46    0       0
1 Read           5904    0    0     5904    0    0       0
1 Write         11451    7    0    11451    7    0       0

Queue Statistics
Node   Free   Clean Write1 WriteN WrtSched Writing DcowPend DcowProc RcpyRev
0      210319 2900624      0  10762      953     328        0        0       0
1      210788 2900511      0  10519     1024     246        0        0       0

Temporary and Page Credits
Node  Node0  Node1 Node2 Node3 Node4 Node5 Node6 Node7
0      0 224264   ---   ---   ---   ---   ---   ---
1 224388      0   ---   ---   ---   ---   ---   ---

Page Statistics
-CfcDirty-- ---CfcMax--- ----DelAck-----
Node FC NL   SSD FC NL    SSD FC NL       SSD
0  0  0 13841  0  0 460800  0  0     19370716
1  0  0 13594  0  0 460800  0  0    129271135
```

Use the following to get a description of each column in the above examples:

```
clihelp -col statcmp
```

statcpu

Syntax

```
statcpu [options <arg>]
```

Description

The `statcpu` command displays CPU statistics for all nodes.

Options

`-d <secs>`

Specifies the interval, in seconds, that statistics are sampled using an integer from 1 through 2147483. If no interval is specified, the option defaults to an interval of two seconds.

`-iter <number>`

Specifies that CPU statistics are displayed a specified number of times as indicated by the number argument using an integer from 1 through 2147483647.

`-t`

Show only the totals for all the CPUs on each node

Specifiers

None.

Authority

Any role in the system

Usage

None.

Examples

The following example displays two iterations of CPU statistics for all nodes:

```
cli% statcpu -iter 2
15:11:03
node,cpu  user  sys   idle  intr/s  ctxt/s
0,0          0     0    100
0,1          0     0    100
0,total      0     0    100     162     412

1,0          0     1     99
1,1          0     1     99
1,total      0     1     99     158     269
Press the enter key to stop...

15:11:05
node,cpu user sys idle intr/s ctxt/s
0,0         0    0  100
0,1        19    7   74
0,total    10    4   87    183    444

1,0         0    0  100
```

```
1,1        0    0  100
1,total    0    0  100    158    235
Press the enter key to stop...
```

Use the following to get a description of each column in the above examples:

```
clihelp -col statcpu
```

statfcoe

Syntax

```
statfcoe [options <arg>]
```

Description

The `statfcoe` command displays Fibre Channel over Ethernet statistics.

Options

`-d <secs>`

Looping delay in seconds <secs>. The default is 2.

`-iter <number>`

The command stops after a user-defined <number> of iterations.

`-nodes <nodelist>`

Lists of nodes for which the ports are included.

`-slots <slotlist>`

Lists of PCI slots for which the ports are included.

`-ports <portlist>`

Lists of ports which are included. Lists are specified in a comma-separated manner such as: -ports 1,2 or -ports 1.

`-counts`

Shows the counts. The default is to show counts/sec.

`-fullcounts`

Shows the values for the full list of counters instead of the default packets and KBytes for the specified protocols. The values are shown in three columns:

- Current—Counts since the last sample.
- CmdStart—Counts since the start of the command.
- Begin—Counts since the port was reset.

`-prev`

Shows the differences from the previous sample.

`-begin`

Shows the values from when the system was last initiated.

Specifiers

None.

Authority

Any role in the system

Usage

None.

Examples

The following example displays one iteration of FCoE counts:

```
cli% statfcoe -iter 1 -counts -begin
```

Use the following to get a description of each column in the above examples:

```
clihelp -col statfcoe
```

statfs

Syntax

```
statfs [options <arg>]
```

Description

The `statfs` command displays statistics for File Persona.

Options

`-av`

Displays Antivirus statistics.

`-block`

Displays statistics for Block Devices.

`-fpg`

Displays statistics for File Provisioning Group.

`-net`

Displays statistics for networking.

`-nfs`

Displays statistics on NFS filesystems.

`-snapshot`

Displays statistics for snapshots.

`-smb`

Displays statistics on SMB filesystems.

`-cpu`

Displays statistics for the CPU.

`-mem`

Displays memory statistics.

`-http`

Displays HTTP statistics.

`-iter <number>`

Specifies that the statistics are to stop after the indicated number of iterations using an integer from 1 through 2147483647.

-d <secs>

Specifies the interval in seconds that statistics are sampled from using an integer from 1 through 2147483. If no count is specified, the command defaults to 60 seconds. Information will only be updated every 60 seconds, times of less will redisplay same information.

-node <nodeid>[,<nodeid>]...

Specifies the node on which to report statistics.

-verbose

Specifies that all statistics will be displayed. Each statistic will be displayed on a line each.

Specifiers

None.

Authority

Any role in the system

Usage

Access to all domains is required to run this command.

Examples

The following example displays 5 iterations of NFS file systems statistics at the default sample period:

```
cli% statfs -nfs -iter 5
```

Use the following to get a description of each column in the above examples:

```
clihelp -col statfs
```

statiscsi

Syntax

```
statiscsi [options <arg>]
```

Description

The `statiscsi` command displays the iSCSI statistics.

Options

-d <secs>

Looping delay in seconds `<secs>`. The default is 2.

-iter <number>

The command stops after a user-defined `<number>` of iterations.

-nodes <nodelist>

Lists of nodes for which the ports are included.

-slots <slotlist>

Lists of PCI slots for which the ports are included.

-ports <portlist>

Lists of ports for which the ports are included. Lists are specified in a comma-separated manner such as: `-ports 1,2` or `-ports 1`.

-counts

Shows the counts. The default is to show counts/sec.

-fullcounts

Shows the values for the full list of counters instead of the default packets and KBytes for the specified protocols. The values are shown in three columns:

- `Current`—Counts since the last sample.
- `CmdStart`—Counts since the start of the command.
- `Begin`—Counts since the port was reset.

This option cannot be used with the `-prot` option. If the `-fullcounts` option is not specified, the metrics from the start of the command are displayed.

-prot <prot>[,<prot>,...]

Shows the statistics for the specified protocols. The available protocols are:

- `Eth`—Ethernet.
- `IP`—Internet Protocol (IP).
- `TCP`—Transmission Control Protocol (TCP).
- `iSCSI`—iSCSI.
- `all`—All protocols (default).

This option cannot be used with the `-fullcounts` option.

-prev

Shows the differences from the previous sample.

-begin

Shows the values from when the system was last initiated.

Specifiers

None.

Authority

Any role in the system

Usage

The statistics for the TCP, IP, and iSCSI protocols listed are the combined values for both ports of the HBA. Each port reports the combined stats for both ports for these protocols. The total reported at the bottom for these protocols is therefore twice the actual values.

Examples

The following example displays one iteration of iSCSI counts (with Ethernet and transmission control protocols):

```
cli% statiscsi -iter 1 -counts -prot Eth,TCP -begin
13:24:38 03/29/06 --------------------From last port reset---------------------
                  ------Receive------ ----Transmit----   -------Total-------
port  Protocol      Pkts    KBytes     Pkts  KBytes      Pkts    KBytes Errs
1:3:1      Eth  766373.0 1066175.9 394473.0 55171.4 1160846.0 1121347.2  1.0
1:3:1      TCP  722917.0 1016532.4 394441.0 30238.3 1117358.0 1046770.7  0.0
1:3:2      Eth       0.0       0.0      0.0     0.0       0.0       0.0  0.0
1:3:2      TCP  722917.0 1016532.4 394441.0 30238.3 1117358.0 1046770.7  0.0
```

```
------------------------------------------------------------------------------
Total      Eth  766373.0 1066175.9 394473.0 55171.4 1160846.0 1121347.2  1.0
Total      TCP 1445834.0 2033064.8 788882.0 60476.6 2234716.0 2093541.4  0.0*
```

The following example displays basic iSCSI statistics collection:

```
cli% statiscsi
16:37:59 04/06/06 ----Receive----   ---Transmit----  -----Total-----
port  Protocol    Pkts/s KBytes/s   Pkts/s KBytes/s  Pkts/s  KBytes/s Errs/s
0:4:1      Eth   15633.2   1116.7  30115.9  45164.4 45749.1   46281.1    0.0
0:4:1       IP   15632.7    522.6  30116.4  44020.8 45749.1   44543.3    0.0
0:4:1      TCP   15632.7     22.3  30116.4  43057.0 45749.1   43079.4    0.0
0:4:1    iSCSI     547.7      0.0   1477.3  50452.8  2025.0   50452.8    0.0
0:4:2      Eth       0.5      0.0      0.0      0.0     0.5       0.0    0.0
0:4:2       IP   15632.1    522.6  30115.2  44019.0 45747.2   44541.5    0.0
0:4:2      TCP   15632.1     22.3  30115.2  43055.3 45747.2   43077.6    0.0
0:4:2    iSCSI     547.7      0.0   1478.7  50509.7  2026.4   50509.7    0.0
1:3:1      Eth   11307.2    807.0  21348.2  32048.1 32655.4   32855.1    0.0
1:3:1       IP   11309.2    377.4  21345.7  31233.2 32654.9   31610.5    0.0
1:3:1      TCP   11309.2     15.5  21344.7  30548.7 32653.9   30564.1    0.0
1:3:1    iSCSI     459.3      0.0   1247.9  42630.9  1707.2   42630.9    0.0
1:3:2      Eth       0.5      0.0      0.0      0.0     0.5       0.0    0.0
1:3:2       IP   11308.5    377.3  21344.5  31231.3 32653.0   31608.7    0.0
1:3:2      TCP   11308.5     15.5  21343.5  30546.9 32652.0   30562.3    0.0
1:3:2    iSCSI     459.3      0.0   1247.8  42628.4  1707.1   42628.4    0.0
------------------------------------------------------------------------------
Total      Eth   26941.4   1923.7  51464.1  77212.5  78405.5  79136.2    0.0
Total       IP   53882.5   1799.8 102921.7 150504.2 156804.2 152304.1    0.0
Total      TCP   53882.5     75.6 102919.7 147207.9 156802.2 147283.5    0.0
Total    iSCSI    2014.0      0.0   5451.7 186221.9   7465.7 186221.9    0.0
Press the enter key to stop...
```

statiscsisession

Syntax

`statiscsisession [options <arg>]`

Description

The `statiscsisession` command displays the iSCSI session statistics.

Options

`-d <secs>`

Looping delay in seconds <secs>. The default is 2.

`-iter <number>`

The command stops after a user-defined <number> of iterations.

`-nodes <nodelist>`

Lists of nodes for which the ports are included.

`-slots <slotlist>`

Lists of PCI slots for which the ports are included.

`-ports <portlist>`

Lists of ports for which the ports are included. Lists are specified in a comma-separated manner such as: -ports 1,2 or -ports 1.

-counts

Shows the counts. The default is to show counts/sec.

By default, the differences from initiating the command are shown. The following options change that behavior:

-prev

Shows the differences from the previous sample.

-begin

Shows the values from when the system was last initiated.

Specifiers

None.

Authority

Any role in the system

Usage

None.

Examples

The following example displays the session statistics for the iSCSI:

```
cli% statiscsisession
08:17:47 06/05/2015 --From start of statiscsisession command---
                    ----PDUs/s---- --KBytes/s--- ----Errs/s----
port -iSCSI_Name- TPGT Cmd Resp Total  Tx  Rx Total Digest TimeOut VLAN
-------------------------------------------------------------------------
Total             -    - 0.0  0.0   0.0 0.0 0.0   0.0    0.0     0.0    -
Press the enter key to stop...
```

statld

Syntax

```
statld [options <arg>] [<LD_name|pattern>...]
```

Description

The `statld` command displays read/write (I/O) statistics about LDs in a timed loop.

Options

-vv {<VV_name>|<pattern>}...

Shows only logical disks that are mapped to virtual volumes with names matching any of the names or patterns specified. Multiple volumes or patterns can be repeated using a comma separated list.

-domain {<domain_name>|<pattern>}...

Shows only logical disks that are in domains with names matching any of the names or patterns specified. Multiple domain names or patterns can be repeated using a comma separated list.

-rw

Specifies that the display includes separate read and write data. If not specified, the total is displayed.

-d <secs>

Specifies the interval in seconds that statistics are sampled from using an integer from 1 through 2147483. If no count is specified, the command defaults to 2 seconds.

-iter <number>

Specifies that the statistics are to stop after the indicated number of iterations using an integer from 1 through 2147483647.

-begin

Specifies that I/O averages are computed from the system start time. If not specified, the average is computed since the first iteration of the command.

-idlep

Specifies the percent of idle columns in the output.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

-filt <fspec>

Specifies that statistics that fall below the threshold as specified by the <fspec> arguments, where <fspec> is <type>, <op>, <meas>, <val>, are filtered out and not displayed. The <type>, <op>, <meas>, and <val> arguments are separated with one comma.

- <type>

  The type argument can be specified with one of the following arguments:

  curs

  Specifies that only devices with current statistics above the threshold are displayed.

  avgs

  Specifies that only devices with average statistics above the threshold are displayed.

  maxs

  Specifies that only devices with maximum values above the threshold are displayed.

- <op>

  The operation argument can be specified with one of the following arguments:

  r

  Specifies that read-only statistics are displayed.

  w

  Specifies that write-only statistics are displayed.

  t|rw

  Specifies that statistics for read and write totals are displayed.

- <meas>

  The meas argument can be specified with one of the following arguments:

iops

Specifies that I/O operations per second are displayed. If this argument is used, the minimum threshold value must be specified using the `<val>` argument.

bw

Specifies that statistics for bandwidth in bytes per second are displayed. If this argument is used, the minimum threshold value must be specified using the `<val>` argument.

svct

Specifies that statistics for service time in milliseconds are displayed.

size

Specifies that statistics for I/O operations in bytes are displayed.

- `<val> [k|K]|[m|M]|[g|G]`

  Specifies the minimum threshold using any integer. The integer can optionally be followed with k or K to indicate a multiple of 1000, m or M to indicate a multiple of 1,000,000, or g or G to indicate a multiple of 1,000,000,000.

An example of this option is `-filt curs,r,iops,10k`. This indicates that statistics for a device that has more than 10,000 current read-only I/O operations are displayed.

-ni

Specifies that statistics for only non-idle devices are displayed. This option is shorthand for the option `-filt curs,t,iops,0`.

Specifiers

<LD_name|pattern>...

Only statistics are displayed for the specified LD or pattern. Multiple LDs or patterns can be repeated.

Authority

Any role in the system

Usage

- If no option is specified, the command defaults to display statistics totals.
- For this command KB = 1000 bytes.
- If the `<LD_name_or_pattern>` specifier is used, then logical disks with names that match any of the patterns are listed, otherwise all logical disks are listed. These patterns are glob-style patterns. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.)

Examples

The following example displays one iteration of I/O statistics for all LDs:

```
cli% statld -iter 1
08:40:45 06/04/2015 r/w I/O per second   KBytes per sec      Svt ms      IOSz KB
             Ldname      Cur  Avg  Max   Cur   Avg  Max   Cur   Avg   Cur   Avg Qlen
        admin.usr.0   t    0    0    0     4     4    4 11.41 11.41   8.2   8.2    0
      .srdata.usr.0   t    1    1    1     4     4    4  9.59  9.59   4.1   4.1    0
             log0.0   t    0    0    0     0     0    0  0.00  0.00   0.0   0.0    0
```

Use the following to get a description of each column in the above examples:

```
clihelp -col statld
```

statlink

Syntax

```
statlink [options <arg>]
```

Description

The `statlink` command displays statistics for link utilization for all nodes in a timed loop.

Options

`-d <sec>`

Specifies the interval, in seconds, that statistics are sampled using an integer from 1 through 2147483. If no interval is specified, the option defaults to an interval of two seconds.

`-iter <number>`

Specifies that internode link statistics are displayed a specified number of times (from 1 through 2147483647).

`-detail`

Displays detailed information regarding the Queue statistics.

Authority

Any role in the system

Usage

For HPE 3PAR 10000 systems, the Queue column is only shown with -detail option. Also for HPE 3PAR 10000 systems, the -From- and -To- headings indicate the Node,ASIC, and Queue columns from/to which the statistics are collected.

Examples

Display one iteration of statistics for the internode link:

```
cli% statlink -iter
Node Q ToNode Cur       Avg         Max     Cur  Avg    Max Cur Avg
0   CM      0 15134     14911       15134 61951 61038 61951 4.1 4.1
0   PCI0    0 22336     21852       22336 79831 77128 79831 3.6 3.5
0   PCI1    1 15312     12103       19805 62106 63198 77324 4.0 4.1
0   L0      1 10339     10307       10339 69673 69384 69673 6.7 6.7
1   CM      1 16372     16189       16372 67020 66269 67020 4.1 4.1
1   PCI0    1 18384     17899       18384 75285 73305 75285 4.1 4.1
1   PCI1    1 5329      5336        5344    984 985   987   0.2 0.2
1   L3      0 10326     10295       10326 69346 69150 69346 6.7 6.7
```

Get a description of each column in the above examples:

