



SGI InfiniteStorage 4000 Series and 5000 Series Command Line Interface and Script Commands

(ISSM 10.86)

The information in this document supports the SGI InfiniteStorage 4000 series and 5000 series storage systems (ISSM 10.86). Refer to the table below to match your specific SGI InfiniteStorage product with the model numbers used in this document.

SGI Model #	NetApp Model
TP9600H	6091
TP9700F	6091
IS4500F	6091
TP9600F	3994 and 3992
IS4000H	3994
IS350	3992
IS220	1932 1333 DE1300
IS4100	4900
IS-DMODULE16-Z	FC4600
IS-DMODULE60	DE6900
IS4600	7091
IS-DMODULE12 & IS2212 (JBOD)	DE1600
IS-DMODULE24 & IS2224 (JBOD)	DE5600
IS-DMODULE60-SAS	DE6600
IS5012	E2600
IS5024	E2600
IS5060	E2600
IS5512	E5400
IS5524	E5400
IS5560	E5400
IS5600	E5500

Copyright information

Copyright © 1994–2012 NetApp, Inc. All rights reserved. Printed in the U.S.A.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987)0

Trademark information

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, FAServer, FastStak, FilerView, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the Web at www.ibm.com/legal/copytrade.shtml.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the U.S.A. and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the U.S.A. and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the U.S.A. and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

Table of Contents

Chapter 1	Formatting the Commands	1
	Structure of a CLI Command	1
	Interactive Mode	2
	CLI Command Wrapper Syntax	2
	Command Line Terminals	4
	Alert Severities Commands	8
	AutoSupport Bundle Collection Commands	9
	Structure of a Script Command	11
	Synopsis of the Script Commands	12
	Recurring Syntax Elements	14
	Naming Conventions	22
	Formatting CLI Commands	23
	Formatting Rules for Script Commands	24
	Usage Guidelines	26
	Detailed Error Reporting	26
	Exit Status	27
	Adding Comments to a Script File	28
	Firmware Compatibility Levels	29
Chapter 2	Script Commands	31
	Command Reference - Listed by Function	32
	Asynchronous Mirroring Commands	32
	AutoSupport Bundle Collection Commands	32
	Consistency Group Commands	33
	Controller Commands	33
	Core Dump Commands	34
	Database Commands	34
	Disk Pool Commands	35
	Drive Commands	35

Feature Management Commands	35
Host Topology Commands	36
iSCSI Commands	36
Core Dump Commands	37
Session Command	37
Snapshot (Legacy) Commands	37
Snapshot Group Commands	37
Snapshot Image Commands	37
Snapshot Volume Commands	37
SSD Cache Commands	38
Storage Array Commands	38
Support Bundle Collection Commands	40
Synchronous Mirroring Commands	40
Tray Commands	41
Uncategorized Commands	41
Volume Commands	41
Volume Commands for Disk Pools	42
Volume Copy Commands	42
Volume Group Commands	42
Command Reference – Listed Alphabetically	44
Activate Asynchronous Mirroring	44
Activate Host Port	44
Activate iSCSI Initiator	44
Activate Storage Array Firmware	45
Activate Synchronous Mirroring	45
Add Drives to SSD Cache	50
Add Member to Consistency Group	50
Add Volume to Asynchronous Mirror Group	53
Autoconfigure Storage Array	55

Autoconfigure Storage Array Hot Spares	60
Cancel Asynchronous Mirror Group Role Reversal	60
Change SSD Cache Application Type	61
Check Asynchronous Mirror Group Consistency	62
Check Repository Consistency	62
Check Storage Array Connectivity	64
Check Synchronous Mirroring Status	65
Check Volume Parity	66
Clear Asynchronous Mirroring Fault	67
Clear Drive Channel Statistics	69
Clear Storage Array Configuration	69
Clear Storage Array Core Dump	70
Clear Storage Array Event Log	71
Clear Storage Array Firmware Pending Area	71
Clear Storage Array Recovery Mode	71
Clear Volume Reservations	72
Clear Volume Unreadable Sectors	73
Configure Automatic Support Bundle Collection	73
Convert Snapshot (Legacy) Volume to Snapshot Group	74
Create Asynchronous Mirror Group	75
Create Consistency Group	78
Create Consistency Group Snapshot Image	83
Create Consistency Group Snapshot Volume	84
Create Consistency Group Snapshot Volume Mapping	88
Create Disk Pool	89
Create Host	94
Create Host Group	95
Create Host Port	96
Create iSCSI Initiator	97

Create RAID Volume (Automatic Drive Select)	98
Create RAID Volume (Free Extent Based Select)	104
Create RAID Volume (Manual Drive Select).	107
Create Read-Only Snapshot Volume	112
Create Snapshot (Legacy) Volume	114
Create Snapshot Group	123
Create Snapshot Image	130
Create Snapshot Volume.	130
Create SSD Cache	133
Create Storage Array Security Key.	135
Create Synchronous Mirroring	137
Create Volume in Disk Pool	139
Create Volume Copy.	143
Create Volume Group	146
Deactivate Asynchronous Mirroring.	151
Deactivate Synchronous Mirroring.	151
Delete Asynchronous Mirror Group	151
Delete Consistency Group.	152
Delete Consistency Group Snapshot Image	153
Delete Consistency Group Snapshot Volume.	154
Delete Disk Pool	155
Delete Host	155
Delete Host Group.	156
Delete Host Port	157
Delete iSCSI Initiator	157
Delete Snapshot (Legacy) Volume	157
Delete Snapshot Group	158
Delete Snapshot Image	159
Delete Snapshot Volume.	161

Delete SSD Cache	161
Delete Volume.	162
Delete Volume from Disk Pool.	163
Delete Volume Group	165
Diagnose Controller	165
Diagnose Synchronous Mirroring	167
Disable AutoSupport at the EMW Level SMcli Version	167
Disable External Security Key Management	168
Disable Snapshot (Legacy)	168
Disable Storage Array Feature	170
Display Automatic Support Bundle Collection Configuration	172
Display Automatic Support Bundle Collection Schedule	172
Download Drive Firmware	172
Download Environmental Card Firmware	174
Download Storage Array Drive Firmware	175
Download Storage Array Firmware/NVSRAM	176
Download Storage Array NVSRAM.	177
Download Tray Configuration Settings	177
Enable AutoSupport at the EMW Level SMcli Version.	178
Enable Controller Data Transfer	178
Enable Disk Pool Security.	179
Enable External Security Key Management	179
Enable or Disable SSD Cache for a Volume.	180
Enable Storage Array Feature	181
Enable Volume Group Security	184
Establish Asynchronous Mirrored Pair	184
Export Storage Array Security Key	185
Import Storage Array Security Key	186
Increase Capacity of Volume in Disk Pool	187

Initialize Thin Volume	189
Load Storage Array DBM Database	191
Locate SSD Cache	193
Recopy Volume Copy	193
Recover RAID Volume	194
Re-create External Security Key	198
Re-create Snapshot (Legacy)	199
Re-create Synchronous Mirroring Repository Volume	201
Remove Drives from SSD Cache	204
Remove Incomplete Asynchronous Mirrored Pair from Asynchronous Mirror Group	204
Remove Member Volume from Consistency Group	205
Remove Synchronous Mirroring	206
Remove Volume Copy	206
Remove Volume from Asynchronous Mirror Group	207
Remove Volume LUN Mapping	208
Rename Snapshot Volume	209
Rename SSD Cache	209
Repair Volume Parity	210
Replace Drive	211
Reset Asynchronous Mirror Group Statistics	212
Reset Controller	213
Reset iSCSI IP Address	214
Reset Storage Array Battery Install Date	214
Reset Storage Array Diagnostic Data	215
Reset Storage Array Infiniband Statistics Baseline	216
Reset Storage Array iSCSI Baseline	216
Reset Storage Array RLS Baseline	217
Reset Storage Array SAS PHY Baseline	217
Reset Storage Array SOC Baseline	217

Reset Storage Array Volume Distribution	218
Resume Asynchronous Mirror Group	218
Resume Consistency Group Snapshot Volume	219
Resume Snapshot (Legacy) Rollback	220
Resume Snapshot Image Rollback	220
Resume Snapshot Volume	221
Resume SSD Cache	223
Resume Synchronous Mirroring	223
Revive Drive	224
Revive Snapshot Group	225
Revive Snapshot Volume	225
Revive Volume Group	226
Save Asynchronous Mirror Group Statistics	226
Save Controller NVSRAM	229
Save Drive Channel Fault Isolation Diagnostic Status	229
Save Drive Log	230
Save Storage Array Configuration	231
Save Storage Array Core Dump	232
Save Storage Array DBM Database	233
Save Storage Array DBM Validator Information File	234
Save Storage Array Diagnostic Data	236
Save Storage Array Events	237
Save Storage Array Firmware Inventory	238
Save Storage Array InfiniBand Statistics	239
Save Storage Array iSCSI Statistics	240
Save Storage Array Performance Statistics	240
Save Storage Array RLS Counts	241
Save Storage Array SAS PHY Counts	242
Save Storage Array SOC Counts	242

Save Storage Array State Capture	243
Save Storage Array Support Data	243
Save Tray Log	244
Schedule Automatic Support Bundle Collection Configuration	245
Set Alert Severities	247
Set Asynchronous Mirror Group	248
Set Consistency Group Attributes	252
Set Consistency Group Snapshot Volume	254
Set Controller	255
Set Controller Service Action Allowed Indicator	259
Set Disk Pool	260
Set Disk Pool (Modify Disk Pool)	263
Set Drawer Service Action Allowed Indicator	264
Set Drive Channel Status	266
Set Drive Hot Spare	266
Set Drive Service Action Allowed Indicator	267
Set Drive State	268
Set Foreign Drive to Native	269
Set Host	270
Set Host Channel	271
Set Host Group	272
Set Host Port	272
Set iSCSI Initiator	273
Set iSCSI Target Properties	274
Set Read-Only Snapshot Volume to Read/Write Volume	275
Set Session	277
Set Snapshot (Legacy) Volume	279
Set Snapshot Group Attributes	283
Set Snapshot Group Media Scan	284

Set Snapshot Group Repository Volume Capacity	285
Set Snapshot Group Schedule	287
Set Snapshot Volume Media Scan	291
Set Snapshot Volume Repository Volume Capacity	291
Set Storage Array	293
Set Storage Array AutoSupport Bundle Disable	297
Set Storage Array AutoSupport Bundle Enable	297
Set Storage Array ICMP Response	298
Set Storage Array iSNS Server IPv4 Address	298
Set Storage Array iSNS Server IPv6 Address	299
Set Storage Array iSNS Server Listening Port	300
Set Storage Array iSNS Server Refresh	300
Set Storage Array Learn Cycle	301
Set Storage Array Redundancy Mode	302
Set Storage Array Security Key	302
Set Storage Array Time	303
Set Storage Array Tray Positions	303
Set Storage Array Unnamed Discovery Session	304
Set Synchronous Mirroring	305
Set Thin Volume Attributes	307
Set Tray Alarm	309
Set Tray Identification	309
Set Tray Service Action Allowed Indicator	310
Set Volume Attributes for a Volume in a Volume Group	312
Set Volume Attributes for a Disk Pool	319
Set Volume Copy	324
Set Volume Group	325
Set Volume Group Forced State	327
Set Volume Mapping	328

Show Alert Severities	330
Show Asynchronous Mirror Group	330
Show Asynchronous Mirror Group Synchronization Progress	331
Show Cache Backup Device Diagnostic Status	332
Show Cache Memory Diagnostic Status	332
Show Consistency Group	333
Show Consistency Group Snapshot Image	334
Show Controller	335
Show Controller Diagnostic Status	338
Show Controller NVSRAM	339
Show Current iSCSI Sessions	339
Show Disk Pool	340
Show Drive	341
Show Drive Channel Statistics	343
Show Drive Download Progress	344
Show Drive Performance Statistics	345
Show Host Interface Card Diagnostic Status	346
Show Host Ports	347
Show Snapshot Group	347
Show Snapshot Image	349
Show Snapshot Volumes	351
Show SSD Cache	352
Show SSD Cache Statistics	353
Show Storage Array	356
Show Storage Array Auto Configure	361
Show Storage Array Core Dump	365
Show Storage Array DBM Database	365
Show Storage Array Host Topology	366
Show Storage Array LUN Mappings	367

Show Storage Array Negotiation Defaults	368
Show Storage Array Unconfigured iSCSI Initiators	368
Show Storage Array Unreadable Sectors	368
Show String	369
Show Synchronous Mirroring Volume Candidates	369
Show Synchronous Mirroring Volume Synchronization Progress	370
Show Thin Volume	371
Show Volume	372
Show Volume Action Progress	374
Show Volume Copy	375
Show Volume Copy Source Candidates	375
Show Volume Copy Target Candidates	376
Show Volume Group	376
Show Volume Group Export Dependencies	377
Show Volume Group Import Dependencies	378
Show Volume Performance Statistics	378
Show Volume Reservations	380
Start Asynchronous Mirroring Synchronization	381
Start Cache Backup Device Diagnostic	382
Start Cache Memory Diagnostic	384
Start Consistency Group Snapshot Rollback	386
Start Controller Diagnostic	387
Start Controller Trace	389
Start Disk Pool Locate	391
Start Drive Channel Fault Isolation Diagnostics	392
Start Drive Channel Locate	393
Start Drive Initialize	394
Start Drive Locate	394
Start Drive Reconstruction	395