```
clihelp -col statlink
```

statpd

Syntax

```
statpd [options]
```

Description

The `statpd` command displays the read/write (I/O) statistics for PDs in a timed loop.

Options

-w <WWN>

Specifies that statistics for a particular PD identified by WWN are displayed.

-nodes <node_list>

Specifies that the display is limited to specified nodes and physical disks connected to those nodes. The node list is specified as a series of integers separated by commas (e.g. 1,2,3). The list can also consist of a single integer. If the node list is not specified, all disks on all nodes are displayed.

-slots <slot_list>

Specifies that the display is limited to specified PCI slots and physical disks connected to those PCI slots. The slot list is specified as a series of integers separated by commas (e.g. 1,2,3). The list can also consist of a single integer. If the slot list is not specified, all disks on all slots are displayed.

-ports <port_list>

Specifies that the display is limited to specified ports and physical disks connected to those ports. The port list is specified as a series of integers separated by commas (e.g. 1,2,3). The list can also consist of a single integer. If the port list is not specified, all disks on all ports are displayed.

-devinfo

Indicates the device disk type and speed.

-rw

Specifies that the display includes separate read and write data. If not specified, the total is displayed.

-d <secs>

Specifies the interval in seconds that statistics are sampled from using an integer from 1 through 2147483. If no count is specified, the command defaults to 2 seconds.

-iter <number>

Specifies that the statistics are to stop after the indicated number of iterations using an integer from 1 through 2147483647.

-begin

Specifies that I/O averages are computed from the system start time. If not specified, the average is computed since the first iteration of the command.

-idlep

Specifies the percent of idle columns in the output.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (`<col>`). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (`<dir>`) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

-filt <fspec>

Specifies that statistics that fall below the threshold as specified by the `<fspec>` arguments, where `<fspec>` is `<type>`, `<op>`, `<meas>`, `<val>`, are filtered out and not displayed. The `<type>`, `<op>`, `<meas>`, and `<val>` arguments are separated with one comma.

- <type>

  The type argument can be specified with one of the following arguments:

  curs

  Specifies that only devices with current statistics above the threshold are displayed.

  avgs

  Specifies that only devices with average statistics above the threshold are displayed.

  maxs

  Specifies that only devices with maximum values above the threshold are displayed.

- <op>

  The operation argument can be specified with one of the following arguments:

  r

  Specifies that read-only statistics are displayed.

  w

  Specifies that write-only statistics are displayed.

  t|rw

  Specifies that statistics for read and write totals are displayed.

- <meas>

  The meas argument can be specified with one of the following arguments:

  iops

  Specifies that I/O operations per second are displayed. If this argument is used, the minimum threshold value must be specified using the <val> argument.

  bw

  Specifies that statistics for bandwidth in bytes per second are displayed. If this argument is used, the minimum threshold value must be specified using the <val> argument.

  svct

  Specifies that statistics for service time in milliseconds are displayed.

  size

  Specifies that statistics for I/O operations in bytes are displayed.

- <val> [k|K]|[m|M]|[g|G]

  Specifies the minimum threshold using any integer. The integer can optionally be followed with k or K to indicate a multiple of 1000, m or M to indicate a multiple of 1,000,000, or g or G to indicate a multiple of 1,000,000,000.

An example of this option is -filt curs,r,iops,10k. This indicates that statistics for a device that has more than 10,000 current read-only I/O operations are displayed.

-ni

Specifies that statistics for only non-idle devices are displayed. This option is shorthand for the option -filt curs,t,iops,0.

-devsvtime

Specifies the time spent in the drive, not including queue time. Shows the average and current time in ms.

-p <pattern>

Specifies that statistics for a pattern of physical disks are displayed.

Patterns are used to filter and select the disks from which the statistics are collected. If specified multiple times, each instance of the specified pattern adds additional candidate disks matching the pattern.

The following arguments can be specified as patterns for this option:

An item is specified as an integer, a comma-separated list of integers, or a range of integers specified from low to high.

-nd <item>

Specifies one or more nodes. Nodes are identified by one or more integers (item). Multiple nodes are separated with a single comma (e.g. 1,2,3). A range of nodes is separated with a hyphen (e.g. 0-7). The primary path of the disks must be on the specified node(s).

-st <item>

Specifies one or more PCI slots. Slots are identified by one or more integers (item). Multiple slots are separated with a single comma (e.g. 1,2,3). A range of slots is separated with a hyphen (e.g. 0-7). The primary path of the disks must be on the specified PCI slot(s).

-pt <item>

Specifies one or more ports. Ports are identified by one or more integers (item). Multiple ports are separated with a single comma (e.g. 1,2,3). A range of ports is separated with a hyphen (e.g. 0-4). The primary path of the disks must be on the specified port(s).

-cg <item>

Specifies one or more drive cages. Drive cages are identified by one or more integers (item). Multiple drive cages are separated with a single comma (e.g. 1,2,3). A range of drive cages is separated with a hyphen (e.g. 0-3). The specified drive cage(s) must contain disks.

-mg <item>

Specifies one or more drive magazines. The "1." or "0." displayed in the `CagePos` column of `showpd` output indicating the side of the cage is omitted when using the `-mg` option. Drive magazines are identified by one or more integers (item). Multiple drive magazines are separated with a single comma (e.g. 1,2,3). A range of drive magazines is separated with a hyphen (e.g. 0-7). The specified drive magazine(s) must contain disks.

-pn <item>

Specifies one or more disk positions within a drive magazine. Disk positions are identified by one or more integers (item). Multiple disk positions are separated with a single comma (e.g. 1,2,3). A range of disk positions is separated with a hyphen (e.g. 0-3). The specified position(s) must contain disks.

-dk <item>

Specifies one or more physical disks. Disks are identified by one or more integers(item). Multiple disks are separated with a single comma (e.g. 1,2,3). A range of disks is separated with a hyphen (e.g. 0-3). Disks must match the specified ID(s).

-tc_gt <number>

Specifies that physical disks with total chunklets greater than the number specified be selected.

-tc_lt <number>

Specifies that physical disks with total chunklets less than the number specified be selected.

-fc_gt <number>

Specifies that physical disks with free chunklets greater than the number specified be selected.

-fc_lt <number>

Specifies that physical disks with free chunklets less than the number specified be selected.

-devid <model>

Specifies that physical disks identified by their models be selected. Models can be specified in a comma-separated list. Models can be displayed by issuing the "`showpd -i`" command.

-devtype <type>

Specifies that physical disks must have the specified device type (FC for Fast Class, NL for Nearline, SSD for Solid State Drive, SCM for Storage Class Memory) to be used. Device types can be displayed by issuing the `showpd` command.

-rpm <number>

Disks must be of the specified speed. Device speeds are shown in the RPM column of the `showpd` command. For FC and NL drives, the number corresponds to both a performance measure and actual rotational speed. For SSD & SCM drives, this field is "N/A".

Disks that satisfy all of the specified characteristics are used.

For example `-p -fc_gt 60 -fc_lt 230 -nd 2` specifies all the disks that have greater than 60 and less than 230 free chunklets and that are connected to node 2 through their primary path.

Specifiers

None.

Authority

Any role in the system

Usage

For this command KB = 1000 bytes.

Examples

The following example displays one iteration of I/O statistics for all physical disks:

```
cli% statpd -iter 1
02:37:19 06/11/2015 r/w I/O per second KBytes per sec      Svt ms   IOSz KB       Idle %
     ID        Port      Cur  Avg  Max  Cur  Avg  Max   Cur   Avg  Cur  Avg Qlen Cur Avg
      0       1:0:1   t    0    0    0    0    0    0  0.00  0.00  0.0  0.0    0 100 100
...
------------------------------------------------------------------------------------------
     48               t    2    2        28   28       9.43  9.43 14.3 14.3    0 100 100
```

Use the following to get a description of each column in the above examples:

```
clihelp -col statpd
```

statport

Syntax

```
statport [options]
```

Description

The `statport` command displays read/write (I/O) statistics for ports.

Options

-both | -ctl | -data

Specifies that both control and data transfers are displayed (`-both`), only control transfers are displayed (`-ctl`), or only data transfers are displayed (`-data`). If this option is not specified, only data transfers are displayed.

-nodes <node_list>

Specifies that the display is limited to specified nodes and physical disks connected to those nodes. The node list is specified as a series of integers separated by commas (e.g. 1,2,3). The list can also consist of a single integer. If the node list is not specified, all disks on all nodes are displayed.

-slots <slot_list>

Specifies that the display is limited to specified PCI slots and physical disks connected to those PCI slots. The slot list is specified as a series of integers separated by commas (e.g. 1,2,3). The list can also consist of a single integer. If the slot list is not specified, all disks on all slots are displayed.

-ports <port_list>

Specifies that the display is limited to specified ports and physical disks connected to those ports. The port list is specified as a series of integers separated by commas (e.g. 1,2,3). The list can also consist of a single integer. If the port list is not specified, all disks on all ports are displayed.

-host|-disk|-rcfc|-peer

Specifies to display only host ports (target ports), only disk ports (initiator ports), only Fibre Channel Remote Copy configured ports, or only Fibre Channel ports for Data Migration. If no option is specified, all ports are displayed.

The `-rcfc` option includes only statistics for Remote Copy over Fibre Channel ports related to cached READ requests. These cached READ requests will timeout after 60 seconds if not used to service Remote Copy requests so higher service times may be detected on this port.

-rcip

Includes only statistics for Ethernet configured Remote Copy ports. This option cannot be used with any other port mode selection options.

-fs

Includes only statistics for File Persona ports.

-rw

Specifies that the display includes separate read and write data. If not specified, the total is displayed.

-d <secs>

Specifies the interval in seconds that statistics are sampled from using an integer from 1 through 2147483. If no count is specified, the command defaults to 2 seconds.

-iter <number>

Specifies that the statistics are to stop after the indicated number of iterations using an integer from 1 through 2147483647.

-begin

Specifies that I/O averages are computed from the system start time. If not specified, the average is computed since the first iteration of the command.

-idlep

Specifies the percent of idle columns in the output.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (`<col>`). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (`<dir>`) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

-filt <fspec>

Specifies that statistics that fall below the threshold as specified by the <fspec> arguments, where <fspec> is <type>, <op>, <meas>, <val>, are filtered out and not displayed. The <type>, <op>, <meas>, and <val> arguments are separated with one comma.

- <type>

  The type argument can be specified with one of the following arguments:

  curs

  Specifies that only devices with current statistics above the threshold are displayed.

  avgs

  Specifies that only devices with average statistics above the threshold are displayed.

  maxs

  Specifies that only devices with maximum values above the threshold are displayed.

- <op>

  The operation argument can be specified with one of the following arguments:

  r

  Specifies that read-only statistics are displayed.

  w

  Specifies that write-only statistics are displayed.

  t|rw

  Specifies that statistics for read and write totals are displayed.

- <meas>

  The meas argument can be specified with one of the following arguments:

  iops

  Specifies that I/O operations per second are displayed. If this argument is used, the minimum threshold value must be specified using the <val> argument.

  bw

  Specifies that statistics for bandwidth in bytes per second are displayed. If this argument is used, the minimum threshold value must be specified using the <val> argument.

  svct

  Specifies that statistics for service time in milliseconds are displayed.

  size

  Specifies that statistics for I/O operations in bytes are displayed.

- <val> [k|K]|[m|M]|[g|G]

Specifies the minimum threshold using any integer. The integer can optionally be followed with k or K to indicate a multiple of 1000, m or M to indicate a multiple of 1,000,000, or g or G to indicate a multiple of 1,000,000,000.

An example of this option is `-filt curs,r,iops,10k`. This indicates that statistics for a device that has more than 10,000 current read-only I/O operations are displayed.

`-ni`

Specifies that statistics for only non-idle devices are displayed. This option is shorthand for the option `-filt curs,t,iops,0`.

Specifiers

None.

Authority

Any role in the system

Usage

- The `-filt` option applies only to data transfers, not control transfers.
- For this command KB = 1000 bytes.
- When a physical port is hosting multiple virtual ports, e.g. when port failover is active, `statport` indicates the aggregate statistics for all virtual ports on the physical port.

Examples

The following example displays one iteration of I/O statistics for all ports:

```
cli% statport -iter 1
03:53:52 07/16/2015 r/w I/O per second       KBytes per sec            Svt ms     IOSz KB
     Port       D/C      Cur  Avg  Max     Cur     Avg     Max      Cur      Avg   Cur   Avg Qlen
    0:0:1      Data   t    2    2    2      34      34      34     13.0     13.0  17.4  17.4    0
    0:0:2      Data   t    0    0    0       4       4       4      8.2      8.2   8.2   8.2    0
    0:2:1      Data   t    0    0    0       0       0       0      0.0      0.0   0.0   0.0    0
    0:2:2      Data   t    0    0    0       0       0       0      0.0      0.0   0.0   0.0    0
    0:3:1      Data   t    0    0    0       0       0       0      0.0      0.0   0.0   0.0    0
    0:3:2      Data   t    0    0    0       0       0       0      0.0      0.0   0.0   0.0    0
    1:0:1      Data   t    0    0    0       4       4       4      9.7      9.7   8.2   8.2    0
    1:0:2      Data   t    2    2    2      34      34      34     11.4     11.4  17.4  17.4    0
    1:3:1      Data   t    0    0    0       0       0       0      0.0      0.0   0.0   0.0    0
    1:3:2      Data   t    0    0    0       0       0       0      0.0      0.0   0.0   0.0    0
------------------------------------------------------------------------------------------------
      10       Data   t    5    5           76      76            11.6     11.6  15.6  15.6    0
```

Use the following to get a description of each column in the above examples:

```
clihelp -col statport
```

statqos

Syntax

```
statqos [options] [{{vvset|domain}:{<name>|<pattern>}|sys:all_others}]...
```

Description

The `statqos` command display the run-time stats of QoS rules that are "on".

Options

-domain {<domainname|pattern>}...

Shows statistics for rules with target objects in domains with names that match one or more of the specified domains or patterns. If -domain is not specified, statistics for all rules with target objects in the current domain are shown. See the currentdomain parameter in the **setclienv** and **showclienv** commands.

-vv {<VV_name>|<pattern>}[,{<VV_name>|<pattern>}]...

Applies only to QoS rules whose targets include virtual volumes with names matching any of the names or patterns specified.

-rw

Specifies that the display includes separate read and write data. If not specified, the total is displayed.

-d <secs>

Specifies the interval in seconds that statistics are sampled from using an integer from 1 through 2147483. If no count is specified, the command defaults to 2 seconds.

-iter <number>

Specifies that the statistics are to stop after the indicated number of iterations using an integer from 1 through 2147483647.

-begin

Specifies that I/O averages are computed from the system start time. If not specified, the average is computed since the first iteration of the command.

-idlep

Specifies the percent of idle columns in the output.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

-filt <fspec>

Specifies that statistics that fall below the threshold as specified by the <fspec> arguments, where <fspec> is <type>, <op>, <meas>, <val>, are filtered out and not displayed. The <type>, <op>, <meas>, and <val> arguments are separated with one comma.

- <type>

  The type argument can be specified with one of the following arguments:

  curs

  Specifies that only devices with current statistics above the threshold are displayed.

  avgs

  Specifies that only devices with average statistics above the threshold are displayed.

maxs

Specifies that only devices with maximum values above the threshold are displayed.

- `<op>`

  The operation argument can be specified with one of the following arguments:

  r

  Specifies that read-only statistics are displayed.

  w

  Specifies that write-only statistics are displayed.

  t|rw

  Specifies that statistics for read and write totals are displayed.

- `<meas>`

  The meas argument can be specified with one of the following arguments:

  iops

  Specifies that I/O operations per second are displayed. If this argument is used, the minimum threshold value must be specified using the <val> argument.

  bw

  Specifies that statistics for bandwidth in bytes per second are displayed. If this argument is used, the minimum threshold value must be specified using the <val> argument.

  svct

  Specifies that statistics for service time in milliseconds are displayed.

  size

  Specifies that statistics for I/O operations in bytes are displayed.

- `<val> [k|K]|[m|M]|[g|G]`

  Specifies the minimum threshold using any integer. The integer can optionally be followed with k or K to indicate a multiple of 1000, m or M to indicate a multiple of 1,000,000, or g or G to indicate a multiple of 1,000,000,000.

An example of this option is `-filt curs,r,iops,10k`. This indicates that statistics for a device that has more than 10,000 current read-only I/O operations are displayed.

-ni

Specifies that statistics for only non-idle devices are displayed. This option is shorthand for the option `-filt curs,t,iops,0`.

Specifiers

{{vvset|domain}:{<name>|<pattern>}|sys:all_others}

The QoS rules to display run-time stats. If none is specified, it displays all active QoS rules. <name> and <pattern> refer to the target object name of the rule, as listed in show {vvset|domain} commands.

Authority

Any role in the system

Usage

- For this command KB = 1000 bytes.
- For all the <pattern> fields, the patterns are glob-style (shell-style) patterns. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.)
- A Domain column may be included by using the **setclienv** command to set listdom to 1. The listdom option is also set to 1 if the CLI was started with the -listdom option, or if the CLI was started with the TPDLISTDOM environment variable set. Run "`cli -h`" and "`setclienv -h`" for details of the environment variables.

Examples

The following example displays one iteration of statistics for QoS:

```
cli% statqos -iter 1
```

Use the following to get a description of each column in the above examples:

```
clihelp -col statqos
```

statrcopy

Syntax

```
statrcopy [options <arg>]
```

Description

The `statrcopy` command displays statistics for Remote Copy volume groups.

Options

`-d <secs>`

Specifies the interval, in seconds, that statistics are sampled using an integer from 1 through 2147483647. If no interval is specified, the option defaults to an interval of two seconds.

`-iter <number>`

Specifies that I/O statistics are displayed a specified number of times as indicated by the num argument using an integer from 1 through 2147483647.

`-u k|m|g`

Displays statistics as kilobytes (k), megabytes (m), or gigabytes (g).

If no unit is specified, the default is kilobytes.

`-hb`

Specifies that the heartbeat round-trip time of the links should be displayed in addition to the link throughput.

Specifiers

None.

Authority

Any role in the system

Usage

- This command requires the HPE 3PAR Remote Copy license. Contact your local service provider for more information.
- This command repeats until directed to stop.

- Within the context of this command, KB is 1000 bytes, MB is 1000 KB, and GB is 1000 MB.
- The numbers displayed by this command might be somewhat less than those displayed with the **statport** command, as `statrcopy` output does not include TCP/IP overhead.

Examples

The following example shows statistics for all Remote Copy links:

```
cli% statrcopy
```

Use the following to get a description of each column in the above examples:

```
clihelp -col statrcopy
```

statrcvv

Syntax

```
statrcvv [options <arg>] [<VV_name>|<pattern>]...
```

Description

The `statrcvv` command displays statistics for Remote Copy volumes in a timed loop.

Options

-d <secs>

Specifies the interval in seconds that statistics are sampled from using an integer from 1 through 2147483. If no count is specified, the command defaults to 2 seconds.

-iter <number>

Specifies that the statistics are to stop after the indicated number of iterations using an integer from 1 through 2147483647.

-t <target_name>...|<pattern>...

Shows only volumes whose group is copied to the specified target name or pattern. Multiple target names or patterns may be specified using a comma-separated list.

-port <port_NSP>...|<pattern>...

Shows only volumes which are copied over the specified port or pattern. Multiple ports or patterns may be specified using a comma-separated list.

-g <group_name>...|<pattern>...

Shows only volumes whose group matches the specified group name or pattern. Multiple group names or patterns may be specified using a comma-separated list.

-async

Shows only volumes which are being copied in asynchronous mode.

-sync

Shows only volumes which are being copied in synchronous mode.

-periodic

Shows only volumes which are being copied in periodic mode.

-primary

Shows only volumes which are in the primary role.

-secondary

Shows only volumes which are in the secondary role.

-targetsum

Specifies that the sums for all volumes of a target are displayed.

-portsum

Specifies that the sums for all volumes on a port are displayed.

-groupsum

Specifies that the sums for all volumes of a group are displayed.

-vvsum

Specifies that the sums for all targets and links of a volume are displayed.

-domainsum

Specifies that the sums for all volumes of a domain are displayed.

-domain <domainname>...|<pattern>...

Shows only the virtual volumes that are in domains with names that match the specified domain name(s) or pattern(s).

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

-filt <fspec>

Specifies that statistics that fall below the threshold as specified by the <fspec> arguments, where <fspec> is <type>, <op>, <meas>, <val>, are filtered out and not displayed. The <type>, <op>, <meas>, and <val> arguments are separated with one comma.

- <type>

  The type argument can be specified with one of the following arguments:

  curs

  Specifies that only devices with current statistics above the threshold are displayed.

  avgs

  Specifies that only devices with average statistics above the threshold are displayed.

  maxs

  Specifies that only devices with maximum values above the threshold are displayed.

- <op>

  The operation argument can be specified with one of the following arguments:

r

Specifies that read-only statistics are displayed.

w

Specifies that write-only statistics are displayed.

t|rw

Specifies that statistics for read and write totals are displayed.

- <meas>

  The meas argument can be specified with one of the following arguments:

  iops

  Specifies that I/O operations per second are displayed. If this argument is used, the minimum threshold value must be specified using the <val> argument.

  bw

  Specifies that statistics for bandwidth in bytes per second are displayed. If this argument is used, the minimum threshold value must be specified using the <val> argument.

  svct

  Specifies that statistics for service time in milliseconds are displayed.

  size

  Specifies that statistics for I/O operations in bytes are displayed.

- <val> [k|K]|[m|M]|[g|G]

  Specifies the minimum threshold using any integer. The integer can optionally be followed with k or K to indicate a multiple of 1000, m or M to indicate a multiple of 1,000,000, or g or G to indicate a multiple of 1,000,000,000.

An example of this option is -filt curs,r,iops,10k. This indicates that statistics for a device that has more than 10,000 current read-only I/O operations are displayed.

-ni

Specifies that statistics for only non-idle devices are displayed. This option is shorthand for the option -filt curs,t,iops,0.

-subset

Shows subset statistics for Asynchronous Remote Copy on a per group basis.

Specifiers

<VV_name|pattern>...

Only statistics are displayed for the specified virtual volume or pattern. Multiple volumes or patterns can be repeated (for example <VV_name> <VV_name>...). If not specified, all virtual volumes which are configured for Remote Copy are listed.

Authority

Any role in the system

Usage

This command requires the HPE 3PAR Remote Copy Software license. Contact your local service provider for more information.

Examples

The following example shows statistics for all synchronous Remote Copy volumes:

```
cli% statrcvv -sync
                    09:48:45 03/30/10 I/O per sec    KBytes per sec    Svt ms Rmt Svt ms IOSz KB
  VVname RCGroup Target Mode  Port Type Cur Avg Max   Cur   Avg   Max  Cur  Avg  Cur   Avg     Cur   Avg
testvv.0 multi.1   amp2 Sync 0:2:1   FC  32  45  50  4181  5839  6534 11.5 11.6  1.3   1.3   131.1 131.1
testvv.0 multi.1   amp2 Sync 1:2:1   FC  49  35  49  6456  4555  6456  9.3 10.4  1.3   1.5   131.1 131.1
testvv.1 multi.1   amp2 Sync 0:2:1   FC  37  38  39  4796  4952  5093 10.0 10.3  1.2   1.1   131.1 131.1
testvv.1 multi.1   amp2 Sync 1:2:1   FC  39  38  39  5164  5044  5164  7.7  7.2  1.0   1.0   131.1 131.1
testvv.2 multi.1   amp2 Sync 0:2:1   FC  45  14  45  5841  1895  5841  9.7  9.9  1.1   1.1   131.1 131.1
testvv.2 multi.1   amp2 Sync 1:2:1   FC  33  70  89  4365  9202 11721  6.9  6.6  1.0   1.4   131.1 131.1
testvv.3 multi.1   amp2 Sync 0:2:1   FC  34  37  42  4427  4907  5449  9.9 10.2  1.1   1.1   131.1 131.1
testvv.3 multi.1   amp2 Sync 1:2:1   FC  44  38  44  5779  5014  5779  7.4  7.2  1.0   0.9   131.1 131.1
testvv.4 multi.1   amp2 Sync 0:2:1   FC  45  48  51  5902  6328  6674 10.9 11.4  1.3   1.3   131.1 131.1
testvv.4 multi.1   amp2 Sync 1:2:1   FC  32  30  32  4181  3913  4181 11.1 11.3  1.5   1.5   131.1 131.1
testvv.8   sync.2   amp3 Sync 0:3:1   IP  33  33  35  4304  4326  4531 16.7 17.4  7.5   7.8   131.1 131.1
testvv.8   sync.2   amp3 Sync 1:3:1   IP  25  23  25  3259  2950  3259 15.5 15.7  7.5   7.7   131.1 131.1
testvv.9   sync.2   amp3 Sync 0:3:1   IP  28  28  28  3689  3669  3691 14.0 14.5  7.6   7.6   131.1 131.1
testvv.9   sync.2   amp3 Sync 1:3:1   IP  29  28  29  3750  3669  3750 12.8 12.8  7.5   7.6   131.1 131.1
-------------------------------------------------------------------------------------------------------------
      14                                 504 506     66093 66262       10.6 10.7  2.6   2.7   131.1 131.1
Press the enter key to stop...
```

Use the following to get a description of each column in the above examples:

```
clihelp -col statrcvv
```

statvlun

Syntax

```
statvlun [options <arg>]
```

Description

The `statvlun` command displays statistics for virtual volumes and LUN host attachments.

Options

`-domain <domain_name>...|<pattern>...`

Shows only VLUNs whose virtual volumes are in domains with names that match one or more of the specified domain names or patterns. Multiple domain names or patterns can be repeated using a comma-separated list.

`-host <hostname>...|<pattern>...`

Shows only VLUNs exported to the specified host(s) or pattern(s).

Multiple host names or patterns can be repeated using a comma-separated list.

`-v <VV_name>...|<pattern>...`

Requests that only logical disks mapped to virtual volumes that match any of the specified names or patterns be displayed. Multiple volume names or patterns can be repeated using a comma-separated list.

`-l <LUN>...|<pattern>...`

Specifies that VLUNs with LUNs matching the specified LUN(s) or pattern(s) are displayed. Multiple LUNs or patterns can be repeated using a comma-separated list.

`-nodes <node_list>`

Specifies that the display is limited to specified nodes and physical disks connected to those nodes. The node list is specified as a series of integers separated by commas (e.g. 1,2,3). The list can also consist of a single integer. If the node list is not specified, all disks on all nodes are displayed.

`-slots <slot_list>`

Specifies that the display is limited to specified PCI slots and physical disks connected to those PCI slots. The slot list is specified as a series of integers separated by commas (e.g. 1,2,3). The list can also consist of a single integer. If the slot list is not specified, all disks on all slots are displayed.

-ports <port_list>

Specifies that the display is limited to specified ports and physical disks connected to those ports. The port list is specified as a series of integers separated by commas (e.g. 1,2,3). The list can also consist of a single integer. If the port list is not specified, all disks on all ports are displayed.

-lw

Lists the host's WWN or iSCSI name. This is especially useful when multiple WWNs or iSCSI names belonging to the same host are visible on the same port.

-domainsum

Specifies that sums for VLUNs are grouped by domain in the display. All VLUNs to unnamed hosts are added and displayed as a single set of data with a "-" host name.

-vvsum

Specifies that sums for VLUNs of the same virtual volume are displayed.

-hostsum

Specifies that sums for VLUNs are grouped by host in the display. All VLUNs to unnamed hosts are added and displayed as a single set of data with a nameless host.

-rw

Specifies that the display includes separate read and write data. If not specified, the total is displayed.

-d <secs>

Specifies the interval in seconds that statistics are sampled from using an integer from 1 through 2147483. If no count is specified, the command defaults to 2 seconds.

-iter <number>

Specifies that the statistics are to stop after the indicated number of iterations using an integer from 1 through 2147483647.

-begin

Specifies that I/O averages are computed from the system start time. If not specified, the average is computed since the first iteration of the command.

-idlep

Specifies the percent of idle columns in the output.

-sortcol <col>[,<dir>][:<col>[,<dir>]...]

Sorts command output based on column number (`<col>`). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (`<dir>`) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

-filt <fspec>

Specifies that statistics that fall below the threshold as specified by the `<fspec>` arguments, where `<fspec>` is `<type>`, `<op>`, `<meas>`, `<val>`, are filtered out and not displayed. The `<type>`, `<op>`, `<meas>`, and `<val>` arguments are separated with one comma.

- `<type>`

The type argument can be specified with one of the following arguments:

curs

Specifies that only devices with current statistics above the threshold are displayed.

avgs

Specifies that only devices with average statistics above the threshold are displayed.

maxs

Specifies that only devices with maximum values above the threshold are displayed.

- `<op>`

The operation argument can be specified with one of the following arguments:

r

Specifies that read-only statistics are displayed.

w

Specifies that write-only statistics are displayed.

t|rw

Specifies that statistics for read and write totals are displayed.

- `<meas>`

The meas argument can be specified with one of the following arguments:

iops

Specifies that I/O operations per second are displayed. If this argument is used, the minimum threshold value must be specified using the `<val>` argument.

bw

Specifies that statistics for bandwidth in bytes per second are displayed. If this argument is used, the minimum threshold value must be specified using the `<val>` argument.

svct

Specifies that statistics for service time in milliseconds are displayed.

size

Specifies that statistics for I/O operations in bytes are displayed.

- `<val> [k|K]|[m|M]|[g|G]`

Specifies the minimum threshold using any integer. The integer can optionally be followed with k or K to indicate a multiple of 1000, m or M to indicate a multiple of 1,000,000, or g or G to indicate a multiple of 1,000,000,000.

An example of this option is `-filt curs,r,iops,10k`. This indicates that statistics for a device that has more than 10,000 current read-only I/O operations are displayed.

-ni

Specifies that statistics for only non-idle devices are displayed. This option is shorthand for the option `-filt curs,t,iops,0`.

The VV filtering option include:

`-p <pattern>`

Pattern for matching VVs to show. (see below for description of `<pattern>`) If the `-p` option is specified multiple times, each instance of `<pattern>` adds additional candidate VVs that match that pattern.

A <pattern> is one or more of the following:

-vmname {<vmname>|<pattern>}[,{<vmname>|<pattern>}...]

Shows only VVs whose vmname matches one or more of the vvname, patterns.

-vmid <vmid>[,<vmid>...]

Shows only VVs whose vmid matches one or more of the vmids.

-vmhost {<vmhost>|<pattern>}[,{<vmhost>|<pattern>}...]

Shows only VVs whose vmhost matches one or more of the vmhost, patterns.

-vvolstate {bound|unbound}

Shows only VVs whose vvolstate matches the specified state — bound or unbound.

-vvolsc {<vvset>|<pattern>}[,{<vvset>|<pattern>}...]

Shows only VVs whose storage container (vvset) name matches one or more of the vvset, patterns.

Specifiers

None.

Authority

Any role in the system

Usage

- For this command KB = 1000 bytes.
- For all the <name_or_pattern> fields, the patterns are glob-style (shell-style) patterns. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.)
- A Domain column may be included by using the **setclienv** command to set listdom to 1. The listdom option is also set to 1 if the CLI was started with the `-listdom` option, or if the CLI was started with the TPDLISTDOM environment variable set. Run "`cli -h`" and "`setclienv -h`" for details of the environment variables.

Examples

The following example displays one iteration of statistics for virtual volumes and LUN host attachments:

```
cli% statvlun -iter 1

             07:57:42 06/29/2015 r/w I/O per second       KBytes per sec    Svt ms   IOSz KB
Lun        VVname     Host  Port      Cur  Avg  Max    Cur    Avg    Max  Cur  Avg  Cur  Avg Qlen
  0 TPVV8280_0002 shrapnel 0:1:1   t    0  680 2809      0  44541 184102 0.00 0.95  0.0 65.5    0
  1 TPVV8280_0003 shrapnel 0:1:1   t    0  680 2831      0  44541 185509 0.00 1.03  0.0 65.5    0
  2   VV8280_0006 shrapnel 0:1:1   t 3356  342 3356 219946  22385 219946 0.86 0.81 65.5 65.5    4
  3   VV8280_0007 shrapnel 0:1:1   t    0    0    0      0      0      0 0.00 0.00  0.0  0.0    0
  0 TPVV8280_0002 shrapnel 1:1:1   t    0  680 2809      0  44541 184066 0.00 0.95  0.0 65.5    0
  1 TPVV8280_0003 shrapnel 1:1:1   t    0  680 2831      0  44541 185541 0.00 1.08  0.0 65.5    0
  2   VV8280_0006 shrapnel 1:1:1   t 3356  342 3356 219969  22385 219969 0.71 0.66 65.5 65.5    4
  3   VV8280_0007 shrapnel 1:1:1   t    0    0    0      0      0      0 0.00 0.00  0.0  0.0    0
------------------------------------------------------------------------------------------------
                8                 t 6713 3402      439914 222932        0.78 0.95 65.5 65.5    8
```

Use the following to get a description of each column in the above examples:

```
clihelp -col statvlun
```

statvv

Syntax

```
statvv [options <arg>] [<VV_name_or_pattern>...]
```

Description

The `statvv` command displays statistics for virtual volumes in a timed loop.

Options

`-domain <domainname>...|<pattern>...`

Shows only the virtual volumes that are in domains with names that match the specified domain name(s) or pattern(s).

`-rw`

Specifies that the display includes separate read and write data. If not specified, the total is displayed.

`-d <secs>`

Specifies the interval in seconds that statistics are sampled from using an integer from 1 through 2147483. If no count is specified, the command defaults to 2 seconds.

`-iter <number>`

Specifies that the statistics are to stop after the indicated number of iterations using an integer from 1 through 2147483647.

`-sortcol <col>[,<dir>][:<col>[,<dir>]...]`

Sorts command output based on column number (<col>). Columns are numbered from left to right, beginning with 0. At least one column must be specified. In addition, the direction of sorting (<dir>) can be specified as follows:

- inc—Sort in increasing order (default).
- dec—Sort in decreasing order.

Multiple columns can be specified and separated by a colon (:). Rows with the same information in them as earlier columns will be sorted by values in later columns.

`-filt <fspec>`

Specifies that statistics that fall below the threshold as specified by the <fspec> arguments, where <fspec> is <type>, <op>, <meas>, <val>, are filtered out and not displayed. The <type>, <op>, <meas>, and <val> arguments are separated with one comma.

- `<type>`

  The type argument can be specified with one of the following arguments:

  `curs`

  Specifies that only devices with current statistics above the threshold are displayed.

  `avgs`

  Specifies that only devices with average statistics above the threshold are displayed.

  `maxs`

  Specifies that only devices with maximum values above the threshold are displayed.

- `<op>`

  The operation argument can be specified with one of the following arguments:

  `r`

  Specifies that read-only statistics are displayed.

w

Specifies that write-only statistics are displayed.

t|rw

Specifies that statistics for read and write totals are displayed.

- `<meas>`

  The meas argument can be specified with one of the following arguments:

  iops

  Specifies that I/O operations per second are displayed. If this argument is used, the minimum threshold value must be specified using the <val> argument.

  bw

  Specifies that statistics for bandwidth in bytes per second are displayed. If this argument is used, the minimum threshold value must be specified using the <val> argument.

  svct

  Specifies that statistics for service time in milliseconds are displayed.

  size

  Specifies that statistics for I/O operations in bytes are displayed.

- `<val> [k|K]|[m|M]|[g|G]`

  Specifies the minimum threshold using any integer. The integer can optionally be followed with k or K to indicate a multiple of 1000, m or M to indicate a multiple of 1,000,000, or g or G to indicate a multiple of 1,000,000,000.

An example of this option is `-filt curs,r,iops,10k`. This indicates that statistics for a device that has more than 10,000 current read-only I/O operations are displayed.

-ni

Specifies that statistics for only non-idle devices are displayed. This option is shorthand for the option -filt curs,t,iops, 0.

The VV filtering option include:

`-p <pattern>`

Pattern for matching VVs to show. (see below for description of <pattern>) If the `-p` option is specified multiple times, each instance of <pattern> adds additional candidate VVs that match that pattern.

A <pattern> is one or more of the following:

-vmname {<vmname>|<pattern>}[,{<vmname>|<pattern>}...]

Shows only VVs whose vmname matches one or more of the vvname, patterns.

-vmid <vmid>[,<vmid>...]

Shows only VVs whose vmid matches one or more of the vmids.

-vmhost {<vmhost>|<pattern>}[,{<vmhost>|<pattern>}...]

Shows only VVs whose vmhost matches one or more of the vmhost, patterns.

-vvolstate {bound|unbound}

Shows only VVs whose vvolstate matches the specified state — bound or unbound.

-vvolsc {<vvset>|<pattern>}[,{<vvset>|<pattern>}...]

Shows only VVs whose storage container (vvset) name matches one or more of the vvset, patterns.

Specifiers

<VV_name|pattern>...

Only statistics are displayed for the specified virtual volume or pattern. Multiple volumes or patterns can be repeated (for example <VV_name> <VV_name>...). If not specified, all virtual volumes are listed.

Authority

Any role in the system

Usage

- If a <VV_name> or <pattern> are specified, then virtual volumes with names matching any of the patterns are listed. Otherwise all virtual volumes are listed. These patterns are glob-style patterns. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.)
- Virtual volumes may be accessed externally by hosts and internally by the prefetcher. Virtual volume data measured by this command include accesses by the prefetcher.
- In addition to external accesses by hosts, virtual volumes can be read internally by the system read-ahead prefetcher. The `histvv` data includes read-ahead accesses from the prefetcher that can cause the read data to appear more than seen by the hosts. Use the `histvlun -vvsum` command to see data for only accesses from the host (see **histvlun**).
- Volumes for VMware virtual machines (vVols), which are hidden by default, may be included by using the **setclienv** command to set matchbulkobjs to "1".
- Additionally, vVols can displayed if the CLI was started with the -matchbulkobjs option, or if the CLI was started with the TPDMATCHBULKOBJS environment variable set to "1". Run "`cli -h`" and "`setclienv -h`" for details of the environment variables. In addition, using the pattern matching options -vmname, -vmid, -vmhost, -vvolstate, or -vvolsc will display vVols.
- For this command KB = 1000 bytes.

Examples

The following example displays I/O statistics for all virtual volumes:

```
cli% statvv -iter 1
15:31:21 05/31/07 r/w I/O per second KBytes per sec Svt ms IOSz KB
                 VVname      Cur Avg Max Cur Avg Max  Cur   Avg   Cur Avg Qlen
                 admin t     0   0   0    0    0 0   0.0   0.0   0.0 0.0 0
                 test2 t     0   0   0    0    0 0   0.0   0.0   0.0 0.0 0
    ------------------------------------------------------------------------
                 total t     0   0        0    0      0.0 0.0 0.0 0.0    0
```

Use the following to get a description of each column in the above examples:

```
clihelp -col statvv
```

statwsapi

Syntax

```
statwsapi [options]
```

Description

Displays statistics for Web Services API (WSAPI).

Options

-d *<sec>*

Specifies the interval, in seconds, that statistics are sampled using an integer from 1 through 2147483. If no interval is specified, the option defaults to an interval of 5 seconds.

-filter *<IP_pattern>*[:*<object_pattern>*[:*<action_pattern>*]]

Used with `-ip` or `-op` to select IP address, object name, and action with glob patterns. If not specified the patterns are "*".

-iter *<number>*

Specifies that statistics are displayed a specified number of times (from 1 through 2147483647).

-ip

Displays detailed statistics sorted by IP address.

-multi

Uses with `-ip` or `-op` to display operation statistics in multiple tables, based on operation object name or source IP address instead of one table for all objects from all source IP addresses.

-op

Displays detailed statistics presented by object/action.

-reset

Resets all statistics.

Specifiers

None.

Authority

Any role in the system

Usage

None.

Examples

Displaying the current Web Services API server statistics:

```
cli% statwsapi
```

Displaying detailed statistics sorted by IP address for all IP addresses, objects, and actions:

```
cli% statwsapi -ip
```

Displaying detailed statistics sorted by Object/Action for all IP addresses, objects starting with **v**, and all actions:

```
cli% statwsapi -op -filter *:v*
```

Displaying detailed statistics presented by object/action for IP addresses starting with 192 and actions starting with **q**

```
cli% statwsapi -ip -filter 192*:*:q*
```

Displaying a description of each column in the output from the above examples:

```
clihelp -col statwsapi
```

Stop Commands

stopcim

Syntax

```
stopcim [options]
```

Description

The `stopcim` command stops the CIM server from servicing CIM requests.

Options

`-f`

Specifies that the operation is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

`-x`

Specifies that the operation terminates the server immediately without graceful shutdown notice.

Specifiers

None.

Authority

- Super, Service
- Any role granted the `cim_stop` right

Usage

- Access to all domains is required to run this command.
- By default, the CIM server is not started until the **`startcim`** command is issued.

Examples

The following example stops the CIM server:

```
cli% stopcim
Are you sure you want to stop CIM server?
select q=quit y=yes n=no: y
CIM server stopped successfully.
```

The following example stops the CIM server immediately without graceful shutdown notice and confirmation:

```
cli% stopcim -f -x
CIM server stopped successfully.
```

stopfs

Syntax

```
stopfs [options] [<nodeid>,]...
```

```
stopfs -remove [options] <n:s:p> <n:s:p> [{<n:s:p> <n:s:p>} ...]
```

Description

The `stopfs` command stops or removes File Persona.

Options

`-remove`

Disables and removes File Persona from the specified nodes.

`-f`

Forces the command. The command completes the process without prompting for confirmation.

Specifiers

`<n:s:p>`

Specifies the node and port on which File Persona will be removed.

This is valid only for removal and used only with the -remove option.

Nodes must be in the matched pairs specified at install time.

`<nodeid>`

Specifies the node on which File Persona is to be stopped. If no node is specified, all nodes with File Persona configured will be used.

Authority

Super, Service

Usage

- Access to all domains is required to run this command.
- The `stopfs` command will stop File Persona on a system node or set of nodes. If the extra option, -remove, is present File Persona will be disabled and removed from the specified node pair.
- If the -enable option is provided File Persona is started on all or specified nodes.
- File Persona must have previously installed onto all or any specified nodes.

Examples

The following example stops File Persona on nodes 0 and 1:

```
cli% stopfs -f 0,1
```

The following example stops and removes File Persona from node pair 0 and 1:

```
cli% stopfs -f -remove 0:2:1 1:2:1
```

stopfsarchive

Syntax

```
stopfsarchive scan [-pause] [-f] [-fpg <fpgname>] -fstore <fstorename> <vfs>
<jobid>
```

Description

The `stopfsarchive` command is used to stop/pause the retention validation scan on the retained/WORM'ed file.

Options

`-fpg <fpgname>`

Specifies the FPG that contains the VFS. This is required if VFS exists under multiple FPGs.

-fstore <fstorename>

Name of the file store under VFS.

-f

Suppresses confirmation from the user before stopping the scan job.

-pause

Pauses the retention validation job.

Specifiers

<vfs>

Specifies the name of the VFS.

<jobid>

Job Id of the scan to pause or stop.

Subcommands

scan

Stops or pauses the already running data validation scan.

Authority

- Super, Service
- Any role granted the `fsarchive_stop` right.

Usage

Access to all domains is required to run this command.

Examples

The following example stops retention validation scan with the provided job Id running on the file store:

```
cli% stopfsarchive scan -fpg fpg1 -fstore fstore1 vfs1 10
```

The following example pauses the retention validation scan process with the provided job Id running on the file store.

```
cli% stopfsarchive scan -pause -fpg fpg1 -fstore fstore1 vfs1 10
```

stopfsav

Syntax

```
stopfsav svc
```

```
stopfsav scan [-pause] [-fpg <fpgname>] [-fstore <fstore>] <vfs_name>
<scan_id>
```

Description

The `stopfsav` command is used to both stop the antivirus service as well as to stop or pause a running scan.

Options

-fpg <fpgname>

The name of the File Provisioning Group in which the VFS was created.

-fstore <fstore>

Specifies the File Store name, using up to 31 characters.

-pause

Pauses, rather than stop, the specified scan.

Specifiers

<vfs_name>

The Virtual File Server (VFS) name, scan to Stop or Pause.

<scan_id>

The scan task identifier, as provided by "`startfsav scan`" command, and displayed by "`showfsav scan <vfs>`" command.

Subcommands

svc

Specifies antivirus service.

scan

Stops and pauses currently running antivirus scan.

Authority

- Super, Edit
- Any role granted the `fsav_stop` right

Usage

- Access to all domains is required to run this command.
- The "`startfsav scan -resume <scan_id> <vfs>`" command can be used to resume the paused scan (see **startfsav**).
- When the antivirus service is stopped manually by using "`stopfsav svc`" command, adding and removing of the VSE cannot be performed using "`setfsav vse`" command (see **setfsav**).
- The antivirus service needs to be started using "`startfsav svc`" command (see **startfsav**).

Examples

Stopping an antivirus scan with the scan_id AvScanTask_12:

```
cli% stopfsav scan test_vfs AvScanTask_12
```

Pausing an antivirus scan with the scan_id AvScanTask_20:

```
cli% stopfsav scan -pause test_vfs AvScanTask_20
```

stopfsnapclean

Syntax

```
stopfsnapclean [-pause] <fpgname>
```

Description

The `stopfsnapclean` command stops or pauses an on-demand reclamation task on a file provisioning group.

Options

-pause

Specifies to pause a reclamation task.

Specifiers

<fpgname>

Specifies the name of the file provisioning group.

Authority

- Super, Service
- Any role granted the `fsnapclean_stop` right

Usage

- Access to all domains is required to run this command.
- There can be only one reclamation task running on a file provisioning group. If we pause reclamation task, it will still be counted.

Examples

The following example stops a running reclamation task on samplepool file provisioning group:

```
cli% stopfsnapclean samplepool
```

If the task is not running, the following output is displayed:

```
cli% stopfsnapclean samplepool
    No reclamation task running on FPG samplepool
```

The following example pauses a running reclamation task on samplepool file provisioning group

```
cli% stopfsnapclean -pause samplepool
```

stopfsndmp

Syntax

`stopfsndmp`

Description

The `stopfsndmp` command is used to stop both NDMP service and ISCSI service.

Options

None.

Specifiers

None.

Authority

- Super, Edit
- Any role granted the `fsndmp_stop` right

Usage

- Access to all domains is required to run this command.
- NDMP and ISCSI services can be started by the **startfsndmp** command.
- The status of the NDMP and ISCSI service can be checked by "`showfsndmp -conf`" command (see **showfsndmp**).

Examples

Stopping NDMP and ISCSI services:

```
cli% stopfsndmp
```

stoprcopy

Syntax

```
stoprcopy [options]
```

Description

Stops the Remote Copy service and optionally stops any started Remote Copy volume groups.

Options

-f

Suppresses user interaction with the command. Does not ask for confirmation for either the `-clear` or `-stopgroups` options.

-stopgroups

Specifies that any started Remote Copy volume groups are stopped.

-clear

Specifies that configuration entries affiliated with the stopped mode are deleted.

-keepalua

Keeps the ALUA state of the local volume from changing. Used only with the `-clear` option. Required for use during Peer Motion migrations of Remote Copy Groups. See, *Volume migration in a Remote Copy Primary group* in the *HPE 3PAR Peer Motion User Guide* for details.

Specifiers

None.

Authority

- Super, Edit
- Any role granted the `rcopy_stop` right

Usage

- Access to all domains is required to run this command.
- This command requires the HPE 3PAR Remote Copy license. Contact your local service provider for more information.
- If the `-stopgroups` option is not used, all groups must already be stopped.
- Unless the `-stopgroups` option is used, the command will fail if there are any Remote Copy groups that are started. If the `-clear` option has been used, the configuration is completely erased. Consequently, Remote Copy

operations cannot be restarted using only the **startrcopy** command. The configuration must be rebuilt. Therefore, the -clear and -stopgroups option requires confirmation with the -f option, the TPDFORCE environment variable, or by interactively typing "y".

Examples

The following example disables the Remote Copy functionality of all primary Remote Copy volume groups:

```
cli% stoprcopy -stopgroups
```

stoprcopygroup

Syntax

```
stoprcopygroup [option <arg>] [<group_name> | <pattern>...]
```

Description

Stops the Remote Copy functionality for the specified Remote Copy volume groups.

Options

-nosnap

In synchronous mode, this option turns off snapshot creation.

This is useful if you are running `removercopygroup` to remove Remote Copy.

In asynchronous periodic mode, this option deletes any current synchronization snapshots. Using this option requires a full resynchronization if you intend to restart the group later.

No coordinated snapshot is taken upon group stop if you use this option in asynchronous streaming mode. If you restart the group, the system uses the most recent coordinated snapshot for resync. If there is no coordinated snapshot available, a full sync occurs.

-t <targetname>

Specifies the target on which to stop the group.

-f

Suppresses user interaction with the command. Without this option, the command requires confirmation before proceeding with its operation.

-pat

Specifies treating patterns as glob-style, and all Remote Copy groups matching the specified pattern are stopped. Use the `-pat` option to specify a list of patterns. Requires confirmation to proceed with the command unless you specify the `-f` option. Stop a maximum of 512 Remote Copy groups using this option. Required if you use the *<pattern>* specifier.

-remote

This option can be used to stop the group from the secondary array. It is ignored when issued from the primary array. Behaves the same way as if the command is run on the primary array without this option. When run from the secondary array, the -t <target> option is ignored. On SLD and 3DC-PP configurations, this option is not supported.

Specifiers

<group_name>

The name of the Remote Copy volume group to stop.

<pattern>

Specifies a glob-style pattern. (For more information, enter the `clihelp sub,globpat` command, or see **Glob-style patterns**.) This specifier can be repeated to stop multiple groups. If this specifier is not used, the *<group_name>* specifier must be used.

Authority

- Super, Edit
- Any role granted the `rcopygroup_stop` right

Usage

- In sync mode, this command creates snapshots that are used for synchronizing the primary and secondary groups if the **`startrcopygroup`** command is run later.
- If a group target has the mirror_config policy set and the group is a primary group, then this command will be mirrored to that target and the corresponding secondary group will be stopped. If the mirror_config policy is set and the group is a secondary, then this command will fail if the user does not specify the -remote option.
- If the mirror_config policy is not set and this command is issued on a secondary group, then the corresponding primary group will also be stopped as a result of this command.

Examples

The following example stops Remote Copy for Group1:

```
cli% stoprcopygroup Group1
```

The following example stops Remote Copy Groups that start with the name testgroup:

```
cli% stoprcopygroup -pat testgroup*
```

stopvasa

Syntax

```
stopvasa [options]
```

Description

The `stopvasa` command stops the VASA Provider server from servicing HTTPS requests.

Options

`-f`

Specifies that the operation is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

Specifiers

None.

Authority

- Super, Service
- Any role granted the `vasa_stop` right

Usage

- Access to all domains is required to run this command.
- Use the **startvasa** command to start the VASA Provider server.

Examples

The following example shows how to stop the VASA Provider server:

```
cli% stopvasa
Are you sure you want to stop the VASA Provider server?
select q=quit y=yes n=no: y
VASA Provider server stopped successfully.
```

The following example shows how to stop the VASA Provider server immediately without a confirmation notice:

```
cli% stopvasa -f
VASA Provider server stopped successfully.
```

stopwsapi

Syntax

```
stopwsapi [options]
```

Description

Stops the Web Services API server from servicing HTTP and HTTPS requests, and stops the 3PAR GUI.

Options

-f

Specifies that the operation is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

Specifiers

None.

Authority

- Super, Service
- Any role granted the `wsapi_stop` right

Usage

- To run this command, you must have access to all domains.
- The Web Services API server defaults to `started`.

Examples

The following example shows how to stop the Web Services API server and 3PAR GUI:

```
cli% stopwsapi
Are you sure you want to stop the Web Services API server?

select q=quit y=yes n=no: y
Web Services API server and 3PAR GUI
 stopped successfully.
```

The following example shows how to stop the Web Services API server and 3PAR GUI immediately without a confirmation notice:

```
cli% stopwsapi -f
Web Services API server and 3PAR GUI
 stopped successfully.
```

Sync Command

syncrcopy

Syntax

```
syncrcopy [options] [<group_name> | <pattern>...]
```

Description

The `syncrcopy` command manually synchronizes remote copy volume groups.

Options

`-w`

Waits for synchronization to complete before returning to a command prompt.

The -w option is not supported if the group is a secondary group. If `syncrcopy` command is run on secondary group with -w specified, the command will fail.

`-n`

Do not save resynchronization snapshot. This option is only relevant for asynchronous periodic mode groups.

`-ovrd`

Forces full synchronization without prompting for confirmation, even if volumes are already synchronized. This option is only relevant for synchronous mode volume groups and asynchronous streaming mode volume groups and can be used to re-synchronize volumes that have become inconsistent.

`-t <targetname>`

Indicates that only the group on the specified target is synchronized. If this option is not used, by default, the `syncrcopy` command will affect all of a group's targets.

`-pat`

Specifies that the patterns specified are treated as glob-style patterns and all remote copy groups matching the specified pattern will be synced. The -pat option can specify a list of patterns. This option must be used if the <pattern> specifier is used.

Specifiers

`<group_name>`

Specifies the name of the remote copy volume group to be synchronized.

This name can be obtained using the **`showrcopy`** command.

`<pattern>`

Specifies a glob-style pattern. This specifier can be repeated to sync multiple groups. If the <pattern> specifier is not used, the <group_name> specifier must be used.

Authority

- Super, Edit
- Any role granted the `rcopy_sync` right

Usage

- This command requires the HPE 3PAR Remote Copy Software license. Contact your local service provider for more information.
- Mode (synchronous, asynchronous periodic, or asynchronous streaming) is set using the `creatercopygroup` command. For information about modes and creating Remote Copy volume groups, see **`creatercopygroup`**.
- For information about setting targets, see **`creatercopytarget`**.
- Using the -n option requires a full synchronization at the next sync.

Examples

The following example specifies that remote copy volume group Group1 should be synchronized with its corresponding secondary volume group:

```
cli% syncrcopy Group1
Synchronization request issued for group Group1
```

The following example specifies that all remote copy volume groups that start with the name testgroup should be synchronized with their corresponding secondary volume group:

```
cli% syncrcopy  -pat testgroup*
```

Tune Commands

tunepd

Syntax

The syntax for the `tunepd` command can be one of the following:

```
tunepd [options <arg>] maxsvct <msecs>|highest
tunepd [options <arg>] avgsvct <msecs>|highest
```

Description

The `tunepd` command identifies physical disks with high service times and optionally executes load balancing.

Options

`-nodes <node_list>`

Specifies that the display is limited to specified nodes and physical disks connected to those nodes. The node list is specified as a series of integers separated by commas (e.g. 1,2,3). The list can also consist of a single integer. If the node list is not specified, all disks on all nodes are displayed.

`-slots <slot_list>`

Specifies that the display is limited to specified PCI slots and physical disks connected to those PCI slots. The slot list is specified as a series of integers separated by commas (e.g. 1,2,3). The list can also consist of a single integer. If the slot list is not specified, all disks on all slots are displayed.

`-ports <port_list>`

Specifies that the display is limited to specified ports and physical disks connected to those ports. The port list is specified as a series of integers separated by commas (e.g. 1,2,3). The list can also consist of a single integer. If the port list is not specified, all disks on all ports are displayed.

`-vv <VV_name>`

Specifies that the physical disks used by the indicated virtual volume name are included for statistic sampling.

`-d <seconds>`

Specifies the interval, in seconds, that statistics are sampled using an integer from 1 through 2147483. If no interval is specified, the option defaults to 30 seconds.

`-iter <number>`

Specifies that I/O statistics are sampled a specified number of times as indicated by the number argument using an integer greater than 0. If 0 is specified, I/O statistics are looped indefinitely. If this option is not specified, the command defaults to 1 iteration.

`-freq <minutes>`

Specifies the interval, in minutes, that the command enters standby mode between iterations using an integer greater than 0. If this option is not specified, the number of iterations is looped indefinitely.

`-vvlayout`

Specifies that the layout of the virtual volume is displayed. If this option is not specified, the layout of the virtual volume is not displayed.

-portstat

Specifies that statistics for all disk ports in the system are displayed. If this option is not specified, statistics for ports are not displayed.

-pdstat

Specifies that statistics for all physical disk, rather than only those with high service times, are displayed. If this option is not specified, statistics for all disks are not displayed.

-chstat

Specifies that chunklet statistics are displayed. If not specified, chunklet statistics are not displayed. If this option is used with the `-movech` option, either on or force must be specified.

-maxpd <number>

Specifies that only the indicated number of physical disks with high service times are displayed. If this option is not specified, 10 physical disks are displayed.

-movech auto|manual

Specifies that if any disks with unbalanced loads are detected that chunklets are moved from those disks for load balancing.

- `auto`

  Specifies that the system chooses source and destination chunklets. If not specified, you are prompted for selecting the source and destination chunklets.

- `manual`

  Specifies that the source and destination chunklets are manually entered.

Specifiers

maxsvct <msecs>|highest

Specifies that either the maximum service time threshold (<msecs>) that is used to discover over-utilized physical disks, or the physical disks that have the highest maximum service times (highest). If a threshold is specified, then any disk whose maximum service time exceeds the specified threshold is considered a candidate for load balancing.

avgsvct <msecs>|highest

Specifies that either the average service time threshold (<msecs>) that is used to discover over-utilized physical disks, or the physical disks that have the highest average service time (highest). If a threshold is specified, any disk whose average service time exceeds the specified threshold is considered a candidate for load balancing.

Authority

- Super, Edit
- Any role granted the `pd_tune` right

Usage

- Access to all domains is required to run this command.
- Use of this command requires a System Tuner license. Contact your local HPE representative for information.
- The system does not allow multiple `tunepd` processes with the -chstat on or -chstat force options to run concurrently.
- If one or more chunklet statistics collection processes are running, the `tunepd` command exits to prevent possible `tunepd -chstat` on processes running simultaneously.
- Use the "`-chstat force`" option only when no other `tunepd` processes are running.

- When the "`-movech auto`" or "`-movech manual`" option is specified, the system only identifies (auto mode) or recommends (manual mode) source chunklets for which there are destination chunklets retaining the availability of the source chunklets' logical disks.
- Separate instances of the `tunepd` command can be issued to identify load balancing candidates and relocate chunklets, or execute both tasks with one instance of the `tunepd` command.
- Previous options -n <node_list>, -s <slot_list> and -p <port_list> have been deprecated in 2.2.3 and replaced with -nodes <node_list>, -slots <slot_list> and -ports <port_list>.
- If the -nodes, -slots, -ports or -vv options are not specified, all physical disks are included.

Examples

In the following example, physical disks with average service times exceeding 50 milliseconds are identified and their chunklets automatically relocated to re-balance the physical disks' load.

```
cli% tunepd -vvlayout -chstat -movech auto avgsvct 50
```

tunesys

Change the layout of a storage system.

Syntax

```
tunesys [options]
tunesys -convert {-tpvv | -reduce} [options] <VV | VVset>[:<dst_cpg>]...
tunesys -convert {-tpvv | -reduce} [options] -cpglist
<src_cpg>[:<dst_cpg>]...
tunesys -convert {-tpvv | -reduce} -allvvs [options]
```

Description

Analyze and detect poor layout and disk utilization across an entire storage system. The command runs a series of low level operations to re-balance resources on the system.

Options

Inter-node tuning options:

`-cpg <CPG name or patterns>`

Limits the scope of a `tunesys` operation to the named CPG(s). The specified CPGs must all be in the same domain as the user. If this option is specified the intra-node (`tunenodech`) phase is not run. `-chunkpct` and `-tunenodech` cannot be used with this option.

`-nodepct <percentage>`

Controls the detection of utilization imbalances between nodes. If any node has a PD devtype where the average utilization is more than `<percentage>` less than the average for that devtype, then detailed VV level analysis is performed. VVs which are poorly balanced between nodes will have a tune generated to correct the imbalance. `<percentage>` must be between 1 and 100. The default value is 3.

`-spindlepct <percentage>`

Specifies the percentage difference between node pairs that can exist before `tunesys` warns that an imbalance exists. The percentage difference calculated between node pairs must be less than spindlepct. `<percentage>` must be between 1 and 200. 200 is the least restrictive and would allow the `tunesys` to not warn with any difference in the number of PDs, while 1 is the most restrictive. 0 cannot be specified as this would always generate a warning. The default for `<percentage>` is 50 (allow for a 50% difference).

-force

Bypasses top-level inter-node balance checks and force detailed analysis of every VV. This option can be used to complete the re-balance of a relatively well balanced system where only a few volumes are unbalanced.

-slth <threshold>{g|G|t|T}

Slices threshold. Volumes above this size will be tuned in slices. <threshold> must be in multiples of 128GiB. Minimum is 128GiB. Default is 2TiB. Maximum is 16TiB.

-slsz <size>{g|G|t|T}

Slices size. Size of slice to use when volume size is greater than <threshold>. <size> must be in multiples of 128GiB. Minimum is 128GiB. Default is 2TiB. Maximum is 16TiB.

Intra-node tuning options:

-chunkpct <percentage>

Controls the detection of any imbalance in PD chunklet allocation between PDs owned by individual nodes. If a PD has utilization of more than <percentage> less than the average for that device type, then that disk can potentially be tuned. <percentage> must be between 1 and 100. The default value is 10. This option cannot be used with the -cpg option.

-devtype <FC|NL|SSD>

Only tune the specified device type. Applies to the intra-node tune phase only and must be used with the -tunenodech option. Multiple devtypes can be specified. If -devtype is not used, all devtypes will be tuned when -tunenodech is specified.

-fulldiskpct <percentage>

This option is used in the intra-node tuning phase.

If a PD has more than <percentage> of its capacity utilized, chunklet movement is used to reduce its usage to <percentage> before LD tuning is used to complete the rebalance. For example, if a PD is 98% utilized and <percentage> is 90, chunklets will be redistributed to other PDs until the utilization is less than 90%. If <percentage> is less than the devtype average then the calculated average will be used instead. <percentage> must be between 1 and 100. The default value is 90.

-maxchunk <number>

Specifies the maximum number of chunklets which can be moved from any PD in a single operation. <number> must be between 1 and 8. The default value is 8.

-tunenodech

Specifies that only intra-node rebalancing should be performed.

LD tuning options:

-ss

Triggers LD re-tuning for any LD where the stepsize value does not match the parent CPG.

VV conversion options:

-convertdedupversion <source>:<destination>[,<source>:<destination>]...

Convert the dedup version of TDVVs. It requires a comma separated list of source and destination CPG pairs, in the format <source>:<destination>. The destination CPGs cannot have any DDS or dedup VVs on the ld version.

One TDVV is converted at a time. Conversion is started only after ensuring that system has sufficient space.

Conversion fails if the dedup VVs have snapshots or are part of remote copy groups if -inplace option is not specified..

-inplace <source CPG>[,<source CPG>]...

Update dedup version from TDVV2 to TDVV3 within same source CPG. It requires a comma separated list of source CPGs. The source CPG must have DDS volume of dedup version TDVV2. Specify this option only with –convertdedupversion.

-partial

Converts as many VVs as possible. If not specified, the command fails when it is unable to convert all VVs in CPG.

-convert <-tpvv | -reduce> <-allvvs | -cpglist>

This option is used to convert the VV type. It requires a comma-separated list of <VVs | VVs set> or source and destination CPG pairs, in the format <VV | VV set:destination CPG> or <source CPG:destination CPG>.

The default the VV list is in the format <VV | VV set:destination CPG> or <VV | VV set>. It converts 1 to 8 VVs at time in given time stamp.

-reduce

This option is used to convert the given VV list into TDVV+COMPR volumes

-tpvv

This option is used to convert the given VV list into TPVV volumes.

-cpglist <cpglist>

This option provides a list of source and destination CPGs. It requires a comma-separated list of CPGs pairs, with the format <source>:<destination>. This option is only used with -convert option.

-allvvs

This option is used to consider all valid volumes for conversion into same CPG.

-stagger <seconds>

Minimum time difference in seconds between launching next conversion. This option is only used with –convert option.

-startt <YYYY-MM-DD-HH:MM>

Date and time to start conversions, into the format YYYY-MM-DD-HH:MM. This option is only used with –convert and --endt options.

-endt <YYYY-MM-DD-HH:MM>

Date and time to stop initiating conversions, into the format YYYY-MM-DD-HH:MM. This option is only used with –convert and –startt options.

-pri <low|med|high|auto>

This option is used to set conversion's task priority. It should be high or med or low or auto. The default priority is auto. The volumes with priority set to high will convert faster than other volumes with medium and low priority. If the priority is set to auto, then the array controls the speed of the operation based on the load in the array for all flash models. Auto priority behaves like medium priority for the converged flash models of the arrays containing Nearline (NL) and Fast Class (FC) drives.

Cleaning and compacting options:

-cleanwait <value>

Maximum number of minutes to wait for chunklet cleaning after each tune. <value> must be between 0 (tunes will be started immediately) and 720 (12 hours). The default value is 120 (2 hours).

-compactmb <value>{g|G|t|T}

Used in the inter-node and LD tuning phases. Once tunes have moved an amount of space greater than <value> the source CPG will be compacted. <value> can be between 0 (compact after every tune) and 2TiB. The default is 512GiB.

General tuning/conversion options:

-dr

Specifies that the command is a dry run and that the system will not be tuned. The result of the analysis will be displayed.

-f

Do not ask for confirmation.

-maxtasks <number>

Specifies the maximum number of individual inter-node tune and conversion tasks which the tunesys command can run simultaneously. <number> must be between 1 and 8. The default value is 2.

-maxnodetasks <number>

Specifies the maximum number of tunenodech tasks which the tunesys command can run simultaneously. <number> must be between 1 and 8. The default value is 1.

-noinitialcompacts

Do not compact CPGs before inter-node analysis phase.

-waittask

Waits for all tasks created by this command to complete before returning.

Specifiers

None.

Authority

- Super, Edit
- Any role granted the sys_tune right

Usage

- By default, tunesys compacts all CPGs and waits for chunklet cleaning before analyzing the system utilization balance and beginning the tuning process. If a dry run analysis is requested or the -noinitialcompacts option is used, no initial compaction takes place and the analysis will be based on the current free space.
- When the -cpg option is used, if the user space and snap space of a VV are from different CPGs, then only the space associated with the named CPG will be tuned.
- When the -cpg option is used, OR the user does not have All Domain privileges, intra-node analysis and tuning will not take place and the -maxchunk option is ignored. Intra-node tunes need access across all CPGs and all domains to work correctly.
- When canceling a tunesys command task, the canceltask command can return before a cancellation is completed. Therefore, resources reserved for the task might not be immediately available. This can prevent actions such as restarting the canceled tunesys command task for a few minutes.
- tunesys runs between 1 and 4 phases depending on the options specified.

- **Phase 1**—Analyzes and re-balances space usage between nodes (referred to as inter-node tuning).
- **Phase 2**—Analyzes and re-balances space between disks associated with each node (referred to as intra-node tuning).
- **Phase 3**—LD re-layout tunes. LDs where the current characteristics differ from the associated CPG are re-laid out to match the CPG.
- **Phase 4**—Conversion of TDVVx to the current dedup version in different CPGs or convert given VV list into same or different CPG.

- When `tunesys` is running, `showtask` will show the current and maximum number of phases for the task. The steps displayed indicate the number of MiB remaining to tune and the total number of MiB to tune.
- For example, a `showtask` step and phase display of 1/3 4096/65536 indicates that `tunesys` is in phase 1 of 3 (inter-node), it has completed 4096 MiB (4 GiB) of tunes out of an estimated total of 65536 MiB (64GiB).
- Only one `tunesys` task can be active at a time. A dry run can be performed while `tunesys` is active.
- The spindlepct percentage is compared against the following calculation:
  - max_disks = maximum number of disks associated with a node
  - min_disks = minimum number of disks associated with a node
- % Difference = (max_disks - min_disks) / (max_disks + min_disks) * 100
- The Difference value will generate a warning message if it is greater than the `-spindlepct` value.
- On StoreServ 7200/c and 8200 systems, `tunesys` is started automatically when new disks are detected by the `admithw` command.

Examples

In the following example:

- VVs will be re-balanced between nodes if there is an imbalance in space usage between nodes of more than 5% (nodepct value)
- PD chunklets will be re-balanced on each node if any disk has usage less than 20% than the average chunklet usage (chunkpct value).
- A maximum of four tune tasks will be run at a time (maxtasks value).
- A maximum of five chunklets will be moved at a time (maxchunk value).

```
cli% tunesys -f -nodepct 5 -chunkpct 20 -maxchunk 5 -maxtasks 4
Task 999 started
```

- Convert TDVVx to the current dedup version in two CPGs. The first does the conversions of all the TDVVx VVs from cpg1 into cpg2. The second does it from cpg3 to cpg5.

```
cli% tunesys -f -convertdedupversion cpg1:cpg2,cpg3:cpg5
        Task 1234 started
```

- Convert given VV list into TDVVs.

```
cli% tunesys -f -convert -tdvv volume1,volume1,volume3 Task 2345 started
```

tunevv

Syntax

The tunevv command uses one of the following syntax conventions:

```
tunevv usr_cpg <cpg> [options] <VV_name>
tunevv snp_cpg <cpg> [options] <VV_name>
tunevv restart [options] <VV_name>
tunevv rollback [options] <VV_name>
```

Description

Change the properties of a virtual volume that was created with the **createvv** command by associating it with a different CPG.

Options

-f

Forces the command. The command completes without prompting for confirmation.

-waittask

Specifies that the command will wait for any created tasks to complete.

-dr

Specifies that the command is a dry run and that no logical disks or virtual volumes are actually tuned. Cannot be used with the -tpvv -dedup, -full, -compr or -reduce options.

-pri <high|med|low|auto>

Specifies the priority of migration of a volume or a volume set. If this option is not specified, the default priority will be auto. The volumes with priority set to high will migrate faster than other volumes with medium and low priority. If the priority is set to auto, then the array controls the speed of the operation based on the load in the array for all flash models. Auto priority behaves like medium priority for the converged flash models of the arrays containing Nearline (NL) and Fast Class (FC) drives.

-cnt <count>

Specifies the number of identical virtual volumes to tune using an integer from 1 through 999. If not specified, one virtual volume is tuned. If not specified, one virtual volume is tuned. If the -cnt option is specified, then the subcommands, restart and rollback are not permitted.

Use the following options with the <usr_cpg> subcommand specified:

-tpvv

Indicates that the VV should be converted to a thin provision virtual volume. Cannot be used with the -dedup or -full options.

-tdvv

This option is deprecated, see -dedup.

-dedup

Indicates that the VV should be converted to a thin provision virtual volume that shares logical disk space with other instances of this volume type. Cannot be used with the -tpvv or -full options.

-full

Indicates that the VV should be converted to a fully provisioned virtual volume. Cannot be used with the -tpvv, -dedup, or -compr options.

-compr

Indicates that the VV should be converted to a compressed virtual volume. Cannot be used with the -full option.

-keepvv <VV_name>

Indicates that the original logical disks should be saved under a new virtual volume with the given name. Can only be used with the -tpvv, -dedup, -full, or -compr or options.

-src_cpg <CPG_name>

Indicates that only regions of the VV that are part of the specified CPG should be tuned to the destination USR or SNP CPG. This option is recommended when a VV belongs to an AO configuration and will avoid disrupting any optimizations already performed.

-slth <threshold>{g|G|t|T}

Slices threshold. Volumes above this size will be tuned in slices. <threshold> must be in multiples of 128GiB. Minimum is 128GiB. Default is 16TiB. Maximum is 16TiB.

-slsz <size>{g|G|t|T}

Slices size. Size of slice to use when volume size is greater than <threshold>. <size> must be in multiples of 128GiB. Minimum is 128GiB. Default is 2TiB. Maximum is 16TiB.

Specifiers

<VV_name>

Specifies the name of the existing virtual volume.

Subcommands

usr_cpg <cpg>

Moves the logical disks being used for user space to the specified CPG.

snp_cpg <cpg>

Moves the logical disks being used for snapshot space to the specified CPG.

restart

Restarts a tunevv command call that was previously interrupted because of component failure, or because of user initiated cancellation. This cannot be used on TPVVs or TDVVs.

rollback

Returns to a previously issued tunevv operation call that was interrupted. The canceltask command needs to run before the rollback. This cannot be used on TPVVs or TDVVs.

Authority

- Super, Edit
- Any role granted the vv_tune right

Usage

- Use of this command requires a Dynamic Optimization license. Contact your local HPE representative for information.
- There is a limit on the amount of data that can be tuned in a tunevv command. This is equivalent to ten 16TB virtual volumes. If this limit is exceeded, the command will be rejected.

- When the amount of space occupied by the user or snap space of a volume is greater than 16TiB (default), the volume is treated as a large volume and the tuning process is different. The volume is tuned in slices, each slice being 2TiB (default) in size. This allows a volume of up to 128TiB in size to be tuned between CPGs with much less free space required at the start of the tune. This is especially important when tuning a large volume to change its RAID type, set size, etc. back to the same class of disk as the amount of space required by the tune at the start will be the size of a slice rather than the equivalent of the whole space being tuned.
- The size threshold at which a volume is tuned in slices is controlled by the `-slth` option. The size of each slice can be changed using the `-slsz` option.
- It is only possible to tune one large volume at a time. The `-cnt` option cannot be used with volumes greater than 16TiB in size.
- When canceling a `tunevv` command task, the **`canceltask`** command can return before a cancellation is completed. Therefore, resources reserved for the task might not be immediately available. This can prevent actions like restarting the canceled tunevv command task. Scripts should use the **`waittask`** command in order to ensure orderly completion of the `tunevv` command cancellation before taking other actions.
- When the `-tpvv`, `-dedup`, `-full`, or `-compr` options for the `usr_cpg` subcommand are specified, the tune will automatically rollback on a failure. These options do not support VVs with remote copy. These options will only convert VVs using snapshots if the `-keepvv` option is used, but the snapshots will reside in the VV specified by the `-keepvv` option.
- When a VV is part of an AO configuration, regions can be spread between up to three different CPGs as part of the optimization process. Tuning all of the VV will result in any existing AO optimizations being discarded. The use of the `-src_cpg <CPG_name>` option will prevent this happening by only tuning the parts of the VV belonging to the source CPG to the specified USR or SNP CPG. This will maintain any changes AO has performed. A warning message will be displayed if a VV in an AO configuration is tuned without the use of the `-src_cpg` option.

Examples

In the following example, the logical disks used for user space are moved to CPG cpg_sn1.0_p for virtual volume nf_st_tp_22.0:

```
cli% tunevv usr_cpg cpg_sn1.0_p nf_st_tp_22.0
Task 999 started
```

In the following example, the logical disks used for user space are moved to CPG cpg_sn1.0_p for virtual volume nf_st_tp_22.0 and the virtual volume is converted to a tpvv:

```
cli% tunevv usr_cpg cpg_sn1.0_p -tpvv nf_st_tp_22.0
Task 1000 started
```

In the following example only the VV regions belonging to CPG cpg_FC_01 are tuned to the destination CPG cpg_SSD_01:

```
cli% tunevv usr_cpg cpg_SSD_01 -src_cpg cpg_FC_01 AO_vv_01
Task 1001 started
```

Update Commands

updatesnapspace

Syntax

```
updatesnapspace [<VV_name>...|<pattern>...]
```

Description

The `updatesnapspace` command starts a non-cancelable task to update the snapshot space usage accounting. The snapshot space usage displayed by "`showvv -hist`" is not necessarily the current usage and the SpaceCalcTime column will show when it was last calculated. This command causes the system to start calculating current snapshot space usage. If one or more VV names or patterns are specified, only the specified VVs will be updated. If none are specified, all VVs will be updated.

The command will return immediately, displaying the task ID associated with the update. To wait for completion, use the **`waittask`** command.

Options

None.

Specifiers

`<VV_name>...`

Specifies the virtual volume name to update. This specifier can be repeated to specify multiple virtual volumes. This specifier is not required.

`<pattern>...`

Specifies that the virtual volumes matching the specified glob-style pattern is updated. (For more information, enter `clihelp sub,globpat` or see **Glob-style patterns**.) This specifier can be repeated. This specifier is not required.

Authority

- Super, Edit
- Any role granted the `snapspace_update` right

Usage

- If one or more virtual volume names or patterns are specified, only the specified virtual volumes are updated. If no volume names or patterns are specified, all virtual volumes are updated.
- To check the snapshot space usage, use the "`showvv -s`" command (see **`showvv`**).

Examples

The following example displays the actual snapshot space used by all VVs:

```
cli% updatesnapspace
Task 2 has been started to calculate actual space usage.
```

Issuing a `waittask` command displays the information about the task in process or if it has been processed:

```
cli% waittask 2
Task 2 done
```

The following example displays the task information using `-d <task_id>`:

```
cli% showtask -d 2
Id Type Name Status Phase Step ----------StartTime--------- ---------FinishTime---------
2  snapspace_accounting ss_accounting Done 0/0 0/0 Mon Sep 22 18:43:35 PDT
2014 Mon Sep 22 18:43:35 PDT 2014

Detailed status:
{Thu Sep 22 18:43:35 PDT 2014} Created task.
{Thu Sep 22 18:43:35 PDT 2014} Started snapshot usage data collection process for VVs
{Thu Sep 22 18:43:35 PDT 2014} Updated snapshot usage data for VV emaildb
{Thu Sep 22 18:43:35 PDT 2014} Updated snapshot usage data for VV www
{Thu Sep 22 18:43:35 PDT 2014} Updated snapshot usage data for VV snapname
{Thu Sep 22 18:43:35 PDT 2014} Finished snapshot usage data collection process
```

updatevv

Syntax

`updatevv [options] <VV_name | VV_set> ...`

Description

The `updatevv` command updates a snapshot Virtual Volume (VVs) with a new snapshot.

Options

`-ro`

Specifies that if the specified virtual volume (<VV_name>) is a read- write snapshot the snapshot's read-only parent volume is also updated with a new snapshot if the parent volume is not a member of a virtual volume set. If this option is not specified, the virtual volumes, as indicated with the <VV_name> specifier, are replaced by a new snapshot.

See Notes for additional information.

`-f`

Specifies that the command is forced. If this option is not used, the command requires confirmation before proceeding with its operation.

Specifiers

`<VV_name | VV_set> ...`

Specifies the name(s) of the snapshot virtual volume(s) or virtual volume set(s) to be updated. They must all be of the same type (read-only or read-write). When a <VV_set> or multiple <VV_name>s are specified group-consistent snapshots are taken.

See **`createvvgroup`** for more information. The vv set name must start with "set:".

Authority

- Super, Edit, Browse
- Any role granted the `vv_update` right

Usage

- For each snapshot VV name specified, the `updatevv` command first verifies that the user is in the correct domain and has privileges of edit or higher for the VVs in the associated domain(s). If that fails then it verifies the user is in the Access Control List (ACL) and has permission to run the `updatevv` command on the specified VVs.
- The `updatevv` command puts the snapshot in standby mode when updating the VV maps. Hosts will see disk resets when in standby mode. The command then updates the snapshots as follows:

- If the -ro option is specified and the specified VV name is a read-write snapshot, the VV read-only parent is also updated with a new snapshot.
- If the read-only parent has multiple read-write snapshots, the `updatevv` operation will not update the original parent. The updated read-write snapshot will be created under a new RO snapshot.

- When setting up the ACL for virtual volume sets, only their associated virtual volumes must be stored in the ACL. This limitation is in place because the system will only check permissions to run `updatevv` on the Virtual Volumes associated with the virtual volume set.
- When setting up the ACL for a read-write snapshot, this snapshot's read-only parent volume may require an entry in the ACL as well. This limitation stems from the fact that updating a read-only snapshot with the -ro option will also update the snapshot's parent volume as well. If the read-write snapshot is updated without the -ro option, then its read-only parent volume does not require an entry in the ACL, since the parent volume will not be updated.
- Running concurrent `updatevv` sessions for VV sets or list of VVs which have common VVs can have unpredictable results.
- Running `updatevv` on VVs that are in use can produce unpredictable results. Hosts and applications usually cache data or save states and often do not handle VVs changing underneath them.

Examples

The following example displays the snapshot update of snapshot VV avvro:

```
cli% updatevv -f avvro
Updating VV avvro
```

The following example will update all VVs in VV set vv_set:

```
cli% updatevv set:vv_set
```

Upgrade Commands

upgradecage

Syntax

```
upgradecage [options <arg>] <cagename>...
upgradecage -a
upgradecage -parallel [-wait]
upgradecage -status [<cagename>...]
```

Description

Downloads new firmware into the specified cage.

Options

-a

Specifies that all drive cages are upgraded at one time with new firmware (do not specify *<cagename>*).

-parallel [-wait]

All drive cages are upgraded in parallel by interface card domain. If -wait is specified, the command will not return until the upgrades are completed. Otherwise, the command returns immediately and completion of the upgrade can be monitored with the -status option.

-status [*<cagename>* ...]

Prints status of the current `upgradecage` operation in progress or the last executed `upgradecage` operation. Specifying one or more a `<cagename>` variables returns a filtered result displaying only the specified cages.

Specifiers

<cagename>

Specifies the cage name as displayed in the **Name** column of the `showcage` command results.

Authority

- Super, Service
- Any role granted the `cage_upgrade` right

Usage

- Running this command requires access to all domains.
- Before executing the `upgradecage` command, issue the **showcage** command to obtain the names of the drive cages in the system.
- When you issue the `upgradecage` command, the drive cage becomes temporarily degraded as the system upgrades each interface card.

Examples

The following example shows the upgrade of cages with firmware level of 2.11 to firmware level 2.12:

```
cli% upgradecage -a
Upgrading cage cage0 from rev 2.11 to revision in file
/opt/tpd/fw/cage/gbod/Oban-Release-2.12-build-001.bin
```

```
                                  Upgrade of cage cage0 status: succeeded
        Beginning test after upgrade for cage0
        cage0 passed test after upgrade
        Skipping cage cage1 already up to date at rev 2.12
```

The following example of using -parallel and -status:

```
cli% upgradecage -parallel
 Upgrading cages in parallel. Please monitor the progress with "upgradecage -status".

cli% upgradecage -status

Cage  OverallState StartTime               StopTime  Current State     Progress (%)    Pending Completed
FailureDetail
cage0 ongoing      2017-10-24 14:11:10 PDT --        sepB    downloading     100     --      sepA      --
cage1 ongoing      2017-10-24 14:11:10 PDT --        sepB    downloading        94     --      sepA      --
```

upgradepd

Syntax

```
upgradepd [-f] [-skiptest] {-a | -w <WWN> ... | <PD_ID> ...}
```

Description

The upgradepd command upgrades the physical disk firmware.

Options

-f

Upgrades the physical disk firmware without requiring confirmation.

-skiptest

Skips the 10 second diagnostic test normally completed after each physical disk upgrade.

-a

Specifies that all physical disks with valid IDs and whose firmware is not current are upgraded. If this option is not specified, then either the -w option or PD_ID specifier must be issued on the command line.

-w <WWN>...

Specifies that the firmware of either one or more physical disks, identified by their WWNs, is upgraded. If this option is not specified, then either the -a option or PD_ID specifier must be issued on the command line.

Specifiers

<PD_ID>...

Specifies that the firmware of either one or more physical disks identified by their IDs (PD_ID) is upgraded. If this specifier is not used, then the -a or -w option must be issued on the command line.

Authority

- Super, Service
- Any role granted the pd_upgrade right

Usage

- Access to all domains is required to run this command.
- Physical disks can be upgraded while I/O is occurring.
- If a disk with RAID-0 chunklets is upgraded, I/O to those chunklets at the time of the upgrade results in data loss. This is not applicable to other RAID types.
- Specify the -w option when upgrading un-admitted physical disks.

Examples

The following example displays the upgrade of all physical disks:

```
cli% upgradepd -a
```

Wait Command

waittask

Syntax

```
waittask -v <task_ID>
waittask [<task_ID>...]
```

Description

The `waittask` command asks the CLI to wait for a task to complete before proceeding. The command automatically notifies you when the specified task is finished.

Options

`-v <task_ID>`

Displays the detailed status of the task specified by <task_ID> as it executes. When the task completes, this command exits.

`-q`

Quiets; do not report the end state of the tasks, only wait for them to exit.

Specifiers

`<task_ID>`

Indicates one or more tasks to wait for using their task IDs. When no task IDs are specified, the command waits for all non-system tasks to complete. To wait for system tasks, <task_ID> must be specified.

Authority

Any role in the system

Usage

- See the *HPE 3PAR Command Line Interface Administrator Guide* for additional information and examples regarding task management and task management commands.
- This command returns an error if any of the tasks it is waiting for fail.

Examples

The following example shows how to wait for a task using the task ID. When successful, the command returns only after the task completes.

```
cli% waittask 1
Task 1 done
```

Mapping roles and rights

NOTE: The `show`, `hist`, and `stat` commands can be used by any role in the system. The `Audit` role cannot access the HPE 3PAR CLI.

In this chapter:

- **3PAR_AO Role**
- **3PAR_RM Role**
- **Audit Role**
- **Basic_edit Role**
- **Browse Role**
- **Create Role**
- **Edit Role**
- **Service Role**
- **Super Role**

3PAR_AO Role

Used internally by HPE for Adaptive Optimization operations.

Table 7: 3PAR_AO rights

Right	Description of right
aocfg_create	Create Adaptive Optimization configuration
aocfg_remove	Remove an Adaptive Optimization configuration
aocfg_set	Set properties on an Adaptive Optimization configuration
aocfg_start	Start executing an Adaptive Optimization configuration
cotask_set	Executes a CO approved request or cancels a pending dual sign on request
cpg_compact	Consolidate space in a Common Provisioning Group
password_checkown	Check a user's own password
password_setown	Set a user's own password

Table Continued

Right	Description of right
reg_move	Move virtual volume regions between logical disks
sched_create	Create a schedule for tasks
sched_setown	Set a user's own scheduled task properties
sshkey_remove	Remove users public SSH keys
sshkey_set	Set the public SSH key for a user

3PAR_RM Role

Used internally by HPE for operations required by Recovery Manager.

Table 8: `3PAR_RM` rights

Right	Description of right
cotask_set	Executes a CO approved request or cancels a pending dual sign on request
groupsv_create	Create a virtual volume group
host_set	Set host properties
ld_remove	Remove a logical disk
password_checkown	Check a user's own password
password_setown	Set a user's own password
rcopy_start	Start the remote copy service
rcopy_stop	Stop the remote copy service
rcopy_sync	Synchronize remote copy groups
rcopygroup_create	Create a remote copy group
rcopygroup_set	Set remote copy group properties
rcopygroup_start	Start remote copy group
rcopygroup_stop	Stop a remote copy group

Table Continued

Right	Description of right
rcopytarget_admit	Add a target to a remote copy volume
rcopytarget_create	Create a remote copy target
rcopytarget_set	Set remote copy target properties
sched_create	Create a schedule for tasks
sched_setown	Set a user's own scheduled task properties
sshkey_remove	Remove a user’s public SSH keys
sshkey_set	Set the public SSH key for a user
sv_create	Create a virtual copy volume
sv_promote	Promote a virtual copy volume
sv_remove	Remove virtual copy of a volume
sv_set	Set properties of a virtual copy of a volume
svgrp_promote	Promote multiple virtual copy volumes
task_start	Run commands in the background
user_setown	Set a user's own properties
vlun_create	Export a virtual volume as a SCSI LUN
vlun_remove	Remove a VLUN
vv_create	Create a virtual volume
vvcopy_remove	Remove physical copy of a volume
vvcopy_set	Set properties of a physical copy of a volume
vvset_create	Create a virtual volume set
vvset_set	Set virtual volume set properties

Audit Role

For security scanners to perform a scan of the HPE 3PAR OS file system. An audit user has no access to the CLI.

Table 9: `Audit` rights

Right	Description of right
audit_chroot	Secure access to a read only chroot of the InForm OS file system

Basic_edit Role

Rights are similar to Edit role, but more restricted, specifically in the ability to remove objects such as volumes, VLUNs, and hosts.

Table 10: `basic_edit` rights

Right	Description of right
alert_set	Set alert properties
cotask_set	Executes a CO approved request or cancels a pending dual sign on request
cpg_compact	Consolidate space in a Common Provisioning Group
cpg_create	Create Common Provision Group
domain_create	Create a domain
domain_moveto	Move objects from one domain to another
domainset_create	Create a domain set
domainset_set	Set domain set properties
groupsv_create	Create a virtual volume group
groupvvcopy_create	Create a physical copy volume group
host_create	Create a host
host_set	Set host properties
hostset_create	Create paths to a system host set
hostset_set	Set host set properties
ld_remove	Remove a logical disk
password_checkown	Check a user's own password

Table Continued

Right	Description of right
password_setown	Set a user's own password
rcopy_start	Start the remote copy service
rcopy_stop	Stop the remote copy service
rcopy_sync	Synchronize remote copy groups
rcopygroup_create	Create a remote copy group
rcopygroup_remove	Remove a remote copy group
rcopygroup_set	Set remote copy group properties
rcopygroup_start	Start remote copy group
rcopygroup_stop	Stop a remote copy group
rcopylink_admit	Add remote copy links
rcopylink_dismiss	Dismiss remote copy links
rcopytarget_admit	Add a target to a remote copy volume
rcopytarget_create	Create a remote copy target
rcopytarget_dismiss	Dismiss a remote copy target from a group
rcopytarget_remove	Remove a remote copy target
rcopytarget_set	Set remote copy target properties
rcopyvv_admit	Add a volume to a remote copy volume group
rcopyvv_dismiss	Dismiss a virtual volume from a remote copy group
sched_create	Create a schedule for tasks
sched_setown	Set a user's own scheduled task properties
snapspace_update	Update snapshot space accounting information
sshkey_remove	Remove user's public SSH keys

Table Continued

Right	Description of right
sshkey_set	Set the public SSH key for a user
sv_create	Create a virtual copy volume
sv_promote	Promote a virtual copy volume
sv_remove	Remove virtual copy of a volume
sv_set	Set properties of a virtual copy of a volume
svgrp_promote	Promote multiple virtual copy volumes
user_setown	Set a user's own properties
vlun_create	Export a virtual volume as a SCSI LUN
vv_create	Create a virtual volume
vv_grow	Grow a virtual volume
vv_tune	Tune a virtual volume
vvcopy_create	Create physical copy volume
vvcopy_promote	Promote a physical volume back to its base volume
vvcopy_remove	Remove physical copy of a volume
vvcopy_set	Set properties of a physical copy of a volume
vvset_create	Create a virtual volume set
vvset_set	Set virtual volume set properties

Browse Role

Rights are limited to read-only access.

Table 11: `browse` rights

Right	Description of right
password_checkown	Check a user's own password
password_setown	Set a user's own password
save_pel	Save PEL counters
sshkey_remove	Remove user's public SSH keys
sshkey_set	Set the public SSH key for a user
task_wait	Wait for task completion
user_setown	Set a user's own properties
vv_update	Remove old virtual copy and create a new virtual copy of a volume

CO Role

Rights to approve the Compliance WORM changes.

Table 12: CO rights

Right	Description of right
corequest_remove	Remove the specified request(s) from queue
corequest_set	Approve, deny, execute, or cancel a dual sign-on request, or modify queue size
password_checkown	Check a user's own password
password_setown	Set a user's own password
sshkey_remove	Remove users public SSH keys
sshkey_set	Set the public SSH key for a user
user_setown	Set a user's own properties

Create Role

create Rights are limited to creation of objects such as volumes, CPGs, hosts, remote copy groups, remote copy targets, and schedules.

Table 13: `create` rights

Right	Description of right
cotask_set	Executes a CO approved request or cancels a pending dual sign on request
cpg_create	Create Common Provision Group
domainset_annotate	Annotate the comment of a domain set
domainset_create	Create a domain set
groupsv_create	Create a virtual volume group
groupvvcopy_create	Create a physical copy volume group
host_create	Create a host
hostset_annotate	Annotate the comment of a host set
hostset_create	Create paths to a system host set
password_checkown	Check a user's own password
password_setown	Set a user's own password
rcopy_start	Start the remote copy service
rcopy_stop	Stop the remote copy service
rcopy_sync	Synchronize remote copy groups
rcopygroup_create	Create a remote copy group
rcopygroup_set	Set remote copy group properties
rcopygroup_start	Start remote copy group
rcopygroup_stop	Stop a remote copy group
rcopytarget_admit	Add a target to a remote copy volume
rcopytarget_create	Create a remote copy target
rcopytarget_set	Set remote copy target properties

Table Continued

Right	Description of right
sched_create	Create a schedule for tasks
sched_setown	Set a user's own scheduled task properties
sshkey_remove	Remove user's public SSH keys
sshkey_set	Set the public SSH key for a user
sv_create	Create a virtual copy volume
user_setown	Set a user's own properties
vlun_create	Export a virtual volume as a SCSI LUN
vv_create	Create a virtual volume
vvcopy_create	Create physical copy volume
vvset_annotate	Annotate the comment of a virtual volume set
vvset_create	Create a virtual volume set

Edit Role

Rights are granted to most operations, such as for creating, editing, and removing virtual volumes.

Table 14: `edit` rights

Right	Description of right
ch_move	Move chunklets from one physical disk to another
ch_movetospare	Move chunklets from a physical disk to spare space
cotask_set	Execute a CO approved request or cancel a pending dual sign-on request
cpg_compact	Consolidate space in a Common Provisioning Group
cpg_create	Create a Common Provision Group
cpg_remove	Remove a Common Provision Group
cpg_set	Set Common Provision Group properties

Table Continued

Right	Description of right
domain_moveto	Move objects from one domain to another
domainset_annotate	Annotate the comment of a domain set
domainset_create	Create a domain set
domainset_set	Set domain set properties
fpg_create	Create a file provisioning group
fpg_grow	Grow a file provisioning group
fpg_remove	Remove a file provisioning group
fpg_set	Set properties of a file provisioning group
fs_sset	Sets or modifies File Services related properties
fsarchive_remove	Remove retention setting, retained files, or validation jobs
fsarchive_set	Set profile, expiry date, and legal hold
fsaudit_set	Set File Access Audit settings
fsav_set	Configure antivirus properties for file services
fsav_start	Start AV service or start/resume on demand AV scan task
fsav_stop	Stop AV service or stop/pause on demand AV scan task
fshare_create	Create file shares
fshare_remove	Remove file shares
fshare_set	Set or modify file share properties
fsip_create	Assign network config to a Virtual File Server
fsip_remove	Remove network config of a Virtual File Server
fsip_set	Modify network config of a Virtual File Server
fsnap_create	Create file store snapshots

Table Continued

Right	Description of right
fsnap_remove	Remove file store snapshots
fsndmp_set	Set or modify NDMP related properties
fsquota_set	Set user, group, and fstore quotas for a Virtual File Server
fsroute_set	Create, modify, and remove target routes
fstore_create	Create File Store
fstore_remove	Remove File Store
fstore_set	Set File Store properties
groupsv_create	Create a virtual volume group
groupvvcopy_create	Create a physical copy volume group
host_create	Create a host
host_remove	Remove a host
host_set	Set host properties
hostset_annotate	Annotate the comment of a host set
hostset_create	Create paths to a system host set
hostset_set	Set host set properties
ld_compact	Consolidate space in a logical disk
ld_create	Create logical disk
ld_remove	Remove a logical disk
password_checkown	Check a user's own password
password_setown	Set a user's own password
pd_movetospare	Move used chunklets from a physical disk to spare space
pd_tune	Tune a physical disk

Table Continued

Right	Description of right
qos_set	Set QoS rule properties
rclink_check	Remote copy link validity checks
rcopy_start	Start the remote copy service
rcopy_stop	Stop the remote copy service
rcopy_sync	Synchronize remote copy groups
rcopygroup_create	Create a remote copy group
rcopygroup_remove	Remove a remote copy group
rcopygroup_set	Set remote copy group properties
rcopygroup_start	Start remote copy group
rcopygroup_stop	Stop a remote copy group
rcopylink_admit	Add remote copy links
rcopylink_dismiss	Dismiss remote copy links
rcopytarget_admit	Add a target to a remote copy volume
rcopytarget_create	Create a remote copy target
rcopytarget_dismiss	Dismiss a remote copy target from a group
rcopytarget_remove	Remove a remote copy target
rcopytarget_set	Set remote copy target properties
rcopyvv_admit	Add a volume to a remote copy volume group.
rcopyvv_dismiss	Dismiss a virtual volume from a remote copy group
reg_move	Move virtual volume regions between logical disks
save_pel	Save PEL counters
sched_create	Create a schedule for tasks

Table Continued

Right	Description of right
sched_setown	Set a user's own scheduled task properties
snapspace_update	Update snapshot space accounting information
space_free	Free space in a specific Virtual Volume
sshkey_remove	Remove user's public SSH keys
sshkey_set	Set the public SSH key for a user
statch_set	Set statistics collection mode on specified chunklets
statpdch_set	Set statistics collection mode on PD chunklets
sv_create	Create a virtual copy volume
sv_promote	Promote a virtual copy volume
sv_remove	Remove virtual copy of a volume
sv_set	Set properties of a virtual copy of a volume
svgrp_promote	Promote multiple virtual copy volumes
sys_tune	Tune the system layout
task_cancel	Cancel a task
task_remove	Remove a task
task_set	Edit tasks
task_start	Run commands in the background
task_wait	Wait for task completion
user_setown	Set a user's own properties
useracl_set	Set user Access Control List properties
vfs_create	Create VFS
vfs_remove	Remove VFS

Table Continued

Right	Description of right
vfs_set	Set VFS properties
vlun_create	Export a virtual volume as a SCSI LUN
vlun_remove	Remove a VLUN
vv_check	Virtual volume validity checks
vv_create	Create a virtual volume
vv_grow	Grow a virtual volume
vv_import	Import a remote virtual volume to local storage
vv_remove	Remove a base virtual volume
vv_set	Set virtual volume properties
vv_tune	Tune a virtual volume
vv_update	Remove old virtual copy and create a new virtual copy of a volume
vvcopy_create	Create physical copy volume
vvcopy_promote	Promote a physical volume back to its base volume
vvcopy_remove	Remove physical copy of a volume
vvcopy_set	Set properties of a physical copy of a volume
vvset_annotate	Annotate the comment of a virtual volume set
vvset_create	Create a virtual volume set
vvset_set	Set virtual volume set properties

Security Role

security_admin rights are granted to create and remove users except super users.

Right	Description of right
password_checkown	check a user's own password

Table Continued

password_setown	Set a user's own password
sshkey_remove	Remove users public SSH keys
user_create	Create a user
user_remove	Remove a user
user_setany	Set any user properties
user_setown	Set a user's own properties

Service Role

Rights are limited to operations required to service the storage server.

Table 15: `service` rights

Right	Description of right
alert_remove	Remove alerts
alert_set	Set alert properties
battery_set	Set battery properties
cage_locate	Locate a drive cage
cage_set	Set cage properties
cage_upgrade	Upgrade cage firmware
ch_move	Move chunklets from one physical disk to another
ch_movetospare	Move chunklets from a physical disk to spare space
cim_set	Set CIM server properties
cim_start	Start the CIM server
cim_stop	Stop the CIM server
control_encryption_status	Show encryption status
cotask_set	Execute a CO approved request or cancel a pending dual sign-on request

Table Continued

Right	Description of right
cpg_set	Set Common Provision Group properties
date_set	Set system date and time properties
event_debug	Debug events
fs_set	Set or modify File Services related properties
fs_start	Initialize and start File Services
fs_stop	Stop File Services
fsarchive_start	Start or resume validation jobs
fsarchive_stop	Stop or pause validation jobs
fsconf_backup	Backup configuration for file services
fsconf_restore	Restore configuration for file services
fsnapclean_start	Start or resume an on-demand reclamation task
fsnapclean_stop	Stop or pause an on-demand reclamation task
fsndmp_start	Start iSCSI and NDMP Services
fsndmp_stop	Stop iSCSI and NDMP Services
host_service	Service a host
ld_check	Logical disk validity checks
ld_create	Create logical disk
ld_remove	Remove a logical disk
ld_start	Start a logical disk
license_set	Set the license key
mag_control	Take magazines or disks on or off a loop
mag_service	Service a magazine or physical disk

Table Continued

Right	Description of right
net_set	Set administration network interface configuration
node_locate	Locate a controller node by blinking its LEDs
node_service	Service a controller node
node_set	Set controller node properties
node_shutdown	Shut down a controller node
ntp_set	Set the system NTP server
password_checkown	Check a user's own password
password_setown	Set a user's own password
pd_admit	Admit a physical disk to the system
pd_check	Physical disk diagnostic checks
pd_control	Spin up or spin down a physical disk
pd_dismiss	Dismiss a physical disk from the system
pd_movereloc	Move relocated chunklets from one physical disk to another
pd_movetospare	Move used chunklets from a physical disk to spare space
pd_set	Set physical disk properties
pd_upgrade	Update physical disk firmware
port_check	Perform diagnostic tests on ports
port_control	Control and configure ports
recoveryauth_status	Show recovery auth status
save_pel	Save PEL counters
sched_create	Create a schedule for tasks
sched_setany	Set any user scheduled task properties

Table Continued

Right	Description of right
sched_setown	Set a user's own scheduled task properties
snmp_check	Check and send SNMPv2 test trap
spare_create	Create spare chunklet
spare_remove	Remove spare chunklets from a PD
sshkey_remove	Remove user's public SSH keys
sshkey_set	Set the public SSH key for a user
sys_locate	Locate a system by blinking its LEDs
sys_set	Set system properties
sys_shutdown	Shut down the system
task_cancel	Cancel a task
task_start	Run commands in the background
task_wait	Wait for task completion
user_remove	Remove a user
vv_check	Virtual volume validity checks
vv_start	Start a virtual volume
wsapi_set	Set the Web Services API server properties
wsapi_start	Start the Web Services API server
wsapi_stop	Stop the Web Services API server

Super Role

Rights are granted to all operations.

Table 16: `super` rights

Right	Description of right
alert_remove	Remove alerts
alert_set	Set alert properties
authparam_set	Set LDAP authentication properties
battery_set	Set battery properties
cage_locate	Locate a drive cage
cage_set	Set cage properties
cage_upgrade	Upgrade cage firmware
ch_move	Move chunklets from one physical disk to another
ch_movetospare	Move chunklets from a physical disk to spare space
cim_set	Set CIM server properties
cim_start	Start the CIM server
cim_stop	Stop the CIM server
control_encryption	Control the encryption
control_encryption_status	Show encryption status
cotask_set	Execute a CO approved request or cancel a pending dual sign-on request
cpg_compact	Consolidate space in a Common Provisioning Group
cpg_create	Create Common Provision Group
cpg_remove	Remove Common Provision Group
cpg_set	Set Common Provision Group properties
date_set	Set system date and time properties
domain_create	Create a domain

Table Continued

Right	Description of right
domain_moveto	Move objects from one domain to another
domain_remove	Remove a domain
domain_set	Set domain properties
domainset_annotate	Annotate the comment of a domain set
domainset_create	Create a domain set
domainset_set	Set domain set properties
event_debug	Debug events
fpg_create	Create a file provisioning group
fpg_grow	Grow a file provisioning group
fpg_remove	Remove a file provisioning group
fpg_set	Set properties of a file provisioning group
fs_set	Set or modify File Services related properties
fs_start	Initialize and start File Services
fs_stop	Stop File Services
fsarchive_remove	Remove retention setting, retained files, or validation jobs
fsarchive_set	Set profile, expiry date, and legal hold
fsarchive_start	Start or resume validation jobs
fsarchive_stop	Stop or pause validation jobs
fsaudit_set	Set File Access Audit settings
fsav_set	Configure antivirus properties for file services
fsav_start	Start AV service or start/resume on demand AV scan task
fsav_stop	Stop AV service or stop/pause on demand AV scan task

Table Continued

Right	Description of right
fsconf_backup	Backup configuration for file services
fsconf_restore	Restore configuration for file services
fsgroup_create	Create File Services local group
fsgroup_remove	Remove File Services local group
fsgroup_set	Modify File Services local group
fshare_create	Creates file shares
fshare_remove	Remove file shares
fshare_set	Set or modify file share properties
fsip_create	Assign network config to a Virtual File Server
fsip_remove	Remove network config of a Virtual File Server
fsip_set	Modify network config of a Virtual File Server
fsnap_create	Create file store snapshots
fsnap_remove	Remove file store snapshots
fsnapclean_start	Start or resume an on-demand reclamation task
fsnapclean_stop	Stop or pause an on-demand reclamation task
fsndmp_set	Set or modify NDMP related properties
fsndmp_start	Start iSCSI and NDMP Services
fsndmp_stop	Stop iSCSI and NDMP Services
fsnetwork_create	Create Network
fsnetwork_remove	Remove Network
fsnetwork_set	Set Network properties
fsquota_set	Set user, group and fstore quotas for a Virtual File Server

Table Continued

Right	Description of right
fsroute_set	Create, modify and remove target routes
fstore_create	Create File Store
fstore_remove	Removes File Store
fstore_set	Set File Store properties
fsuser_create	Create File Services local user
fsuser_remove	Remove File Services local user
fsuser_set	Modify File Services local user
groupsv_create	Create a virtual volume group
groupvvcopy_create	Create a physical copy volume group
host_create	Create a host
host_remove	Remove a host
host_service	Service a host
host_set	Set host properties
hostset_annotate	Annotate the comment of a host set
hostset_create	Create paths to a system host set
hostset_set	Set host set properties
ld_check	Logical disk validity checks
ld_compact	Consolidate space in a logical disk
ld_create	Create logical disk
ld_remove	Remove a logical disk
ld_start	Start a logical disk
license_set	Set the license key

Table Continued

Right	Description of right
mag_control	Take magazines or disks on or off a loop
mag_service	Service a magazine or physical disk
net_set	Set administration network interface configuration
node_locate	Locate a controller node by blinking its LEDs
node_service	Service a controller node
node_set	Set controller node properties
node_shutdown	Shut down a controller node
ntp_set	Set the system NTP server
password_checkany	Check any user password
password_checkown	Check a user's own password