Start Host Interface Card Diagnostic	396
Start iSCSI DHCP Refresh	399
Start Secure Drive Erase	400
Start Snapshot (Legacy) Rollback.	400
Start Snapshot Image Rollback	401
Start SSD Cache Performance Modeling	402
Start Storage Array Configuration Database Diagnostic	403
Start Storage Array Core Dump	406
Start Storage Array iSNS Server Refresh	406
Start Storage Array Locate	407
Start Synchronous Mirroring Synchronization	407
Start Tray Locate.	407
Start Volume Group Defragment	408
Start Volume Group Export	408
Start Volume Group Import	409
Start Volume Group Locate	410
Start Volume Initialization	410
Stop Cache Backup Device Diagnostic	411
Stop Cache Memory Diagnostic	411
Stop Consistency Group Snapshot Rollback.	411
Stop Consistency Group Snapshot Volume	413
Stop Controller Diagnostic	413
Stop Disk Pool Locate.	414
Stop Drive Channel Fault Isolation Diagnostics.	414
Stop Drive Channel Locate.	415
Stop Drive Locate	415
Stop Host Interface Card Diagnostic	415
Stop Pending Snapshot Images on Consistency Group	416
Stop Snapshot Group Pending Snapshot Images	416

Stop Snapshot (Legacy)	417
Stop Snapshot (Legacy) Rollback	418
Stop Snapshot Image Rollback	418
Stop Snapshot Volume	419
Stop SSD Cache Performance Modeling	420
Stop Storage Array Configuration Database Diagnostic	422
Stop Storage Array Drive Firmware Download	422
Stop Storage Array iSCSI Session	423
Stop Storage Array Locate	423
Stop Tray Locate	423
Stop Volume Copy	424
Stop Volume Group Locate.	424
Suspend Asynchronous Mirror Group	424
Suspend SSD Cache	425
Suspend Synchronous Mirroring.	426
Test Alert Severities	427
Test Asynchronous Mirror Group Connectivity	427
Validate Storage Array Security Key	428

Chapter 3	Deprecated Commands and Parameters.	431
	Deprecated Commands.	431
	Deprecated Parameters	431

The command line interface (CLI) is a software application that provides a way to configure and monitor storage arrays. Using the CLI, you can run commands from an operating system prompt, such as the DOS C : prompt, a Linux operating system path, or a Solaris operating system path.

The CLI gives you direct access to a script engine that is a utility in the SANtricity ES Storage Manager software (also referred to as the *storage management software*). The script engine runs commands that configure and manage the storage arrays. The script engine reads the commands, or runs a script file, from the command line and performs the operations instructed by the commands.

The script commands configure and manage a storage array. The script commands are distinct from the CLI commands. You can enter individual script commands, or you can run a file of script commands. When you enter an individual script command, you embed the script command in a CLI command. When you run a file of script commands, you embed the file name in the CLI command.

Structure of a CLI Command

The CLI commands are in the form of a command wrapper and elements embedded into the wrapper. A CLI command consists of these elements:

- A command wrapper identified by the term `SMcli`
- The storage array identifier
- Terminals that define the operation to be performed
- Script commands

The CLI command wrapper is a shell that identifies storage array controllers, embeds operational terminals, embeds script commands, and passes these values to the script engine.

All CLI commands have the following structure:

```
SMcli storageArray terminal script-commands;
```

- `SMcli` invokes the command line interface.
- `storageArray` is the name or the IP address of the storage array.
- `terminal` are CLI values that define the environment and the purpose for the command.
- `script-commands` are one or more script commands or the name of a script file that contains script commands. (The script commands configure and manage the storage array.)

If you enter an incomplete or inaccurate `SMcli` string that does not have the correct syntax, parameter names, options, or terminals, the script engine returns usage information.

Interactive Mode

If you enter `SMcli` and a storage array name but do not specify CLI parameters, script commands, or a script file, the command line interface runs in interactive mode. Interactive mode lets you run individual commands without prefixing the commands with `SMcli`.

In interactive mode, you can enter a single command, view the results, and enter the next command without typing the complete `SMcli` string. Interactive mode is useful for determining configuration errors and quickly testing configuration changes.

To end an interactive mode session, type the operating system-specific command for terminating a program, such as **Control-C** on the UNIX operating system or the Windows operating system. Typing the termination command (**Control-C**) while in interactive mode turns off interactive mode and returns operation of the command prompt to an input mode that requires you to type the complete `SMcli` string.

CLI Command Wrapper Syntax

General syntax forms of the CLI command wrappers are listed in this section. The general syntax forms show the terminals and the parameters that are used in each command wrapper. The conventions used in the CLI command wrapper syntax are listed in the following table.

Convention	Definition
a b	Alternative ("a" or "b")
<i>italicized-words</i>	A terminal that needs user input to fulfill a parameter (a response to a variable)
[...] (square brackets)	Zero or one occurrence (square brackets are also used as a delimiter for some command parameters)
{ ... } (curly braces)	Zero or more occurrences
(a b c)	Choose only one of the alternatives
bold	A terminal that needs a command parameter entered to start an action

```
SMcli host-name-or-IP-address
[host-name-or-IP-address]
[-c "command; {command2};"]
[-n storage-system-name | -w wwID]
[-o outputfile] [-p password] [-R (admin | monitor)]
[-e] [-S] [-quick]
```

```
SMcli host-name-or-IP-address [hostname-or-IP-address]
[-f scriptfile]
[-n storage-system-name | -w wwID]
[-o outputfile] [-p password] [-R (admin | monitor)]
[-e] [-S] [-quick]
```

```

SMcli (-n storage-system-name | -w wwID)
[-c "command; {command2};"]
[-o outputfile] [-p password] [-R (admin | monitor)]
[-e] [-S] [-quick]

SMcli (-n storage-system-name -w wwID)
[-f scriptfile]
[-o outputfile] [-R (admin | monitor)] [-p password]
[-e] [-S] [-quick]

SMcli -a email: email-address
[host-name-or-IP-address1
[host-name-or-IP-address2]]
[-n storage-system-name | -w wwID | -h host-name]
[-I information-to-include] [-q frequency] [-S]

SMcli -x email: email-address
[host-name-or-IP-address1
[host-name-or-IP-address2]]
[-n storage-system-name | -w wwID | -h host-name] [-S]

SMcli (-a | -x) trap: community,
host-name-or-IP-address
[host-name-or-IP-address1 [host-name-or-IP-address2]]
[-n storage-system-name | -w wwID | -h host-name] [-S]

SMcli -d [-w] [-i] [-s] [-v] [-S]

SMcli -m host-name-or-IP-address -F email-address
[-g contactInfoFile] [-S]

SMcli -A [host-name-or-IP-address
[host-name-or-IP-address]]
[-S]

SMcli -X (-n storage-system-name | -w wwID | -h
host-name)

SMcli -?

```

Command Line Terminals

Terminal	Definition
<i>host-name-or-IP-address</i>	<p>Specifies either the host name or the Internet Protocol (IP) address (<i>xxx.xxx.xxx.xxx</i>) of an in-band managed storage array or an out-of-band managed storage array.</p> <ul style="list-style-type: none"> ■ If you are managing a storage array by using a host through in-band storage management, you must use the <i>-n</i> terminal or the <i>-w</i> terminal if more than one storage array is connected to the host. ■ If you are managing a storage array by using out-of-band storage management through the Ethernet connection on each controller, you must specify the <i>host-name-or-IP-address</i> of the controllers. ■ If you have previously configured a storage array in the Enterprise Management Window, you can specify the storage array by its user-supplied name by using the <i>-n</i> terminal. ■ If you have previously configured a storage array in the Enterprise Management Window, you can specify the storage array by its World Wide Identifier (WWID) by using the <i>-w</i> terminal.
<i>-A</i>	<p>Adds a storage array to the configuration file. If you do not follow the <i>-A</i> terminal with a <i>host-name-or-IP-address</i>, auto-discovery scans the local subnet for storage arrays.</p>
<i>-a</i>	<p>Adds a Simple Network Management Protocol (SNMP) trap destination or an email address alert destination.</p> <ul style="list-style-type: none"> ■ When you add an SNMP trap destination, the SNMP community is automatically defined as the community name for the trap, and the <i>host</i> is the IP address or Domain Name Server (DNS) host name of the system to which the trap should be sent. ■ When you add an email address for an alert destination, the <i>email-address</i> is the email address to which you want the alert message to be sent.
<i>-c</i>	<p>Indicates that you are entering one or more script commands to run on the specified storage array. End each command with a semicolon (<i>;</i>). You cannot place more than one <i>-c</i> terminal on the same command line. You can include more than one script command after the <i>-c</i> terminal.</p>
<i>-d</i>	<p>Shows the contents of the script configuration file. The file content has this format:</p> <pre><i>storage-system-name host-name1 host-name2</i></pre>
<i>-e</i>	<p>Runs the commands without performing a syntax check first.</p>
<i>-F</i> (uppercase)	<p>Specifies the email address from which all alerts will be sent.</p>

Terminal	Definition
-f (lowercase)	<p>Specifies a file name that contains script commands that you want to run on the specified storage array. The -f terminal is similar to the -c terminal in that both terminals are intended for running script commands. The -c terminal runs individual script commands. The -f terminal runs a file of script commands.</p> <p>By default, any errors that are encountered when running the script commands in a file are ignored, and the file continues to run. To override this behavior, use the <code>set session errorAction=stop</code> command in the script file.</p>
-g	<p>Specifies an ASCII file that contains email sender contact information that will be included in all email alert notifications. The CLI assumes that the ASCII file is text only, without delimiters or any expected format. Do not use the -g terminal if a <code>userdata.txt</code> file exists.</p>
-h	<p>Specifies the host name that is running the SNMP agent to which the storage array is connected. Use the -h terminal with these terminals:</p> <ul style="list-style-type: none"> ■ -a ■ -x
-I (uppercase)	<p>Specifies the type of information to be included in the email alert notifications. You can select these values:</p> <ul style="list-style-type: none"> ■ <code>eventOnly</code> – Only the event information is included in the email. ■ <code>profile</code> – The event and array profile information is included in the email. <p>You can specify the frequency for the email deliveries using the -q terminal.</p>
-i (lowercase)	<p>Shows the IP address of the known storage arrays. Use the -i terminal with the -d terminal. The file contents has this format:</p> <pre>storage-system-name IP-address1 IPaddress2</pre>
-m	<p>Specifies the host name or the IP address of the email server from which email alert notifications are sent.</p>
-n	<p>Specifies the name of the storage array on which you want to run the script commands. This name is optional when you use a <code>host-name-or-IP-address</code>. If you are using the in-band method for managing the storage array, you must use the -n terminal if more than one storage array is connected to the host at the specified address. The storage array name is required when the <code>host-name-or-IP-address</code> is not used. The name of the storage array that is configured for use in the Enterprise Management Window (that is, the name is listed in the configuration file) must not be a duplicate name of any other configured storage array.</p>

Terminal	Definition
-o	<p>Specifies a file name for all output text that is a result of running the script commands. Use the -o terminal with these terminals:</p> <ul style="list-style-type: none"> ■ -c ■ -f <p>If you do not specify an output file, the output text goes to standard output (stdout). All output from commands that are not script commands is sent to stdout, regardless of whether this terminal is set.</p>
-p	<p>Defines the password for the storage array on which you want to run commands. A password is not necessary if a password has not been set on the storage array.</p> <p>If you set a monitor password for the storage array, the use of the -p parameter is mandatory. Users cannot run any of the non -destructive commands such as the show commands.</p>
-q	<p>Specifies the frequency that you want to receive event notifications and the type of information returned in the event notifications. An email alert notification containing at least the basic event information is always generated for every critical event.</p> <p>These values are valid for the -q terminal:</p> <ul style="list-style-type: none"> ■ everyEvent – Information is returned with every email alert notification. ■ 2 – Information is returned no more than once every two hours. ■ 4 – Information is returned no more than once every four hours. ■ 8 – Information is returned no more than once every eight hours. ■ 12 – Information is returned no more than once every 12 hours. ■ 24 – Information is returned no more than once every 24 hours. <p>Using the -I terminal you can specify the type of information in the email alert notifications.</p> <ul style="list-style-type: none"> ■ If you set the -I terminal to eventOnly, the only valid value for the -q terminal is everyEvent. ■ If you set the -I terminal to either the profile value or the supportBundle value, this information is included with the emails with the frequency specified by the -q terminal.

Terminal	Definition
-quick	<p>Reduces the amount of time that is required to run a single-line operation. An example of a single-line operation is the <code>recreate snapshot volume</code> command. This terminal reduces time by not running background processes for the duration of the command.</p> <p>Do not use this terminal for operations that involve more than one single-line operation. Extensive use of this command can overrun the controller with more commands than the controller can process, which causes operational failure. Also, status updates and configuration updates that are collected usually from background processes will not be available to the CLI. This terminal causes operations that depend on background information to fail.</p>
-R (uppercase)	<p>Defines the user role for the password. The roles can be either:</p> <ul style="list-style-type: none"> ■ <code>admin</code> – The user has privilege to change the storage array configuration. ■ <code>monitor</code> – The user has privilege to view the storage array configuration, but cannot make changes. <p>The <code>-R</code> parameter is valid only when used with the <code>-p</code> parameter, which specifies that you define a password for a storage array.</p> <p>The <code>-R</code> parameter is required only if the dual password feature is enabled on the storage array. The <code>-R</code> parameter is not necessary under these conditions:</p> <ul style="list-style-type: none"> ■ The dual password feature is not enabled on the storage array. ■ Only one admin role is set and the monitor role is not set for the storage array.
-S (uppercase)	<p>Suppresses informational messages describing the command progress that appear when you run script commands. (Suppressing informational messages is also called silent mode.) This terminal suppresses these messages:</p> <ul style="list-style-type: none"> ■ <code>Performing syntax check</code> ■ <code>Syntax check complete</code> ■ <code>Executing script</code> ■ <code>Script execution complete</code> ■ <code>SMcli completed successfully</code>
-s (lowercase)	<p>Shows the alert settings in the configuration file when used with the <code>-d</code> terminal.</p>
-v	<p>Shows the current global status of the known devices in a configuration file when used with the <code>-d</code> terminal.</p>

Terminal	Definition
-w	Specifies the WWID of the storage array. This terminal is an alternate to the -n terminal. Use the -w terminal with the -d terminal to show the WWIDs of the known storage arrays. The file content has this format: <i>storage-system-name world-wide-ID IP-address1 IP-address2</i>
-X (uppercase)	Deletes a storage array from a configuration.
-x (lowercase)	Removes an SNMP trap destination or an email address alert destination. The <i>community</i> is the SNMP community name for the trap, and the <i>host</i> is the IP address or DNS host name of the system to which you want the trap sent.
-?	Shows usage information about the CLI commands.

Alert Severities Commands

The CLI provides special commands that enable you to set alert severities, and to send out a test alert to the Windows Event Log and all configured syslog receivers. The alert severities apply to all of the storage arrays in the entire storage system. The commands are SMcli commands that run only from a command line.

Setting Alert Severities

```
SMcli -alertSeverities (severity |
[severity1, ... severityN])
```

The alert severities values that you can set are the following:

- `critical` – Alerts will be sent
- `warning` – Alerts will be sent
- `informational` – Alerts will not be sent
- `debug` – Alerts will be sent

NOTE The debug value is for Technical Support only. Do not attempt to use this value.

You can set one or more alert severities values. If you set more than one alert severities value, enclose all of the values in square brackets ([]) and separate the values by a comma.

Showing Alert Severities

```
SMcli -alertSeverities
```

This command shows all of the severities for which an alert is sent. This command cannot show information for a specific type of severity.

Sending a Test Alert

```
SMcli -alertTest
```

This command sends out a test alert to the Windows Event Log and all configured syslog receivers.

AutoSupport Bundle Collection Commands

AutoSupport (ASUP) is a feature that enables storage arrays to automatically collect support data into a customer support bundle and send the data to Technical Support. Technical Support can then perform remote troubleshooting and problem analysis with the storage management software. ASUP collects support data to report configuration, subsystem status, and exceptions in near-real time. ASUP messages typically include a collection of system logs files, configuration data (formatted XML and unstructured command output), state data (subsystem up/down, capacity used), performance metrics, and system inventory data. All of the data gathered is collected into a single compressed archive file format (7z).

With the implementation of ASUP, users have two possible methods for collecting support data in a storage array:

- ASUP collection

Data is automatically collected and sent to Technical Support.

- Legacy support bundle collection

Collection of legacy support bundle data is configured by the user at intervals scheduled by the user. Users can then manually send the support bundles to Technical Support.

ASUP operations and legacy support bundle operations are mutually exclusive on a given storage array. When you turn on ASUP you automatically disable legacy support bundle collection. If you want to run legacy support bundle collection, you must turn off ASUP.

In the CLI, ASUP is a nonconfigurable, set it and forget it feature. Using the CLI commands, you can only turn on or turn off ASUP. Once turned on, ASUP automatically reports configuration, subsystem status, and exceptions in near-real time. Because ASUP speeds up troubleshooting and problem analysis, ASUP is the preferred data collection method to use if available on the storage array.

ASUP Messages

ASUP provides these types of messages:

- Event:
 - Sent when a support event occurs on the managed storage array.
 - Includes system configuration and diagnostic information.
 - Includes minimal extent of system configuration information.

- Daily:
 - Sent at midnight, local time of the host.
 - Provides a current set of system event logs and performance data.
 - Places less burden on payload and transmission on the messages originating from Event ASUP messages.
- Weekly:
 - Sent once every week at times that do not impact storage array operations.
 - Includes configuration and system state information.

The storage management software automatically assigns the schedule for each storage array it has discovered.

The storage array uses the internet to send ASUP messages to the ASUP backend. The ASUP backend provides near-real time access to the messages by Technical Support. ASUP requires compliance to the following transport protocol-specific requirements:

- HTTP or HTTPS upload:
- SMTP notifications

ASUP Commands

The CLI ASUP commands in the following table turn on or turn off the ASUP feature for either all of the storage arrays managed at the Enterprise Management Window (EMW) level or for a specific storage array.

<code>SMcli enable autoSupportFeature</code>	Turns on the ASUP feature at the EMW level
<code>SMcli disable autoSupportFeature</code>	Turns off the ASUP feature at the EMW level
<code>set storageArray autoSupportFeature enable</code>	Turns on the ASUP feature for a specific storage array
<code>set storageArray autoSupportFeature disable</code>	Turns off the ASUP feature for a specific storage array

The two "SMcli" commands run at the EMW level. All of the storage arrays being managed that are ASUP capable are can be enabled or disabled using the commands. As shown in the table, these are the complete commands.

The two "set" commands are script commands that you can be use to turn on or turn off ASUP for individual storage arrays. You can run these commands from the script editor in the storage management software GUI, a script file, or from the command line if you use a CLI wrapper as shown in the following example:

```
c:\...\smX\client>smcli 123.45.67.88 123.45.67.89
-c "set storageArray autoSupportFeature enable;"
```

Structure of a Script Command

ASUP Log

The ASUP log file has a detailed list of events encountered during delivery of the ASUP messages. The ASUP log provides information about status, history of transmission activity, and any errors encountered during delivery of the ASUP messages. The log file is available for all ASUP-enabled storage arrays.

The archived log filename is `ASUPMessages.n`, where `n` is an integer from 1 to 5. The log file is located in the `ASUPLog` directory. As the current log file reaches a size limit of 200 KB, the current log file is archived and a new log file is created.

All script commands have the following structure:

```
command operand-data (statement-data)
```

- *command* identifies the action to be performed.
- *operand-data* represents the objects associated with a storage array that you want to configure or manage.
- *statement-data* provides the information needed to perform the command.

The syntax for *operand-data* has the following structure:

```
(object-type | all object-types | [qualifier]
(object-type [identifier] (object-type [identifier] |
object-types [identifier-list]))
```

An object can be identified in four ways:

- Object type – Use when the command is not referencing a specific object.
- `all` parameter prefix – Use when the command is referencing all of the objects of the specified type in the storage array (for example, `allVolumes`).
- Square brackets – Use when performing a command on a specific object to identify the object (for example, `volume [engineering]`).
- A list of identifiers – Use to specify a subset of objects. Enclose the object identifiers in square brackets (for example, `volumes [sales engineering marketing]`).

A qualifier is required if you want to include additional information to describe the objects.

The object type and the identifiers that are associated with each object type are listed in this table.

Table 1 Script Command Object Type Identifiers

Object Type	Identifier
controller	a or b
drive	Tray ID and slot ID

Object Type	Identifier
replacementDrive	Tray ID and slot ID
driveChannel	Drive channel identifier
host	User label
hostChannel	Host channel identifier
hostGroup	User label
hostPort	User label
iscsiInitiator	User label or iSCSI Qualified Name (IQN)
iscsiTarget	User label or IQN
snapshot (legacy)	Volume user label
storageArray	Not applicable
tray	Tray ID
volume	Volume user label or volume World Wide Identifier (WWID) (<i>set</i> command only)
volumeCopy	Target volume user label and, optionally, the source volume user label
volumeGroup	User label Valid characters are alphanumeric, a hyphen, and an underscore.

Statement data is in the form of:

- Parameter = value (such as `raidLevel=5`)
- Parameter-name (such as `batteryInstallDate`)
- Operation-name (such as `redundancyCheck`)

A user-defined entry (such as user label) is called a variable. In the syntax, it is shown in *italic* (such as *trayID* or *volumeGroupName*).

Synopsis of the Script Commands

Because you can use the script commands to define and manage the different aspects of a storage array (such as host topology, drive configuration, controller configuration, volume definitions, and volume group definitions), the actual number of commands is extensive. The commands, however, fall into general categories that are reused when you apply the commands to configure or maintain a storage array. The following table lists the general form of the script commands and a definition of each command.

Table 2 General Form of the Script Commands

Syntax	Description
<code>accept object</code> { <i>statement-data</i> }	Perform the pending operation.
<code>activate object</code> { <i>statement-data</i> }	Sets up the environment so that an operation can take place or performs the operation if the environment is already set up correctly.
<code>autoConfigure storageArray</code> { <i>statement-data</i> }	Automatically creates a configuration that is based on the parameters that are specified in the command.
<code>check object</code> { <i>statement-data</i> }	Starts an operation to report on errors in the object, which is a synchronous operation.
<code>clear object</code> { <i>statement-data</i> }	Discards the contents of some attributes of an object. This operation is destructive and cannot be reversed.
<code>create object</code> { <i>statement-data</i> }	Creates an object of the specified type.
<code>deactivate object</code> { <i>statement-data</i> }	Removes the environment for an operation.
<code>delete object</code>	Deletes a previously created object.
<code>diagnose object</code> { <i>statement-data</i> }	Runs a test and shows the results.
<code>disable object</code> { <i>statement-data</i> }	Prevents a feature from operating.
<code>download object</code> { <i>statement-data</i> }	Transfers data to the storage array or to the hardware that is associated with the storage array.
<code>enable object</code> { <i>statement-data</i> }	Sets a feature to operate.
<code>load object</code> { <i>statement-data</i> }	Transfers data to the storage array or to the hardware that is associated with the storage array. This command is functionally similar to the <code>download</code> command.
<code>recopy object</code> { <i>statement-data</i> }	Restarts a volume copy operation by using an existing volume copy pair. You can change the parameters before the operation is restarted.
<code>recover object</code> { <i>statement-data</i> }	Re-creates an object from saved configuration data and the statement parameters. (This command is similar to the <code>create</code> command.)
<code>recreate object</code> { <i>statement-data</i> }	Restarts a snapshot (legacy) operation by using an existing snapshot (legacy) volume. You can change the parameters before the operation is restarted.

Syntax	Description
<code>remove object</code> { <i>statement-data</i> }	Removes a relationship from between objects.
<code>repair object</code> { <i>statement-data</i> }	Repairs errors found by the <code>check</code> command.
<code>replace object</code> { <i>statement-data</i> }	The specified object replaces an existing object in the storage array.
<code>reset object</code> { <i>statement-data</i> }	Returns the hardware or an object to an initial state.
<code>resume object</code>	Starts a suspended operation. The operation starts where it left off when it was suspended.
<code>revive object</code>	Forces the object from the Failed state to the Optimal state. Use this command only as part of an error recovery procedure.
<code>save object</code> { <i>statement-data</i> }	Writes information about the object to a file.
<code>set object</code> { <i>statement-data</i> }	Changes object attributes. All changes are completed when the command returns.
<code>show object</code> { <i>statement-data</i> }	Shows information about the object.
<code>start object</code> { <i>statement-data</i> }	Starts an asynchronous operation. You can stop some operations after they have started. You can query the progress of some operations.
<code>stop object</code> { <i>statement-data</i> }	Stops an asynchronous operation.
<code>suspend object</code> { <i>statement-data</i> }	Stops an operation. You can then restart the suspended operation, and it continues from the point where it was suspended.

Recurring Syntax Elements

Recurring syntax elements are a general category of parameters and options that you can use in the script commands. [Table 3](#) lists the recurring syntax parameters and the values that you can use with the recurring syntax parameters. The conventions used in the recurring syntax elements are listed in the following table.