password_setany	Change any user password
password_setown	Set a user's own password
pd_admit	Admit a physical disk to the system
pd_check	Physical disk diagnostic checks
pd_control	Spin up or spin down a physical disk
pd_dismiss	Dismiss a physical disk from the system
pd_movereloc	Move relocated chunklets from one physical disk to another
pd_movetospare	Move used chunklets from a physical disk to spare space
pd_set	Set physical disk properties
pd_tune	Tune a physical disk
pd_upgrade	Update physical disk firmware
port_check	Perform diagnostic tests on ports

Table Continued

Right	Description of right
port_control	Control and configure ports
qos_set	Set QoS rule properties
rclink_check	Remote copy link validity checks
rcopy_start	Start the remote copy service
rcopy_stop	Stop the remote copy service
rcopy_sync	Synchronize remote copy groups
rcopygroup_create	Create a remote copy group
rcopygroup_remove	Remove a remote copy group
rcopygroup_set	Set remote copy group properties
rcopygroup_start	Start remote copy group
rcopygroup_stop	Stop a remote copy group
rcopylink_admit	Add remote copy links
rcopylink_dismiss	Dismiss remote copy links
rcopytarget_admit	Add a target to a remote copy volume
rcopytarget_create	Create a remote copy target
rcopytarget_dismiss	Dismiss a remote copy target from a group
rcopytarget_remove	Remove a remote copy target
rcopytarget_set	Set remote copy target properties
rcopyvv_admit	Add a volume to a remote copy volume group
rcopyvv_dismiss	Dismiss a virtual volume from a remote copy group
recoveryauth_status	Show recovery auth status
reg_move	Move virtual volume regions between logical disks

Table Continued

Right	Description of right
save_pel	Save PEL counters
sched_create	Create a schedule for tasks
sched_setany	Set any user scheduled task properties
sched_setown	Set a user's own scheduled task properties
snapspace_update	Update snapshot space accounting information
snmp_check	Check and send SNMPv2 test trap
snmpmgr_add	Add an SNMP manager
snmpmgr_remove	Remove an SNMP manager
snmppw_remove	Remove an SNMP community strings
snmppw_set	Set SNMP community strings
snmpuser_create	Create an SNMP user
snmpuser_remove	Remove an SNMP user
space_free	Free space in a specific Virtual Volume
spare_create	Create spare chunklet
spare_remove	Remove spare chunklets from a PD
spcred_remove	Remove SP credentials
sshbanner_set	Set or remove the SSH banner
sshkey_remove	Remove user's public SSH keys
sshkey_set	Set the public SSH key for a user
statch_set	Set statistics collection mode on specified chunklets
statpdch_set	Set statistics collection mode on PD chunklets
sv_create	Create a virtual copy volume

Table Continued

Right	Description of right
sv_promote	Promote a virtual copy volume
sv_remove	Remove virtual copy of a volume
sv_set	Set properties of a virtual copy of a volume
svgrp_promote	Promote multiple virtual copy volumes
sys_locate	Locate a system by blinking its LEDs
sys_set	Set system properties
sys_shutdown	Shut down the system
sys_tune	Tune the system layout
sysmgr_set	Set System Manager Properties
task_cancel	Cancel a task
task_remove	Remove a task
task_set	Edit tasks
task_start	Run commands in the background
task_wait	Wait for task completion
template_create	Create a template for a VV, CPG, or LD
template_remove	Remove template for a VV, CPG, or LD
template_set	Set template for a VV, CPG, or LD
user_create	Create a user
user_remove	Remove a user
user_setany	Set any user properties
user_setown	Set a user's own properties
useracl_set	Set user Access Control List properties

Table Continued

Right	Description of right
userconn_remove	Remove user connections
vfs_create	Create a VFS
vfs_remove	Remove a VFS
vfs_set	Set VFS properties
vlun_create	Export a virtual volume as a SCSI LUN
vlun_remove	Remove a VLUN
vv_admit	Admit remote virtual volume to the system
vv_check	Virtual volume validity checks
vv_create	Create a virtual volume
vv_grow	Grow a virtual volume
vv_import	Import a remote virtual volume to local storage
vv_remove	Remove a base virtual volume
vv_set	Set virtual volume properties
vv_start	Start a virtual volume
vv_tune	Tune a virtual volume
vv_update	Remove old virtual copy and create a new virtual copy of a volume
vvcopy_create	Create physical copy volume
vvcopy_promote	Promote a physical volume back to its base volume
vvcopy_remove	Remove physical copy of a volume
vvcopy_set	Set properties of a physical copy of a volume
vvset_annotate	Annotate the comment of a virtual volume set
vvset_create	Create a virtual volume set

Table Continued

Right	Description of right
vvset_set	Set virtual volume set properties
wsapi_set	Set the Web Services API server properties
wsapi_set_authenmode	Set the 3PAR UI API Server Authentication Mode
wsapi_start	Start the Web Services API server
wsapi_stop	Stop the Web Services API server
wsapisession_remove	Remove WSAPI session

Command Authorities

List of authorities for Command Line Interface commands.

Command	Authority
addsnmpmgr	Super Any role granted the snmpmgr_add right
admithw	Super, Service
admitpd	Super, Service Any role granted the pd_admit right
admitrcopylink	Super, Edit Any role granted the rcopylink_admit right
admitrcopytarget	Super, Edit Any role granted the rcopytarget_admit right
admitrcopyvv	Super, Edit Any role granted the rcopyvv_admit right
admitvv	Super Any role granted the vv_admit right
canceltask	Super, Service, Edit Any role granted the task_cancel right
checkld	Super, Service Any role granted the ld_check right
checkpassword	Super, Edit, Browse, Service Any role granted the password_checkany right can check the password of any user Any role granted the password_checkown right can only check their own password
checkpd	Super, Service Any role granted the pd_check right

Table Continued

Command	Authority
checkport	Super, Service Any role granted the port_check right
checkrclink	Super, Service Any role granted the rclink_check right
checksnmp	Super, Service Any role granted the snmp_check right
checkvv	Super, Service, Edit Any role granted the vv_check right
compactcpg	Super, Edit Any role granted the cpg_compact right
compactld	Super, Edit Any role granted the ld_compact right
controlencryption	Super Any role granted the control_encryption right Super, Service (for status subcommand only) Any role granted the control_encryption_status right
controliscsiport	Super, Service Any role granted the port_control right
controlmag	Super, Service Any role granted the mag_control right
controlpd	Super, Service Any role granted the pd_control right
controlport	Super, Service Any role granted the port_control right

Table Continued

Command	Authority
controlrecoveryauth	Super Any role granted the recoveryauth_control right Super, Service (for status subcommand only) Any role granted the recoveryauth_status right
controlsecurity	Super Any role granted the security_control right Super, Service (for status option only) Any role granted the security_status_control right
controlsr	Super, Service
createcert	Super
createcpg	Super, Edit Any role granted the cpg_create right
createdomain	Super Any role granted the domain_create right
createdomainset	Super, Edit Any role granted the domainset_create right can create a domain set Any role granted the domainset_set right can add domains to a domain set
createfspn	Super, Edit Any role granted the fspn_create right.
creategroupsv	Super, Edit Any role granted the groupsv_create right
creategroupvvcopy	Super, Edit Any role granted the groupvvcopy_create right
createhost	Super, Edit Any role granted the host_create right

Table Continued

Command	Authority
createhostset	Super, Edit Any role granted the hostset_create right can create a host set Any role granted the hostset_set right can add hosts to a host set
createmaint	Super, Service Any role granted maint_set right
creatercopygroup	Super, Edit Any role granted the rcopygroup_create right
creatercopytarget	Super, Edit Any role granted the rcopytarget_create right
createsched	Super, Service, Edit Any role granted the sched_create right
createsnmpuser	Super Any role granted the snmpuser_create right
createspare	Super, Service Any role granted the spare_create right
createsv	Super, Edit Any role granted the sv_create right
createtemplate	Super Any role granted the template_create right
createuser	Super Any role granted the user_create right
createvlun	Super, Edit Any role granted the vlun_create right
createvv	Super, Edit Any role granted the vv_create right

Table Continued

Command	Authority
createvvcopy	Super, Edit Any role granted the vvcopy_create right
createvvset	Super, Edit Any role granted the vvset_create right can create a volume set Any role granted the vvset_set right can add volumes to a volume set
dismisspd	Super, Service Any role granted the pd_dismiss right
dismissrcopylink	Super, Edit Any role granted the rcopylink_dismiss right
dismissrcopytarget	Super, Edit Any role granted the rcopytarget_dismiss right
dismissrcopyvv	Super, Edit Any role granted the rcopyvv_dismiss right
freespace	Super, Edit Any role granted the space_free right
growvv	Super, Edit Any role granted the vv_grow right
importcert	Super
importvv	Super, Edit Any role granted the vv_import right
locatecage	Super, Service Any role granted the cage_locate right
locatenode	Super, Service Any role granted the node_locate right

Table Continued

Command	Authority
locatesys	Super, Service Any role granted the sys_locate right
movech	Super, Service, Edit Any role granted the ch_move right
movechtospare	Super, Service, Edit Any role granted the ch_movetospare right
movepd	Super, Service, Edit Any role granted the pd_move right
movepdtospare	Super, Service, Edit Any role granted the pd_movetospare right
moverelocpd	Super, Service Any role granted the pd_movereloc right
movetodomain	Super, Edit Any role granted the domain_moveto right
promotegroupsv	Super, Edit Any role granted the svgrp_promote right
promotesv	Super, Edit Any role granted the sv_promote right
promotevvcopy	Super, Edit Any role granted the vvcopy_promote right
removealert	Super, Service Any role granted the alert_remove right
removecert	Super
removecpg	Super, Edit Any role granted the cpg_remove right

Table Continued

Command	Authority
removedomain	Super Any role granted the domain_remove right
removedomainset	Super, Edit Any role granted the domainset_set right can remove a domain set or domains from a domain set
removefspn	Super, Edit Any role granted the fspn_delete right.
removehost	Super, Edit Any role granted the host_remove right
removehostset	Super, Edit Any role granted the hostset_set right can remove a host set or hosts from a host set
removeld	Super, Service, Edit Any role granted the ld_remove right
removercopygroup	Super, Edit Any role granted the rcopygroup_remove right
removercopytarget	Super, Edit Any role granted the rcopytarget_remove right
removesched	Super, Service, Edit Any role granted the sched_setany right can remove any user scheduled tasks Any role granted the sched_setown right can only remove their own scheduled tasks
removesnmpmgr	Super Any role granted the snmpmgr_remove right
removesnmppw	Super Any role granted the snmppw_remove right

Table Continued

Command	Authority
removesnmpuser	Super Any role granted the snmpuser_remove right
removespare	Super, Service Any role granted the spare_remove right
removespcredential	Super Any role granted the spcred_remove right
removesshkey	Super, Service, Edit, Browse Any role granted the sshkey_remove right
removetask	Super, Edit Any role granted the task_remove right
removetemplate	Super Any role granted the template_remove right
removeuser	Super Any role granted the user_remove right
removeuserconn	Super Any role granted the userconn_remove right
removevlun	Super, Edit Any role granted the vlun_remove right
removevv	Super, Edit Any role granted the vv_remove right can remove base volumes Any role granted the vvcopy_remove right can remove physical volumes Any role granted the sv_remove right can remove virtual copies
removevvset	Super, Edit Any role granted the vvset_set right can remove a volume set or volumes from a volume set

Table Continued

Command	Authority
removewsapisession	Super Any role granted the wsapisession_remove right
servicecage	Super, Service
servicehost	Super, Service Any role granted the host_service right
servicemag	Super, Service Any role granted the mag_service right
servicenode	Super, Service Any role granted the node_service right
setalert	Super, Service Any role granted the alert_set right
setauthparam	Super Any role granted the authparam_set right
setbattery	Super, Service Any role granted the battery_set right
setcage	Super, Service Any role granted the cage_set right
setcim	Super, Service Any role granted the cim_set right
setclienv	Super, Service, Edit, Browse Any role granted the clienv_set right
setcpg	Super, Service, Edit Any role granted the cpg_set right
setdate	Super, Service Any role granted the date_set right

Table Continued

Command	Authority
setdomain	Super Any role granted the domain_set right
setdomainset	Super, Edit Any role granted the domainset_annotate right can only annotate the comment of a domain set Any role granted the domainset_set right can set any domain set property
sethost	Super, Edit Any role granted the host_set right
sethostset	Super, Edit Any role granted the hostset_annotate right can only annotate the comment of a host set Any role granted the hostset_set right can set any host set property
setlicense	Super, Service Any role granted the license_set right
setnet	Super, Service Any role granted the net_set right
setnode	Super, Service Any role granted the node_set right
setpassword	Super, Service, Edit, Browse Any role granted the password_setany right can set any user password Any role granted the password_setown right can only set their own password Only users with Super role can set the minimum password length
setpd	Super, Service Any role granted the pd_set right
setqos	Super, Edit Any role granted the qos_set right can set QoS configurations

Table Continued

Command	Authority
setrcopygroup	Super, Edit Any role granted the rcopygroup_set right
setrcopytarget	Super, Edit Any role granted the rcopytarget_set right
setsched	Super, Service, Edit Any role granted the sched_setany right can set any user scheduled task properties Any role granted the sched_setown right can only set their own scheduled task properties
setsnmpmgr	Super Any role granted the snmpmgr_set right
setsnmppw	Super Any role granted the snmppw_set right
setsnmpuserpw	Super Any role granted the snmppw_set right.
setsshkey	Super, Service, Edit, Browse Any role granted the sshkey_set right
setstatch	Super, Edit Any role granted the statch_set right
setstatpdch	Super, Edit Any role granted the statpdch_set right
setsys	Super, Service Any role granted the sys_set right
setsysmgr	Super Any role granted the sysmgr_set right
settask	Super, Edit Any role granted the task_set right

Table Continued

Command	Authority
settemplate	Super Any role granted the template_set right
setuser	Super, Edit, Browse Any role granted the user_setany right can add, remove, and set default domains for any user Any role granted the user_setown right can only set their own default domain
setuseracl	Super, Edit Any role granted the useracl_set right
setvv	Super, Edit Any role granted the vv_set right can edit the properties of base volumes Any role granted the vvcopy_set right can edit the properties of physical copies of volumes Any role granted the sv_set right can edit the properties of virtual copies of volumes
setvvolsc	Super, Edit Any role granted the vvset_set right
setvvset	Super, Edit Any role granted the vvset_annotate right can only annotate the comment of a volume set Any role granted the vvset_set right can set any volume set property
setwsapi	Super, Service Any role granted the wsapi_set right
showencryption	Super, Service Any role granted the control_encryption_status right
showportpel	Save/reset Any role granted by the save_pel right. Everything else, any role in the system.

Table Continued

Command	Authority
showsched	Super, Service, Edit, Browse Any extended role in the system
shutdownnode	Super, Service Any role granted the `node_shutdown` right
shutdownsys	Super, Service Any role granted the `sys_shutdown` right
startcim	Super, Service Any role granted the `cim_start` right
startld	Super, Service Any role granted the `ld_start` right
startnoderescue	Super, Service
startrcopy	Super, Edit Any role granted the `rcopy_start` right
startrcopygroup	Super, Edit Any role granted the `rcopygroup_start` right
starttask	Super, Service, Edit Any role granted the `task_start` right
startvv	Super, Service Any role granted the `vv_start` right
startwsapi	Super, Service Any role granted the `wsapi_start` right
stopcim	Super, Service Any role granted the `cim_stop` right
stoprcopy	Super, Edit Any role granted the `rcopy_stop` right

Table Continued

Command	Authority
stoprcopygroup	Super, Edit Any role granted the rcopygroup_stop right
stopwsapi	Super, Service Any role granted the wsapi_stop right
syncrcopy	Super, Edit Any role granted the rcopy_sync right
tunepd	Super, Edit Any role granted the pd_tune right
tunesys	Super, Edit Any role granted the sys_tune right
tunevv	Super, Edit Any role granted the vv_tune right
updatesnapspace	Super, Edit Any role granted the snapspace_update right
updatevv	Super, Edit, Browse Any role granted the vv_update right
upgradecage	Super, Service Any role granted the cage_upgrade right
upgradepd	Super, Service Any role granted the pd_upgrade right

Websites

General websites

Single Point of Connectivity Knowledge (SPOCK) Storage compatibility matrix

https://www.hpe.com/storage/spock

Storage white papers and analyst reports

https://www.hpe.com/storage/whitepapers

For additional websites, see **Support and other resources**.

Support and other resources

Accessing Hewlett Packard Enterprise Support

- For live assistance, go to the Contact Hewlett Packard Enterprise Worldwide website:

  https://www.hpe.com/info/assistance

- To access documentation and support services, go to the Hewlett Packard Enterprise Support Center website:

  https://www.hpe.com/support/hpesc

Information to collect

- Technical support registration number (if applicable)
- Product name, model or version, and serial number
- Operating system name and version
- Firmware version
- Error messages
- Product-specific reports and logs
- Add-on products or components
- Third-party products or components

Accessing updates

- Some software products provide a mechanism for accessing software updates through the product interface. Review your product documentation to identify the recommended software update method.
- To download product updates:

  Hewlett Packard Enterprise Support Center

  https://www.hpe.com/support/hpesc

  Hewlett Packard Enterprise Support Center: Software downloads

  https://www.hpe.com/support/downloads

  My HPE Software Center

  https://www.hpe.com/software/hpesoftwarecenter

- To subscribe to eNewsletters and alerts:

  https://www.hpe.com/support/e-updates

- To view and update your entitlements, and to link your contracts and warranties with your profile, go to the Hewlett Packard Enterprise Support Center **More Information on Access to Support Materials** page:

  https://www.hpe.com/support/AccessToSupportMaterials

IMPORTANT: Access to some updates might require product entitlement when accessed through the Hewlett Packard Enterprise Support Center. You must have an HPE Passport set up with relevant entitlements.

Remote support

Remote support is available with supported devices as part of your warranty or contractual support agreement. It provides intelligent event diagnosis, and automatic, secure submission of hardware event notifications to Hewlett Packard Enterprise, which will initiate a fast and accurate resolution based on your product's service level. Hewlett Packard Enterprise strongly recommends that you register your device for remote support.

If your product includes additional remote support details, use search to locate that information.

Remote support and Proactive Care information

HPE Get Connected

https://www.hpe.com/services/getconnected

HPE Proactive Care services

https://www.hpe.com/services/proactivecare

HPE Datacenter Care services

https://www.hpe.com/services/datacentercare

HPE Proactive Care service: Supported products list

https://www.hpe.com/services/proactivecaresupportedproducts

HPE Proactive Care advanced service: Supported products list

https://www.hpe.com/services/proactivecareadvancedsupportedproducts

Proactive Care customer information

Proactive Care central

https://www.hpe.com/services/proactivecarecentral

Proactive Care service activation

https://www.hpe.com/services/proactivecarecentralgetstarted

Warranty information

To view the warranty information for your product, see the links provided below:

HPE ProLiant and IA-32 Servers and Options

https://www.hpe.com/support/ProLiantServers-Warranties

HPE Enterprise and Cloudline Servers

https://www.hpe.com/support/EnterpriseServers-Warranties

HPE Storage Products

https://www.hpe.com/support/Storage-Warranties

HPE Networking Products

https://www.hpe.com/support/Networking-Warranties

Regulatory information

To view the regulatory information for your product, view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products*, available at the Hewlett Packard Enterprise Support Center:

https://www.hpe.com/support/Safety-Compliance-EnterpriseProducts

Additional regulatory information

Hewlett Packard Enterprise is committed to providing our customers with information about the chemical substances in our products as needed to comply with legal requirements such as REACH (Regulation EC No 1907/2006 of the European Parliament and the Council). A chemical information report for this product can be found at:

https://www.hpe.com/info/reach

For Hewlett Packard Enterprise product environmental and safety information and compliance data, including RoHS and REACH, see:

https://www.hpe.com/info/ecodata

For Hewlett Packard Enterprise environmental information, including company programs, product recycling, and energy efficiency, see:

https://www.hpe.com/info/environment

Documentation feedback

Hewlett Packard Enterprise is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback (**docsfeedback@hpe.com**). When submitting your feedback, include the document title, part number, edition, and publication date located on the front cover of the document. For online help content, include the product name, product version, help edition, and publication date located on the legal notices page.