Convention	Definition
<code>a b</code>	Alternative ("a" or "b")
<i>italicized-words</i>	A terminal that needs user input to fulfill a parameter (a response to a variable)

Convention	Definition
[...] (square brackets)	Zero or one occurrence (square brackets are also used as a delimiter for some command parameters)
{ ... } (curly braces)	Zero or more occurrences
(a b c)	Choose only one of the alternatives
bold	A terminal that needs a command parameter entered to start an action

Table 3 Recurring Syntax Elements

Recurring Syntax	Syntax Value
<i>raid-level</i>	(0 1 3 5 6)
<i>repository-raid-level</i>	(1 3 5 6)
<i>capacity-spec</i>	<i>integer-literal</i> [KB MB GB TB Bytes]
<i>segment-size-spec</i>	<i>integer-literal</i>
<i>boolean</i>	(TRUE FALSE)
<i>user-label</i>	<i>string-literal</i> Valid characters are alphanumeric, the dash, and the underscore.
<i>user-label-list</i>	<i>user-label</i> { <i>user-label</i> }
<i>create-raid-vol-attr-value-list</i>	<i>create-raid-volume-attribute-value-pair</i> { <i>create-raid-volume-attribute-value-pair</i> }
<i>create-raid-volume-attribute-value-pair</i>	capacity= <i>capacity-spec</i> owner=(a b) cacheReadPrefetch=(TRUE FALSE) segmentSize= <i>integer-literal</i> usageHint= <i>usage-hint-spec</i>
<i>noncontroller-trayID</i>	(0-99)
<i>slotID</i>	(1-32)
<i>portID</i>	(0-127)
<i>drive-spec</i>	<i>trayID</i> , <i>slotID</i> or <i>trayID</i> , <i>drawerID</i> , <i>slotID</i> A drive is defined as two or three interger literal values separated by a comma. Low-density trays require two values. High-density trays, those trays that have drawers, require three values.
<i>drive-spec-list</i>	<i>drive-spec</i> <i>drive-spec</i>
<i>trayID-list</i>	<i>trayID</i> { <i>trayID</i> }

Recurring Syntax	Syntax Value
<i>esm-spec-list</i>	<i>esm-spec</i> { <i>esm-spec</i> }
<i>esm-spec</i>	<i>trayID</i> , (left right)
<i>hex-literal</i>	<i>0xhexadecimal-literal</i>
<i>volumeGroup-number</i>	<i>integer-literal</i>
<i>filename</i>	<i>string-literal</i>
<i>error-action</i>	(stop continue)
<i>drive-channel-identifier</i> (four drive ports per tray)	(1 2 3 4)
<i>drive-channel-identifier</i> (eight drive ports per tray)	(1 2 3 4 5 6 7 8)
<i>drive-channel-identifier-list</i>	<i>drive-channel-identifier</i> { <i>drive-channel-identifier</i> }
<i>host-channel-identifier</i> (four host ports per tray)	(a1 a2 b1 b2)
<i>host-channel-identifier</i> (eight host ports per tray)	(a1 a2 a3 a4 b1 b2 b3 b4)
<i>host-channel-identifier</i> (16 host ports per tray)	(a1 a2 a3 a4 a5 a6 a7 a8 b1 b2 b3 b4 b5 b6 b7 b8)
<i>drive-type</i>	(fibre SATA SAS)
<i>drive-media-type</i>	(HDD SSD unknown allMedia) <i>HDD</i> means hard disk drive. <i>SSD</i> means solid state disk.
<i>repository-spec</i>	<i>instance-based-repository-spec</i> <i>count-based-repository-spec</i>

Recurring Syntax	Syntax Value
<i>instance-based-repository-spec</i>	<pre>(repositoryRAIDLevel =repository-raid-level repositoryDrives= (drive-spec-list) [repositoryVolumeGroupUserLabel =user-label] [trayLossProtect=(TRUE FALSE)¹] [drawerLossProtect=(TRUE FALSE)²] (repositoryVolumeGroup=user-label [freeCapacityArea=integer-literal³])</pre> <p>Specify the repositoryRAIDLevel parameter with the repositoryDrives parameter. Do not specify the RAID level or the drives with the volume group. Do not set a value for the trayLossProtect parameter when you specify a volume group.</p>
<i>count-based-repository-spec</i>	<pre>repositoryRAIDLevel =repository-raid-level repositoryDriveCount=integer-literal [repositoryVolumeGroupUserLabel =user-label] [driveType=drive-type⁴] [trayLossProtect=(TRUE FALSE)¹] [drawerLossProtect=(TRUE FALSE)²] [dataAssurance=(none enabled)⁵] </pre>
<i>wwID</i>	<i>string-literal</i>
<i>gid</i>	<i>string-literal</i>
<i>host-type</i>	<i>string-literal</i> <i>integer-literal</i>
<i>host-card-identifier</i>	(1 2 3 4)
<i>backup-device-identifier</i>	(1 n all) n is a specific slot number. Specifying all includes all of the cache backup devices available to the entire storage array.
<i>nvsram-offset</i>	<i>hex-literal</i>
<i>nvsram-byte-setting</i>	<i>nvsram-value</i> = <i>0xhexadecimal</i> <i>integer-literal</i> The <i>0xhexadecimal</i> value is typically a value from 0x0000 to 0xFFFF.

Recurring Syntax	Syntax Value
<i>nvsram-bit-setting</i>	<i>nvsram-mask, nvsram-value = 0xhexadecimal, 0xhexadecimal integer-literal</i> The <i>0xhexadecimal</i> value is typically a value from 0x0000 to 0xFFFF.
<i>ip-address</i>	(0-255).(0-255).(0-255).(0-255)
<i>ipv6-address</i>	(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF) You must enter all 32 hexadecimal characters.
<i>autoconfigure-vols-attr-value-list</i>	<i>autoconfigure-vols-attr-value-pair</i> { <i>autoconfigure-vols-attr-value-pair</i> }
<i>autoconfigure-vols-attr-value-pair</i>	<i>driveType=drive-type </i> <i>driveMediaType=drive-media-type </i> <i>raidLevel=raid-level </i> <i>volumeGroupWidth=integer-literal </i> <i>volumeGroupCount=integer-literal </i> <i>volumesPerGroupCount=integer-literal</i> ⁶ <i>hotSpareCount=integer-literal </i> <i>segmentSize=segment-size-spec </i> <i>cacheReadPrefetch=(TRUE FALSE)</i> <i>securityType=(none capable </i> <i>enabled)</i> ⁷ <i>dataAssurance=(none enabled)</i> ⁵
<i>create-volume-copy-attr-value-list</i>	<i>create-volume-copy-attr-value-pair</i> { <i>create-volume-copy-attr-value-pair</i> }
<i>create-volume-copy-attr-value-pair</i>	<i>copyPriority=(highest high medium </i> <i>low lowest) </i> <i>targetReadOnlyEnabled=(TRUE FALSE) </i> <i>copyType=(offline online) </i> <i>repositoryPercentOfBase=(20 40 60 </i> <i>120 default) </i> <i>repositoryGroupPreference=(sameAsSource </i> <i>otherThanSource default)</i>
<i>recover-raid-volume-attr-value-list</i>	<i>recover-raid-volume-attr-value-pair</i> { <i>recover-raid-volume-attr-value-pair</i> }
<i>recover-raid-volume-attr-value-pair</i>	<i>owner=(a b) </i> <i>cacheReadPrefetch=(TRUE FALSE) </i> <i>dataAssurance=(none enabled)</i>
<i>cache-flush-modifier-setting</i>	<i>immediate, 0, .25, .5, .75, 1, 1.5, 2,</i> <i>5, 10, 20, 60, 120, 300, 1200, 3600,</i> <i>infinite</i>

Recurring Syntax	Syntax Value
<i>serial-number</i>	string-literal
<i>usage-hint-spec</i>	usageHint=(multiMedia database fileSystem)
<i>iscsiSession</i>	[<i>session-identifier</i>]
<i>iscsi-host-port</i>	(1 2 3 4) The host port number might be 2, 3, or 4 depending on the type of controller you are using.
<i>ethernet-port-options</i>	enableIPv4=(TRUE FALSE) enableIPv6=(TRUE FALSE) IPv6LocalAddress= <i>ipv6-address</i> IPv6RoutableAddress= <i>ipv6-address</i> IPv6RouterAddress= <i>ipv6-address</i> IPv4Address= <i>ip-address</i> IPv4ConfigurationMethod=(static dhcp) IPv4GatewayIP= <i>ip-address</i> IPv4SubnetMask= <i>ip-address</i> duplexMode=(TRUE FALSE) portSpeed=(autoNegotiate 10 100 1000)

Recurring Syntax	Syntax Value
<i>iscsi-host-port-options</i>	IPv4Address= <i>ip-address</i> IPv6LocalAddress= <i>ipv6-address</i> IPv6RoutableAddress= <i>ipv6-address</i> IPv6RouterAddress= <i>ipv6-address</i> enableIPv4=(TRUE FALSE) enableIPv6=(TRUE FALSE) enableIPv4Priority=(TRUE FALSE) enableIPv6Priority=(TRUE FALSE) IPv4ConfigurationMethod= (static dhcp) IPv6ConfigurationMethod= (static auto) IPv4GatewayIP= <i>ip-address</i> IPv6HopLimit= <i>integer</i> IPv6NdDetectDuplicateAddress= <i>integer</i> IPv6NdReachableTime= <i>time-interval</i> IPv6NdRetransmitTime= <i>time-interval</i> IPv6NdTimeOut= <i>time-interval</i> IPv4Priority= <i>integer</i> IPv6Priority= <i>integer</i> IPv4SubnetMask= <i>ip-address</i> IPv4VlanId= <i>integer</i> IPv6VlanId= <i>integer</i> maxFramePayload= <i>integer</i> tcpListeningPort= <i>tcp-port-id</i> portSpeed=(autoNegotiate 1 10)
<i>test-devices-list</i>	<i>test-devices</i> { <i>test-devices</i> }
<i>test-devices</i>	controller=(a b) esms=(<i>esm-spec-list</i>) drives=(<i>drive-spec-list</i>)
<i>snapshot</i> (<i>legacy</i>)- <i>schedule-attribute-value-list</i>	<i>snapshot</i> (<i>legacy</i>)- <i>schedule-attribute-value-pair</i> { <i>snapshot</i> (<i>legacy</i>)- <i>schedule-attribute-value-pair</i> }
<i>time-zone-spec</i>	(GMT+HH:MM GMT-HH:MM) [dayLightSaving=HH:MM]
<i>snapshot</i> (<i>legacy</i>)- <i>schedule-attribute-value-pair</i>	startDate=MM:DD:YY scheduleDay=(<i>dayOfWeek</i> all) startTime=HH:MM scheduleInterval= <i>interger</i> endDate=(MM:DD:YY noEndDate) timesPerDay= <i>interger</i>

¹For tray loss protection to work, each drive in a volume group must be in a separate tray. If you set the `trayLossProtect` parameter to `TRUE` and you have selected more than one drive from any one tray, the storage array returns an error. If you set `trayLossProtect` parameter to `FALSE`, the storage array performs operations, but the volume group that you create might not have tray loss protection.

If you set the `trayLossProtect` parameter to `TRUE`, the storage array returns an error if the controller firmware cannot find drives that will enable the new volume group to have tray loss protection. If you set the `trayLossProtect` parameter to `FALSE`, the storage array performs the operation even if it means that the volume group might not have tray loss protection.

²In trays that have drawers for holding the drives, drawer loss protection determines whether data on a volume is accessible or inaccessible if a drawer fails. To help make sure that your data is accessible, set the `drawerLossProtect` parameter to `TRUE`. For drawer loss protection to work, each drive in a volume group must be in separate drawers. If you have a storage array configuration in which a volume group spans several trays, you must make sure that the setting for drawer loss protection works with the setting for tray loss protection. If you set the `trayLossProtect` parameter to `TRUE`, you must set the `drawerLossProtect` parameter to `TRUE`. If you set the `trayLossProtect` parameter to `TRUE`, and you set the `drawerLossProtect` parameter to `FALSE`, the storage array returns an error message and a storage array configuration will not be created.

³To determine if a free capacity area exists, run the `show volumeGroup` command.

⁴The default drive (drive type) is `fibre` (Fibre Channel).

The `driveType` parameter is not required if only one type of drive is in the storage array. If you use the `driveType` parameter, you also must use the `hotSpareCount` parameter and the `volumeGroupWidth` parameter. If you do not use the `driveType` parameter, the configuration defaults to Fibre Channel drives.

⁵The `dataAssurance` parameter applies to the drives in a volume group. Using the `dataAssurance` parameter, you can specify that protected drives must be selected for a volume group. If you want to set the `dataAssurance` parameter to `enabled`, all of the drives in the volume group must be capable of data assurance. You cannot have a mix of drives that are capable of data assurance and drives that are not capable of data assurance in the volume group.

⁶The `volumesPerGroupCount` parameter is the number of equal-capacity volumes per volume group.

⁷The `securityType` parameter enables you to specify the security setting for a volume group that you are creating. All of the volumes are also set to the security setting that you choose. Available options for setting the security setting include:

- `none` – The volume group is not secure.
- `capable` – The volume group is security capable, but security has not been enabled.
- `enabled` – The volume group is security enabled.

NOTE A storage array security key must already be created for the storage array if you want to set `securityType=enabled`. (To create a storage array security key, use the `create storageArray securityKey` command).

Naming Conventions

- Names can have a maximum of 30 characters.
- You can use any combination of alphanumeric characters, hyphens, and underscores for the names of the following components:
 - Storage arrays
 - Host groups
 - Hosts
 - Volume groups
 - Volumes
 - HBA host ports
- You must use unique names. If you do not use unique names, the controller firmware returns an error.
- If the name contains more than one word, hyphens, or underscores, enclose the name in double quotation marks (“ ”). In some usages, you must also surround the name with square brackets ([]). The description of each parameter indicates whether you need to enclose a parameter in double quotation marks, square brackets, or both.
- The name character string cannot contain a new line.
- On Windows operating systems, you must enclose the name between two back slashes (\) in addition to other delimiters. For example, the following name is used in a command that runs under a Windows operating system:

```
[ \ "Engineering\ " ]
```

- For a UNIX operating system and, when used in a script file, the name appears as in the following example:

```
[ "Engineering" ]
```

- When you enter a World Wide Identifier (WWID) of an HBA host port, some usages require that you surround the WWID with double quotation marks. In other uses, you must surround the WWID with angle brackets (<>). The description of the WWID parameter indicates whether you need to enclose the WWID in double quotation marks or angle brackets.

Entering Numerical Names

When the storage management software automatically configures a storage array, the storage management software assigns names that consist of numerical characters. Names that consist only of numerical characters are valid names. Numerical character names, however, must be treated differently than names that start with alphabetic characters.

When you enter a script command that requires a name, the script engine looks for a name that starts with an alphabetic character. The Script Engine might not recognize the following names:

- Names that are only numbers, such as 1 or 2
- Names that start with a number, such as 1Disk or 32Volume

To enter a name that consists only of numerical characters so that the Script Engine will recognize the name, use a combination of back slashes and double quotation marks. The following are examples of how you can enter names that consist only of numerical characters or start with numerical characters:

- [\`"1"`]
- [\`"1Disk"`]

Formatting CLI Commands

Double quotation marks (" ") that are used as part of a name or label require special consideration when you run the CLI commands and the script commands on a Microsoft Windows operating system.

When double quotation marks (" ") are part of a name or value, you must insert a backslash (\) before each double quotation mark character. For example:

```
-c "set storageArray userLabel=\"Engineering";"
```

In this example, "Engineering" is the storage array name. A second example is:

```
-n \"My"_Array
```

In this example, "My" _Array is the name of the storage array.

You cannot use double quotation marks (" ") as part of a character string (also called string literal) within a script command. For example, you cannot enter the following string to set the storage array name to "Finance" Array:

```
-c "set storageArray userLabel=\"\"Finance\"Array";"
```

In the Linux operating system and the Solaris operating system, the delimiters around names or labels are single quotation marks (' '). The UNIX versions of the previous examples are as follows:

```
-c 'set storageArray userLabel="Engineering";'  
-n "My"_Array
```

In a Windows operating system, if you do not use double quotation marks (" ") around a name, you must insert a caret (^) before each special script character. Special characters are ^, |, <, and >.

Insert a caret before each special script character when used with the terminals -n, -o, -f, and -p. For example, to specify storage array CLI>CLIENT, enter this string:

```
-n CLI^>CLIENT
```

Insert one caret (^) before each special script character when used within a string literal in a script command. For example, to change the name of a storage array to FINANCE_>_PAYROLL, enter the following string:

```
-c "set storageArray  
userLabel=\"FINANCE_^|_PAYROLL\"; "
```

Formatting Rules for Script Commands

Syntax unique to a specific script command is explained in the Notes section at the end of each script command description.

Case sensitivity – The script commands are not case sensitive. You can type the script commands in lowercase, uppercase, or mixed case. (In the following command descriptions, mixed case is used as an aid to reading the command names and understanding the purpose of the command.)

Spaces – You must enter spaces in the script commands as they are shown in the command descriptions.

Square brackets – Square brackets are used in two ways:

- As part of the command syntax.
- To indicate that the parameters are optional. The description of each parameter tells you if you need to enclose a parameter value in square brackets.

Parentheses – Parentheses shown in the command syntax enclose specific choices for a parameter. That is, if you want to use the parameter, you must enter one of the values enclosed in parentheses. Generally, you do not include parentheses in a script command; however, in some instances, when you enter lists, you must enclose the list in parentheses. Such a list might be a list of tray ID values and slot ID values. The description of each parameter tells you if you need to enclose a parameter value in parentheses.

Vertical bars – Vertical bars in a script command indicate “or” and separate the valid values for the parameter. For example, the syntax for the `raidLevel` parameter in the command description appears as follows:

```
raidLevel=(0 | 1 | 3 | 5 | 6)
```

To use the `raidLevel` parameter to set RAID Level 5, enter this value:

```
raidLevel=5
```

Drive locations – The CLI commands that identify drive locations support both high-capacity drive trays and low-capacity drive trays. A high-capacity drive tray has drawers that hold the drives. The drawers slide out of the drive tray to provide access to the drives. A low-capacity drive tray does not have drawers. For a high-capacity drive tray, you must specify the identifier (ID) of the drive tray, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive tray, you need only specify the ID of the drive tray and the ID of the slot in which a drive resides. For a low-capacity drive tray, an alternative method for identifying a location for a drive is to specify the ID of the drive tray, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides. Separate the ID values with a comma. If you enter more than one set of ID values, separate each set of values with a space. Enclose the set of values in parentheses. For example:

```
(1,1 1,2 1,3 1,4 2,1 2,2 2,3 2,4)
```

or, for a high-capacity drive tray, this example:

```
(1,1,1 1,2,2 1,3,3 1,4,4 2,1,1 2,2,2 2,3,3 2,4,4)
```

Italicized terms – Italicized terms in the command indicate a value or information that you need to provide. For example, when you encounter the italicized term:

```
numberOfDrives
```

Replace the italicized term with a value for the number of drives that you want to include with the script command.

Semicolon – Script commands must end with a semicolon (;). You can enter more than one script command on the command line or in a script file. For example, a semicolon is used to separate each script command in the following script file.

```
create volume drives=(0,2 0,3 1,4 1,5 2,6 2,7)
raidLevel=5
userLabel="v1" capacity=2gb owner=a;
create volume volumeGroup=2 userLabel="v2"
capacity=1gb owner=b;
create volume volumeGroup=2 userLabel="v3"
capacity=1gb owner=a;
```

```
create volume drives=(0,4 0,5 1,6 1,7 2,8 2,9)
raidLevel=5
userLabel="v4" capacity=2gb owner=b;
create volume volumeGroup=3 userLabel="v5"
capacity=1gb owner=a;
create volume volumeGroup=3 userLabel="v6"
capacity=1gb owner=b;
```

Usage Guidelines

This list provides guidelines for writing script commands on the command line:

- You must end all commands with a semicolon (;).
- You can enter more than one command on a line, but you must separate each command with a semicolon (;).
- You must separate each base command and its associated primary parameters and secondary parameters with a space.
- The script engine is not case sensitive. You can enter commands by using uppercase letters, lowercase letters, or mixed-case letters.
- Add comments to your scripts to make it easier for you and future users to understand the purpose of the script commands. (For information about how to add comments, see "[Adding Comments to a Script File](#)."

NOTE While the CLI commands and the script commands are not case sensitive, user labels (such as for volumes, hosts, or host ports) are case sensitive. If you try to map to an object that is identified by a user label, you must enter the user label exactly as it is defined, or the CLI commands and the script commands will fail.

Detailed Error Reporting

Data collected from an error encountered by the CLI is written to a file. Detailed error reporting under the CLI works as follows:

- If the CLI must abnormally end running CLI commands and script commands, error data is collected and saved before the CLI finishes.
- The CLI saves the error data by writing the data to a standard file name.
- The CLI automatically saves the data to a file. Special command line options are not required to save the error data.
- You are not required to perform any action to save the error data to a file.
- The CLI does not have any provision to avoid over-writing an existing version of the file that contains error data.

For error processing, errors appear as two types:

- Terminal errors or syntax errors that you might enter
- Exceptions that occur as a result of an operational error

When the CLI encounters either type of error, the CLI writes information that describes the error directly to the command line and sets a return code. Depending on the return code, the CLI also might write additional information about which terminal caused the error. The CLI also writes information about what it was expecting in the command syntax to help you identify any syntax errors that you might have entered.

When an exception occurs while a command is running, the CLI captures the error. At the end of processing the command (after the command processing information has been written to the command line), the CLI automatically saves the error information to a file.

The name of the file to which error information is saved is `excprpt.txt`. The CLI tries to place the `excprpt.txt` file in the directory that is specified by the system property `devmgr.datadir`. If for any reason the CLI cannot place the file in the directory specified by `devmgr.datadir`, the CLI saves the `excprpt.txt` file in the same directory from which the CLI is running. You cannot change the file name or the location. The `excprpt.txt` file is overwritten every time that an exception occurs. If you want to save the information in the `excprpt.txt` file, you must copy the information to a new file or a new directory.

Exit Status

This table lists the exit statuses that might be returned and the meaning of each status.

Status Value	Meaning
0	The command terminated without an error.
1	The command terminated with an error. Information about the error also appears.
2	The script file does not exist.
3	An error occurred while opening an output file.
4	A storage array was not at the specified address.
5	Addresses specify different storage arrays.
6	A storage array name does not exist for the host agent that is connected.
7	The storage array name was not at the specified address.
8	The storage array name was not unique.
9	The storage array name was not in the configuration file.
10	A management class does not exist for the storage array.
11	A storage array was not found in the configuration file.
12	An internal error occurred.
13	Invalid script syntax was found.

Status Value	Meaning
14	The controller was unable to communicate with the storage array.
15	A duplicate argument was entered.
16	An execution error occurred.
17	A host was not at the specified address.
18	The WWID was not in the configuration file.
19	The WWID was not at the address.
20	An unknown IP address was specified.
21	The Event Monitor configuration file was corrupted.
22	The storage array was unable to communicate with the Event Monitor.
23	The controller was unable to write alert settings.
24	The wrong organizer node was specified.
25	The command was not available.
26	The device was not in the configuration file.
27	An error occurred while updating the configuration file.
28	An unknown host error occurred.
29	The sender contact information file was not found.
30	The sender contact information file could not be read.
31	The userdata.txt file exists.
32	An invalid -I value in the email alert notification was specified.
33	An invalid -f value in the email alert notification was specified.

Adding Comments to a Script File

The script engine looks for certain characters or a command to show comments. You can add comments to a script file in three ways:

1. Add text after two forward slashes (//) as a comment until an end-of-line character is reached. If the script engine does not find an end-of-line character in the script after processing a comment, an error message appears, and the script operation is terminated. This error usually occurs when a comment is placed at the end of a script and you have forgotten to press the **Enter** key.

```
// Deletes the existing configuration.
set storageArray resetConfiguration=true;
```

2. Add text between /* and */ as a comment. If the script engine does not find both a starting comment notation and an ending comment notation, an error message appears, and the script operation is terminated.

```
/* Deletes the existing configuration */
set storageArray resetConfiguration=true;
```

Firmware Compatibility Levels

- Use the `show` statement to embed comments in a script file that you want to appear while the script file is running. Enclose the text that you want to appear by using double quotation marks (“”).

```
show "Deletes the existing configuration";  
set storageArray resetConfiguration=true;
```

The script commands and the command parameters do not run under all versions of the controller firmware. The script commands in the following sections list the minimum firmware levels under which the script commands can run. In the script commands, the firmware levels are listed under the heading “Minimum Firmware Level.” This list describes how to interpret the information about the firmware levels.

- If a script command does not list a minimum controller firmware level, the script command and all of the parameters associated with that script command can run under any level of controller firmware.
- A controller firmware number without any explanatory information indicates that the controller firmware level applies to the entire script command and all of the parameters for that script command.
- A controller firmware number that is associated with a parameter indicates the minimum controller firmware level under which the parameter can run.

NOTE The minimum controller firmware level indicates support by the software that releases the command, as well as support by all storage management software that picks up usage. CLI support capabilities depend on the hardware used. When an unsupported command is entered, an error message appears.

Examples of Firmware Compatibility Levels

The `create hostGroup` command has the following section.

Minimum Firmware Level

5.20

This level indicates that the entire script command runs under a minimum of controller firmware version 5.20.

The `show volume` command has the following section.

Minimum Firmware Level

5.00

5.43 adds the `summary` parameter

These notations indicate that the script command and all of the parameters except `summary` run under a minimum of controller firmware version 5.00. The `summary` parameter runs under a minimum of controller firmware version 5.43.

ATTENTION The script commands are capable of damaging a configuration and causing loss of data access if not used correctly – Command operations are performed as soon as you run the commands. Some commands can immediately delete configurations or data. Before using the script commands, make sure that you have backed up all data, and have saved the current configuration so that you can reinstall it if the changes you make do not work.

The description of each script command is intended to provide all of the information that you need to be able to use the command. If, however, you have questions about command usage, these sections provide additional information that can help you use the script commands:

- “[Naming Conventions](#)” lists the general rules for entering the names of storage array entities, such as volumes or drives, with the script commands.
- “[Formatting CLI Commands](#)” lists the general formatting rules that apply to the CLI command wrapper.
- “[Formatting Rules for Script Commands](#)” lists the general formatting rules that apply to the script command syntax.
- “[Firmware Compatibility Levels](#)” explains how to interpret the firmware level information.
- “[Command Reference - Listed by Function](#)” lists the script commands organized into groups related to the physical features, the logical features, and the operational features of the storage array.
- “[Command Reference – Listed Alphabetically](#)” lists the script commands alphabetically and, for each script command, includes script command name, syntax, and parameters.

NOTE Terminology differences – The names of components and features change from time to time; however, the command syntax does not change at the same time. You will notice minor differences between the terminology used to describe components and features and the terminology used in the syntax to describe those same items when used in a command name, a parameter, or a variable.

Command Reference - Listed by Function

Asynchronous Mirroring Commands

Activate Asynchronous Mirroring
Add Volume to Asynchronous Mirror Group
Cancel Asynchronous Mirror Group Role Reversal
Check Asynchronous Mirror Group Consistency
Check Storage Array Connectivity
Clear Asynchronous Mirroring Fault
Create Asynchronous Mirror Group
Deactivate Asynchronous Mirroring
Delete Asynchronous Mirror Group
Establish Asynchronous Mirrored Pair
Remove Incomplete Asynchronous Mirrored Pair from Asynchronous Mirror Group
Remove Volume from Asynchronous Mirror Group
Reset Asynchronous Mirror Group Statistics
Reset iSCSI IP Address
Resume Asynchronous Mirror Group
Save Asynchronous Mirror Group Statistics
Set Asynchronous Mirror Group
Show Asynchronous Mirror Group
Show Asynchronous Mirror Group Synchronization Progress
Start Asynchronous Mirroring Synchronization
Suspend Asynchronous Mirror Group
Test Asynchronous Mirror Group Connectivity

AutoSupport Bundle Collection Commands

Disable AutoSupport at the EMW Level SMcli Version
Enable AutoSupport at the EMW Level SMcli Version
Set Storage Array AutoSupport Bundle Disable
Set Storage Array AutoSupport Bundle Enable

Consistency Group Commands

- Add Member to Consistency Group
- Create Consistency Group
- Create Consistency Group Snapshot Image
- Create Consistency Group Snapshot Volume
- Create Consistency Group Snapshot Volume Mapping
- Delete Consistency Group
- Delete Consistency Group Snapshot Image
- Delete Consistency Group Snapshot Volume
- Remove Member Volume from Consistency Group
- Resume Consistency Group Snapshot Volume
- Set Consistency Group Attributes
- Set Consistency Group Snapshot Volume
- Show Consistency Group
- Show Consistency Group Snapshot Image
- Start Consistency Group Snapshot Rollback
- Stop Consistency Group Snapshot Rollback
- Stop Consistency Group Snapshot Volume
- Stop Pending Snapshot Images on Consistency Group

Controller Commands

- Clear Drive Channel Statistics
- Diagnose Controller
- Enable Controller Data Transfer
- Reset Controller
- Save Controller NVSRAM
- Save Drive Channel Fault Isolation Diagnostic Status
- Set Controller
- Set Controller Service Action Allowed Indicator
- Set Drive Channel Status
- Set Host Channel
- Show Cache Backup Device Diagnostic Status
- Show Cache Memory Diagnostic Status

Show Controller
Show Controller Diagnostic Status
Show Controller NVSRAM
Show Drive Channel Statistics
Show Host Interface Card Diagnostic Status
Start Cache Backup Device Diagnostic
Start Cache Memory Diagnostic
Start Controller Diagnostic
Start Controller Trace
Start Drive Channel Fault Isolation Diagnostics
Start Drive Channel Locate
Start Host Interface Card Diagnostic
Stop Cache Backup Device Diagnostic
Stop Cache Memory Diagnostic
Stop Controller Diagnostic
Stop Drive Channel Fault Isolation Diagnostics
Stop Drive Channel Locate
Stop Host Interface Card Diagnostic

**Core Dump
Commands**

Clear Storage Array Core Dump
Save Storage Array Core Dump
Start Storage Array Core Dump

**Database
Commands**

Clear Storage Array Configuration
Clear Storage Array Recovery Mode
Save Storage Array DBM Database
Save Storage Array DBM Validator Information File
Show Storage Array DBM Database
Start Storage Array Configuration Database Diagnostic
Stop Storage Array Configuration Database Diagnostic

**Disk Pool
Commands**

Create Disk Pool
Delete Disk Pool
Set Disk Pool
Set Disk Pool (Modify Disk Pool)
Show Disk Pool
Start Disk Pool Locate
Stop Disk Pool Locate

Drive Commands

Download Drive Firmware
Replace Drive
Revive Drive
Save Drive Channel Fault Isolation Diagnostic Status
Save Drive Log
Set Drive Hot Spare
Set Drive Service Action Allowed Indicator
Set Drive State
Set Foreign Drive to Native
Show Drive
Show Drive Download Progress
Start Drive Channel Fault Isolation Diagnostics
Start Drive Initialize
Start Drive Locate
Start Drive Reconstruction
Start Secure Drive Erase
Stop Drive Channel Fault Isolation Diagnostics
Stop Drive Locate

**Feature
Management
Commands**

Enable Storage Array Feature
Disable Storage Array Feature

Host Topology Commands

Activate Host Port
Activate iSCSI Initiator
Create Host
Create Host Group
Create Host Port
Create iSCSI Initiator
Delete Host
Delete Host Group
Delete Host Port
Delete iSCSI Initiator
Set Host
Set Host Channel
Set Host Group
Set Host Port
Set iSCSI Initiator
Set iSCSI Target Properties
Show Current iSCSI Sessions
Show Host Ports

iSCSI Commands

Create iSCSI Initiator
Delete iSCSI Initiator
Reset Storage Array iSCSI Baseline
Save Storage Array iSCSI Statistics
Set iSCSI Initiator
Set iSCSI Target Properties
Show Current iSCSI Sessions
Show Storage Array Negotiation Defaults
Show Storage Array Unconfigured iSCSI Initiators
Start iSCSI DHCP Refresh
Stop Storage Array iSCSI Session

**Core Dump
Commands**

Check Repository Consistency

Session Command

Set Session

**Snapshot (Legacy)
Commands**

Create Snapshot (Legacy) Volume

Delete Snapshot (Legacy) Volume

Re-create Snapshot (Legacy)

Resume Snapshot (Legacy) Rollback

Set Snapshot (Legacy) Volume

Start Snapshot (Legacy) Rollback

Stop Snapshot (Legacy)

Stop Snapshot (Legacy) Rollback

**Snapshot Group
Commands**

Convert Snapshot (Legacy) Volume to Snapshot Group

Create Snapshot Group

Revive Snapshot Group

Set Snapshot Group Attributes

Set Snapshot Group Media Scan

Set Snapshot Group Repository Volume Capacity

Set Snapshot Group Schedule

Show Snapshot Group

Stop Snapshot Group Pending Snapshot Images

**Snapshot Image
Commands**

Create Snapshot Image

Delete Snapshot Image

Show Snapshot Image

Start Snapshot Image Rollback

Stop Snapshot Image Rollback

**Snapshot Volume
Commands**

Create Read-Only Snapshot Volume

Create Snapshot Volume

Rename Snapshot Volume

Resume Snapshot Volume

SSD Cache Commands

Revive Snapshot Volume
Set Read-Only Snapshot Volume to Read/Write Volume
Set Snapshot Volume Repository Volume Capacity
Set Snapshot Volume Media Scan
Show Snapshot Volumes

Add Drives to SSD Cache
Change SSD Cache Application Type
Create SSD Cache
Delete SSD Cache
Enable or Disable SSD Cache for a Volume
Locate SSD Cache
Remove Drives from SSD Cache
Rename SSD Cache
Resume SSD Cache
Show SSD Cache
Stop SSD Cache Performance Modeling
Start SSD Cache Performance Modeling
Suspend SSD Cache

Storage Array Commands

Activate Storage Array Firmware
Autoconfigure Storage Array
Autoconfigure Storage Array Hot Spares
Clear Storage Array Configuration
Clear Storage Array Event Log
Clear Storage Array Firmware Pending Area
Create Storage Array Security Key
Disable External Security Key Management
Disable Storage Array Feature
Download Storage Array Drive Firmware
Download Storage Array Firmware/NVSRAM
Download Storage Array NVSRAM

Enable External Security Key Management
Enable Storage Array Feature
Export Storage Array Security Key
Import Storage Array Security Key
Load Storage Array DBM Database"
Re-create External Security Key
Reset Storage Array Battery Install Date
Reset Storage Array Diagnostic Data
Reset Storage Array Infiniband Statistics Baseline
Reset Storage Array iSCSI Baseline
Reset Storage Array RLS Baseline
Reset Storage Array SAS PHY Baseline
Reset Storage Array SOC Baseline
Reset Storage Array Volume Distribution
Save Storage Array Configuration
Save Storage Array DBM Database
Save Storage Array DBM Validator Information File
Save Storage Array Diagnostic Data
Save Storage Array Events
Save Storage Array Firmware Inventory
Save Storage Array InfiniBand Statistics
Save Storage Array iSCSI Statistics
Save Storage Array Performance Statistics
Save Storage Array RLS Counts
Save Storage Array SAS PHY Counts
Save Storage Array SOC Counts
Save Storage Array State Capture
Save Storage Array Support Data
Set Storage Array
Set Storage Array ICMP Response
Set Storage Array iSNS Server IPv4 Address

Set Storage Array iSNS Server IPv6 Address
Set Storage Array iSNS Server Listening Port
Set Storage Array iSNS Server Refresh
Set Storage Array Learn Cycle
Set Storage Array Redundancy Mode
Set Storage Array Security Key
Set Storage Array Time
Set Storage Array Tray Positions
Show Storage Array
Show Storage Array Auto Configure
Show Storage Array Host Topology
Show Storage Array LUN Mappings
Show Storage Array Negotiation Defaults
Show Storage Array Unreadable Sectors
Show Storage Array Unconfigured iSCSI Initiators
Start Secure Drive Erase
Start Storage Array Configuration Database Diagnostic
Start Storage Array iSNS Server Refresh
Start Storage Array Locate
Stop Storage Array Configuration Database Diagnostic
Stop Storage Array Drive Firmware Download
Stop Storage Array iSCSI Session
Stop Storage Array Locate
Validate Storage Array Security Key

**Support Bundle
Collection
Commands**

Configure Automatic Support Bundle Collection
Display Automatic Support Bundle Collection Configuration
Schedule Automatic Support Bundle Collection Configuration

**Synchronous
Mirroring
Commands**

Activate Synchronous Mirroring
Check Synchronous Mirroring Status
Create Synchronous Mirroring

Deactivate Synchronous Mirroring
Diagnose Synchronous Mirroring
Re-create Synchronous Mirroring Repository Volume
Remove Synchronous Mirroring
Resume Synchronous Mirroring
Set Synchronous Mirroring
Show Synchronous Mirroring Volume Candidates
Show Synchronous Mirroring Volume Synchronization Progress
Start Synchronous Mirroring Synchronization
Suspend Synchronous Mirroring

Tray Commands

Download Tray Configuration Settings
Save Tray Log
Set Drawer Service Action Allowed Indicator
Set Tray Alarm
Set Tray Identification
Set Tray Service Action Allowed Indicator
Start Tray Locate
Stop Tray Locate

Uncategorized Commands

Set Storage Array ICMP Response
Set Storage Array iSNS Server IPv4 Address
Set Storage Array iSNS Server IPv6 Address
Set Storage Array iSNS Server Listening Port
Set Storage Array iSNS Server Refresh
Set Storage Array Unnamed Discovery Session
Show Storage Array Negotiation Defaults
Show String

Volume Commands

Check Volume Parity
Clear Volume Reservations
Clear Volume Unreadable Sectors
Create RAID Volume (Automatic Drive Select)

Create RAID Volume (Free Extent Based Select)
Create RAID Volume (Manual Drive Select)
Delete Volume
Remove Volume LUN Mapping
Repair Volume Parity
Set Volume Attributes for a Volume in a Volume Group
Show Volume
Show Volume Action Progress
Show Volume Performance Statistics
Show Volume Reservations
Start Volume Initialization

Volume Commands for Disk Pools

Create Volume in Disk Pool
Delete Volume from Disk Pool
Initialize Thin Volume
Set Volume Attributes for a Disk Pool
Set Thin Volume Attributes
Set Volume Attributes for a Volume in a Volume Group
Set Volume Mapping
Show Thin Volume

Volume Copy Commands

Create Volume Copy
Recopy Volume Copy
Remove Volume Copy
Set Volume Copy
Show Volume Copy
Show Volume Copy Source Candidates
Show Volume Copy Target Candidates
Stop Volume Copy

Volume Group Commands

Create Volume Group
Delete Volume Group
Enable Volume Group Security

Revive Volume Group
Set Volume Group
Set Volume Group Forced State
Show Volume Group
Show Volume Group Export Dependencies
Show Volume Group Import Dependencies
Start Volume Group Defragment
Start Volume Group Export
Start Volume Group Import
Start Volume Group Locate
Stop Volume Group Locate

Command Reference – Listed Alphabetically

Activate Asynchronous Mirroring

This command activates the Asynchronous Mirroring premium feature. After you activate the Asynchronous Mirroring premium feature, you must set up an asynchronous mirror group and an asynchronous mirrored pair.

Syntax

```
activate storageArray feature=asyncRemoteMirror
```

Parameters

None.

Minimum Firmware Level

7.84

Activate Host Port

This command activates an inactive host port that was created when the Host Context Agent (HCA) registered the host port to a host.

Syntax

```
activate hostPort "userLabel"
```

Parameter

Parameter	Description
userLabel	The name of the HCA host port. Enclose the host port name in double quotation marks (" ").

Minimum Firmware Level

7.50

Activate iSCSI Initiator

This command activates an inactive iSCSI initiator that was created when the Host Context Agent (HCA) registered the iSCSI initiator to a host.

Syntax

```
activate iscsiInitiator "iscsiID"
```

Parameter

Parameter	Description
iscsiInitiator	The name of the iSCSI initiator. Enclose the name in double quotation marks (" ").

Minimum Firmware Level

7.50

Activate Storage Array Firmware

This command activates firmware that you have previously downloaded to the pending configuration area on the controllers in the storage array.

Syntax

```
activate storageArray firmware
```

Parameters

None.

Minimum Firmware Level

6.10

Activate Synchronous Mirroring

This command creates the mirror repository volume and activates the Synchronous Mirroring premium feature. When you use this command, you can define the mirror repository volume in one of three ways:

- User-defined drives
- User-defined volume group
- User-defined number of drives

If you choose to define a number of drives, the controller firmware chooses which drives to use for the mirror repository volume.

NOTE In previous versions of this command the feature identifier was `remoteMirror`. This feature identifier is no longer valid and is replaced by `syncMirror`.

Syntax (User-Defined Drives)

```
activate storageArray feature=syncMirror
repositoryRAIDLevel=(1 | 3 | 5 | 6)
repositoryDrives=(trayID1,drawerID1,slotID1 ...
trayIDn,drawerIDn,slotIDn)
repositoryVolumeGroupUserLabel=[volumeGroupName]
driveMediaType=(HDD | SSD | unknown | allMedia)
driveType=(fibre | SATA | SAS)
[trayLossProtect=(TRUE | FALSE)
drawerLossProtect=(TRUE | FALSE)
dataAssurance=(none | enabled)]
```

Syntax (User-Defined Volume Group)

```
activate storageArray feature=syncMirror
repositoryVolumeGroup=volumeGroupName
[freeCapacityArea=freeCapacityIndexNumber]
```

Syntax (User-Defined Number of Drives)

```
activate storageArray feature=syncMirror
repositoryRAIDLevel=(1 | 3 | 5 | 6)
repositoryDriveCount=numberOfDrives
repositoryVolumeGroupUserLabel=[volumeGroupName]
driveMediaType=(HDD | SSD | unknown | allMedia)
driveType=(fibre | SATA | SAS)
[trayLossProtect=(TRUE | FALSE)
drawerLossProtect=(TRUE | FALSE)
dataAssurance=(none | enabled)]
```

Parameters

Parameter	Description
repositoryRAIDLevel	The RAID level for the mirror repository volume. Valid values are 1, 3, 5, or 6.
repositoryDrives	The drives for the mirror repository volume. For high-capacity drive trays, specify the tray ID value, the drawer ID value, and the slot ID value for each drive that you assign to the mirror repository volume. For low-capacity drive trays, specify the tray ID value and the slot ID value for each drive that you assign to the mirror repository volume. Tray ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the tray ID values, the drawer ID values, and the slot ID values in parentheses.

Parameter	Description
repositoryVolumeGroupUserLabel	The alphanumeric identifier (including - and _) that you want to give the new volume group in which the mirror repository volume will be located. Enclose the volume group identifier in square brackets ([]).
repositoryVolumeGroup	The name of the mirror repository volume group where the mirror repository volume is located. (To determine the names of the volume groups in your storage array, run the <code>show storageArray profile</code> command.)
driveMediaType	<p>The type of drive media that you want to use for the mirror repository volume group. Valid drive media are these:</p> <ul style="list-style-type: none"> ■ HDD – Use this option when you have hard drives in the drive tray. ■ SSD – Use this option when you have solid state drives in the drive tray. ■ unknown – Use if you are not sure what types of drive media are in the drive tray. ■ allMedia – Use this option when you want to use all types of drive media that are in the drive tray. <p>Use this parameter when you use the <code>repositoryDriveCount</code> parameter.</p> <p>You must use this parameter when you have more than one type of drive media in your storage array.</p>
driveType	<p>The type of drive that you want to use in the mirror volume. You cannot mix drive types.</p> <p>You must use this parameter when you have more than one type of drive in your storage array.</p> <p>Valid drive types are :</p> <ul style="list-style-type: none"> ■ fibre ■ SATA ■ SAS <p>If you do not specify a drive type, the command defaults to <code>fibre</code>.</p> <p>Use this parameter when you use the <code>repositoryDriveCount</code> parameter.</p>

Parameter	Description
freeCapacityArea	The index number of the free space in an existing volume group that you want to use to create the mirror repository volume. Free capacity is defined as the free capacity between existing volumes in a volume group. For example, a volume group might have the following areas: volume 1, free capacity, volume 2, free capacity, volume 3, free capacity. To use the free capacity following volume 2, you would specify: freeCapacityArea=2 Run the show volumeGroup command to determine if a free capacity area exists.
repositoryDriveCount	The number of unassigned drives that you want to use for the mirror repository volume.
trayLossProtect	The setting to enforce tray loss protection when you create the mirror repository volume. To enforce tray loss protection, set this parameter to TRUE. The default value is FALSE.
drawerLossProtect	The setting to enforce drawer loss protection when you create the mirror repository volume. To enforce drawer loss protection, set this parameter to TRUE. The default value is FALSE.
dataAssurance	The setting to specify that a volume group, and the volumes within the volume group, has data assurance protection to make sure that the data maintains its integrity. When you use this parameter, only protected drives can be used for the volume group. These settings are valid: <ul style="list-style-type: none"> ■ none – The volume group does not have data assurance protection. ■ enabled – The volume group has data assurance protection. The volume group supports protected information and is formatted with protection information enabled.

Notes

The repositoryDrives parameter supports both high-capacity drive trays and low-capacity drive trays. A high-capacity drive tray has drawers that hold the drives. The drawers slide out of the drive tray to provide access to the drives. A low-capacity drive tray does not have drawers. For a high-capacity drive tray, you must specify the

identifier (ID) of the drive tray, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive tray, you need only specify the ID of the drive tray and the ID of the slot in which a drive resides. For a low-capacity drive tray, an alternative method for identifying a location for a drive is to specify the ID of the drive tray, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

If the drives that you select for the `repositoryDrives` parameter are not compatible with other parameters (such as the `repositoryRAIDLevel` parameter), the script command returns an error, and Synchronous Mirroring is not activated. The error returns the amount of space that is needed for the mirror repository volume. You can then re-enter the command, and specify the appropriate amount of space.

If you enter a value for the repository storage space that is too small for the mirror repository volumes, the controller firmware returns an error message that provides the amount of space that is needed for the mirror repository volumes. The command does not try to activate Synchronous Mirroring. You can re-enter the command by using the value from the error message for the repository storage space value.

When you assign the drives, if you set the `trayLossProtect` parameter to `TRUE` and have selected more than one drive from any one tray, the storage array returns an error. If you set the `trayLossProtect` parameter to `FALSE`, the storage array performs operations, but the volume group that you create might not have tray loss protection.

When the controller firmware assigns the drives, if you set the `trayLossProtect` parameter to `TRUE`, the storage array returns an error if the controller firmware cannot provide drives that result in the new volume group having tray loss protection. If you set the `trayLossProtect` parameter to `FALSE`, the storage array performs the operation even if it means that the volume group might not have tray loss protection.

The `drawerLossProtect` parameter defines if data on a volume is accessible if a drawer fails. When you assign the drives, if you set the `drawerLossProtect` parameter to `TRUE` and select more than one drive from any one drawer, the storage array returns an error. If you set the `drawerLossProtect` parameter to `FALSE`, the storage array performs operations, but the volume group that you create might not have drawer loss protection.

You must set the `trayLossProtect` parameter and the `drawerLossProtect` parameter to the same value. Both of the parameters must be either `TRUE` or `FALSE`. If the `trayLossProtect` parameter and the `drawerLossProtect` parameter are set to different values, the storage array returns an error.

Minimum Firmware Level

6.10

7.10 adds RAID Level 6 capability.

7.60 adds the `drawerID` user input, the `driveMediaType` parameter, and the `drawerLossProtect` parameter.

7.75 adds the `dataAssurance` parameter.

Add Drives to SSD Cache

Use this command to increase the capacity of an existing SSD cache by adding additional solid state disks (SSDs).

Syntax

```
set ssdCache [ssdCacheName]  
addDrives=(trayID1,drawerID1,slotID1 ...  
trayIDn,drawerIDn,slotIDn)
```

Parameters

Parameter	Description
<code>ssdCache</code>	The alphanumeric identifier (including - and _) of the SSD cache to which you want to add SSDs. Enclose the identifier in square brackets ([]). If the SSD cache name contains special characters or consists only of numbers, you also must enclose the identifier in double quotation marks (" ") inside square brackets.
<code>addDrives</code>	The drives that you want to add to the SSD cache. For high-capacity drive trays, specify the tray ID value, the drawer ID value, and the slot ID value for each SSD that you want to add. For low-capacity drive trays, specify the tray ID value and the slot ID value for each SSD that you want to add. Tray ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the tray ID values, the drawer ID values, and the slot ID values in parentheses.

Notes

The SSD cache can contain any number of SSDs. The maximum size of the SSD cache is 5 TB, but might be less depending on the size of the controller's primary cache.

Minimum Firmware Level

7.84

Add Member to Consistency Group

This command adds a new base volume as a member to an existing consistency group. You can specify an existing existing repository volume for the new consistency group member, or create a new repository volume. When you create a new repository volume, you identify an existing volume group or an existing disk pool where you want the repository volume.

Syntax for Use With an Existing Repository Volume

```
set consistencyGroup ["consistencyGroupName"]
addCGMemberVolume="baseVolumeName"
repositoryVolume="repos_XXXX"
```

Syntax for Use When Creating a New Repository Volume in a Volume Group

```
set consistencyGroup ["consistencyGroupName"]
addCGMemberVolume="baseVolumeName"
repositoryVolume=( "volumeGroupName"
capacity=capacityValue (KB|MB|GB|TB|bytes) )
```

Syntax for Use When Creating a New Repository Volume in a Disk Pool

```
set consistencyGroup ["consistencyGroupName"]
addCGMemberVolume="baseVolumeName"
repositoryVolume=( "diskPoolName"
capacity=capacityValue (KB|MB|GB|TB|bytes) )
```

Parameters

Parameter	Description
consistencyGroup	The name of the consistency group to which you want to add a new member volume. The new member volume is the base volume for snapshot operations. Enclose the consistency group name in double quotation marks (" ") inside square brackets ([]).
addCGMemberVolume	The name of a base volume that you want to add. Enclose the volume name in square brackets ([]). If the volume name has special characters, you also must enclose the member name in double quotation marks (" "). If the specified volume is an existing repository volume or an existing snapshot volume, the command fails.
repositoryVolume	This parameter performs two functions: <ul style="list-style-type: none">■ In an existing consistency group that has a repository volume, this parameter identifies the repository volume.■ When creating a new repository volume this parameter identifies either a volume group or disk pool in which to create the new repository volume.
capacity	The size of a new repository volume in either a volume group or a disk pool. Size is defined in units of bytes, KB, MB, GB, or TB.

Notes

You can use any combination of alphanumeric characters, underscore (`_`), hyphen (`-`), and pound (`#`) for the names. Names can have a maximum of 30 characters.

A consistency group is a collection of base volumes that are the source volumes for snapshots. You collect the base volumes in a consistency group so that you can perform the same snapshot operations on each of the base volumes. In the context of this command, the term *member* means a base volume for snapshot operations. You can manipulate snapshot images associated with members of a consistency group through batch-style operations, while maintaining consistency across the snapshot images.

Each member volume must have a corresponding repository volume. You can define the relationship between the member volume and the repository volume using the `repositoryVolume` parameter. The `repositoryVolume` parameter can perform one of these actions:

- Identify an existing repository volume that is connected to the consistency group.
- Identify either a volume group or a disk pool in which you want to create a new repository volume.

Adding a member to a consistency group with an existing repository has two purposes:

- You can create an entirely new repository volume by running the command without the `repositoryVolume` parameter. When you run the command without the `repositoryVolume` parameter the command creates a new repository volume in the volume group or disk pool in which all the other repository volumes are stored. An example of this command usage is:

```
set consistencyGroup ["First_Images"]
addCGMemberVolume="Data_020212"
```

- You can reuse an existing repository volume if that repository volume is empty and is not related to another member volume. You might want to do this if you want to maintain a particular sequence or relationship for the repository volumes. To reuse an existing, empty repository volume you need to know the name of the repository volume. To determine the name of the repository volume use the `show allVolumes summary` command. All repository volume names have the form `repos_XXXX` where `XXXX` is a unique identifier generated by the storage management software. An example of this command usage is:

```
set consistencyGroup ["First_Images"]
addCGMemberVolume="Data_020212"
repositoryVolume="repos_0011"
```

You can place the repository volume in a volume group or a disk pool of your choosing. You are not required to have the repository volume in the same location as other repository volumes. To place the repository volume in a volume group or a disk

pool of your choice, use the `repositoryVolume` parameter and identify the volume group or the disk pool and a size for the repository volume. An example of this command usage is:

```
set consistencyGroup ["First_Images"]
addCGMemberVolume="Data_020212"
repositoryVolume=("12" capacity=2 GB)
```

In this example, "12" is the name of an existing volume group or an existing disk pool. The capacity parameter defines the size that you want for the repository volume group.

When you create a new repository volume in either a volume group or a disk pool, you must include parenthesis around the volume group name and capacity, or the disk pool name and capacity.

Minimum Firmware Level

7.83

Add Volume to Asynchronous Mirror Group

This command adds a primary volume to an asynchronous mirror group. This command is valid only on the local storage array that contains the asynchronous mirror group to which you want to add the primary volume. An asynchronous mirror group has a repository volume that is used to save data for all of the point-in-time images that are part of the asynchronous mirror group. Each primary volume in the asynchronous mirror group has a corresponding mirror volume on a remote storage array.

Syntax

```
add volume="volumeName"
asyncMirrorGroup="asyncMirrorGroupName"
remotePassword="password"
(repositoryVolume="repos_xxxx" |
repositoryVolume=(volumeGroupName [capacity=capacityValue])
repositoryVolume=(diskPoolName [capacity=capacityValue]))
```

Parameters

Parameter	Description
<code>volume</code>	The name of the primary volume that you want to add to the asynchronous mirror group. Enclose the volume name in double quotation marks (" ").
<code>asyncMirrorGroup</code>	The name of the asynchronous mirror group that will contain the primary volume that you want to add. Enclose the asynchronous mirror group name in double quotation marks (" ").

Parameter	Description
remotePassword	<p>This parameter is optional. Use this parameter when the remote storage array is password protected. Enclose the password in double quotation marks (" ").</p>
repositoryVolume	<p>The name of the repository volume that will contain the changed data from the primary volume.</p> <p>You have two options for defining the name of a repository volume:</p> <ul style="list-style-type: none"> ■ Use an existing repository volume: name. ■ Create a new repository volume when you run this command. <p>The name of an existing repository volume is comprised of two parts:</p> <ul style="list-style-type: none"> ■ The term <i>repos</i> ■ A four digit numerical identifier that the storage management software assigns to the repository volume name <p>Enclose the name of the existing repository volume in double quotation marks (" ").</p> <p>If you want to create a new repository volume when you run this command you must enter the name of either a volume group or a disk pool in which you want the repository volume. Optionally, you also can define the capacity of the repository volume. If you want to define the capacity you can use these values:</p> <ul style="list-style-type: none"> ■ An integer value that represents a percentage of the base volume capacity ■ A decimal fraction value that represents a percentage of the base volume capacity ■ A specific size for the repository volume. Size is defined in units of bytes, KB, MB, GB, or TB. <p>If you do not define the capacity, the storage management software sets the capacity to 20 percent of the primary volume capacity.</p> <p>The storage management software creates the repository volume and links the repository volume to the primary volume.</p>

Autoconfigure Storage Array

Notes

- The Asynchronous Mirroring premium feature must be enabled and activated on the local and remote storage arrays that will be used for mirror activities.
- The local and remote storage arrays must be connected through a proper Fibre Channel fabric or iSCSI interface.
- The remote storage array must have a volume that has a capacity that is greater than or equal to the capacity of the primary volume on the local storage array. The volume on the remote storage array will be used as the mirror volume.

Minimum Firmware Level

7.84

This command automatically configures a storage array. Before you enter the `autoConfigure storageArray` command, run the `show storageArray autoConfiguration` command. The `show storageArray autoConfiguration` command returns configuration information in the form of a list of valid drive types, RAID levels, volume information, and hot spare information. (This list corresponds to the parameters for the `autoConfigure storageArray` command.) The controllers audit the storage array and then determine the highest RAID level that the storage array can support and the most efficient volume definition for the RAID level. If the configuration that is described by the returned list is acceptable, you can enter the `autoConfigure storageArray` command without any parameters. If you want to modify the configuration, you can change the parameters to meet your configuration requirements. You can change a single parameter or all of the parameters. After you enter the `autoConfigure storageArray` command, the controllers set up the storage array by using either the default parameters or those you selected.

Syntax

```
autoConfigure storageArray
[driveType=(fibre | SATA | SAS)
driveMediaType=(hdd | ssd | allMedia | unknown) |
raidLevel=(0 | 1 | 3 | 5 | 6)
volumeGroupWidth=numberOfDrives
volumeGroupCount=numberOfVolumeGroups
volumesPerGroupCount=numberOfVolumesPerGroup
hotSpareCount=numberOfHotSpares
segmentSize=segmentSizeValue
cacheReadPrefetch=(TRUE | FALSE)
readAheadMultiplier=multiplierValue
securityType=(none | capable | enabled)
dataAssurance=(none | enabled)]
```

Parameters

Parameter	Description
driveType	<p>The type of drives that you want to use for the storage array.</p> <p>You must use this parameter when you have more than one type of drive in your storage array.</p> <p>Valid drive types are :</p> <ul style="list-style-type: none"> ■ fibre ■ SATA ■ SAS <p>If you do not specify a drive type, the command defaults to fibre.</p>
driveMediaType	<p>The type of drive media that you want to use for the storage array.</p> <p>You must use this parameter when you have more than one type of drive media in your storage array.</p> <p>Valid drive media are :</p> <ul style="list-style-type: none"> ■ hdd – Use this option when you have hard drives. ■ ssd – Use this option when you have solid state disks. ■ unknown – Use if you are not sure what types of drive media are in the drive tray. <p>The default value is hdd.</p>
raidLevel	The RAID level of the volume group that contains the drives in the storage array. Valid RAID levels are 0, 1, 3, 5, or 6.
volumeGroupWidth	The number of drives in a volume group in the storage array.
volumeGroupCount	The number of volume groups in the storage array. Use integer values.
volumesPerGroupCount	The number of equal-capacity volumes per volume group. Use integer values.
hotSpareCount	The number of hot spares that you want in the storage array. Use integer values.
segmentSize	The amount of data (in KB) that the controller writes on a single drive in a volume before writing data on the next drive. Valid values are 8, 16, 32, 64, 128, 256, or 512.

Parameter	Description
cacheReadPrefetch	The setting to turn on or turn off cache read prefetch. To turn off cache read prefetch, set this parameter to <code>FALSE</code> . To turn on cache read prefetch, set this parameter to <code>TRUE</code> .
readAheadMultiplier	<p>This parameter defines how many additional data blocks are read into cache. Valid values range from 0 to 65535.</p> <hr/> <p>NOTE This parameter is deprecated and will be removed in a future release of storage management software. For best operation use the <code>cacheReadPrefetch</code> parameter.</p> <hr/>
securityType	<p>The setting to specify the security level when creating the volume groups and all associated volumes. These settings are valid:</p> <ul style="list-style-type: none"> ■ <code>none</code> – The volume group and volumes are not secure. ■ <code>capable</code> – The volume group and volumes are capable of having security set, but security has not been enabled. ■ <code>enabled</code> – The volume group and volumes have security enabled.
dataAssurance	<p>The setting to specify that a volume group, and the volumes within the volume group, has data assurance protection to make sure that the data maintains its integrity. When you use this parameter, only protected drives can be used for the volume group. These settings are valid:</p> <ul style="list-style-type: none"> ■ <code>none</code> – No explicit data assurance is defined. The volume group will be comprised of volumes that may or may not have data assurance, depending on the availability of data assurance enabled drives. Data Assurance is enabled on volumes created on volume groups that support data assurance. If only data assurance drives are available, the volume groups will be data assurance enabled. ■ <code>enabled</code> – The volume group has data assurance protection. The volume group supports protected information and is formatted with protection information enabled.

Notes

Drives and Volume Group

A volume group is a set of drives that are logically grouped together by the controllers in the storage array. The number of drives in a volume group is a limitation of the RAID level and the controller firmware. When you create a volume group, follow these guidelines:

- Beginning with firmware version 7.10, you can create an empty volume group so that you can reserve the capacity for later use.
- You cannot mix drive types, such as SAS, SATA and Fibre Channel, within a single volume group.
- The maximum number of drives in a volume group depends on these conditions:
 - The type of controller
 - The RAID level
- RAID levels include: 0, 1, 10, 3, 5, and 6.
 - In a CDE3992 or a CDE3994 storage array, a volume group with RAID level 0 and a volume group with RAID level 10 can have a maximum of 112 drives.
 - In a CE6998 storage array, a volume group with RAID level 0 and a volume group with RAID level 10 can have a maximum of 224 drives.
 - A volume group with RAID level 3, RAID level 5, or RAID level 6 cannot have more than 30 drives.
 - A volume group with RAID level 6 must have a minimum of five drives.
 - If a volume group with RAID level 1 has four or more drives, the storage management software automatically converts the volume group to a RAID level 10, which is RAID level 1 + RAID level 0.
- If a volume group contains drives that have different capacities, the overall capacity of the volume group is based on the smallest capacity drive.
- To enable tray loss protection, you must create a volume group that uses drives located in at least three drive trays.

Hot Spares

Hot spare drives can replace any failed drive in the storage array. The hot spare must be the same type of drive as the drive that failed (that is, a SAS hot spare cannot replace a Fibre Channel drive). A hot spare must have capacity greater than or equal to any drive that can fail. If a hot spare is smaller than a failed drive, you cannot use the hot spare to rebuild the data from the failed drive. Hot spares are available only for RAID Level 1, RAID Level 3, RAID Level 5, or RAID Level 6.

Segment Size

The size of a segment determines how many data blocks that the controller writes on a single drive in a volume before writing data on the next drive. Each data block stores 512 bytes of data. A data block is the smallest unit of storage. The size of a segment determines how many data blocks that it contains. For example, an 8-KB segment holds 16 data blocks. A 64-KB segment holds 128 data blocks.

When you enter a value for the segment size, the value is checked against the supported values that are provided by the controller at run time. If the value that you entered is not valid, the controller returns a list of valid values. Using a single drive for a single request leaves other drives available to simultaneously service other requests.

If the volume is in an environment where a single user is transferring large units of data (such as multimedia), performance is maximized when a single data transfer request is serviced with a single data stripe. (A data stripe is the segment size that is multiplied by the number of drives in the volume group that are used for data transfers.) In this case, multiple drives are used for the same request, but each drive is accessed only once. For optimal performance in a multiuser database or file system storage environment, set your segment size to minimize the number of drives that are required to satisfy a data transfer request.

Cache Read Prefetch

Cache read prefetch lets the controller copy additional data blocks into cache while the controller reads and copies data blocks that are requested by the host from the drive into cache. This action increases the chance that a future request for data can be fulfilled from cache. Cache read prefetch is important for multimedia applications that use sequential data transfers. The configuration settings for the storage array that you use determine the number of additional data blocks that the controller reads into cache. Valid values for the `cacheReadPrefetch` parameter are `TRUE` or `FALSE`.

Security Type

The `securityType` parameter is valid for drives that are capable of full disk encryption (FDE). With FDE, the controller firmware can create a key and activate the Drive Security feature. The Drive Security feature encrypts data as the data is written to the drive and decrypts the data as the data is read from the drive. Without the key created by the controller, the data written to the drive is inaccessible.

Before you can set the `securityType` parameter to `capable` or `enabled`, you must create a storage array security key. Use the `create storageArray securityKey` command to create a storage array security key. These commands are related to the security key:

- `create storageArray securityKey`
- `set storageArray securityKey`
- `import storageArray securityKey`
- `export storageArray securityKey`

- `start secureErase (drive | drives)`
- `enable volumeGroup [volumeGroupName] security`

Minimum Firmware Level

6.10

7.10 adds RAID Level 6 capability and removes hot spare limits.

7.50 adds the `securityType` parameter.

7.75 adds the `dataAssurance` parameter.

**Autoconfigure
Storage Array Hot
Spares**

This command automatically defines and configures the hot spares in a storage array. You can run this command at any time. This command provides the best hot spare coverage for a storage array.

Syntax

```
autoConfigure storageArray hotSpares
```

Parameters

None.

Notes

When you run the `autoconfigure storageArray hotSpares` command, the controller firmware determines the number of hot spares to create based on the total number and type of drives in the storage array. For Fibre Channel drives, SATA drives, and SAS drives, the controller firmware creates one hot spare for the storage array and one additional hot spare for every 60 drives in the storage array.

Minimum Firmware Level

6.10

**Cancel
Asynchronous
Mirror Group Role
Reversal**

This command cancels a pending role reversal operation between asynchronous mirror groups.

Syntax

```
stop asyncMirrorGroup ["asyncMirrorGroupName"] rolechange
```

Parameter

Parameter	Description
asyncMirrorGroup	The name of the asynchronous mirror group for which you want to cancel the pending role reversal operation. Enclose the asynchronous mirror group name in square brackets ([]). If the asynchronous mirror group name has special characters, you also must enclose the asynchronous mirror group name in double quotation marks (" ").

Minimum Firmware Level

7.84

Change SSD Cache Application Type

This command changes the application type associated with the SSD cache. The application type can be web server, database, or multimedia. Changing the application type changes the block size, subblock size, populate on read threshold, and populate on write threshold for the volumes underlying the SSD cache.

Syntax

```
set ssdCache [ssdCacheName]
usageHint=(webServer|dataBase|fileSystem)
```

Parameters

Parameter	Description
ssdCache	The alphanumeric identifier (including - and _) of the SSD cache for which you want to change the application type. Enclose the identifier in square brackets ([]). If the SSD cache name contains special characters or consists only of numbers, you also must enclose the identifier in double quotation marks (" ") inside square brackets.
usageHint	The values are based on the typical I/O usage pattern of the application that is using the SSD cache. Valid values are webServer, dataBase, or fileSystem.

Minimum Firmware Level

7.84

Check Asynchronous Mirror Group Consistency

The command produces a report based on analysis of the data in the repository. This command applies to an asynchronous mirror group that has underlying repository volumes.

Syntax

```
check asyncMirrorGroup[asyncMirrorGroupName]
repositoryConsistency localVolume=[ "localVolumeName" ]
file="filePath"
```

Parameters

Parameter	Description
asyncMirrorGroup	The alphanumeric identifier (including - and _) of the asynchronous mirror group on which to run a consistency check. Enclose the asynchronous mirror group identifier in square brackets ([]).
localVolume	Specifies the alphanumeric identifier of a local volume that participates in the mirror relationship. You can specify a local volume that is either the source or the destination on a mirror relationship. Enclose the local volume identifier in double quotation marks (" ") within square brackets ([]).
file	The file path and the file name to which you want to save the report that results from the consistency check. Enclose the file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\logs\repoConsistency.txt" This parameter must appear last, after any of the optional parameters.

Notes

The report generated by this command is intended for use by Technical Support to analyze problems.

Minimum Firmware Level

7.84

Check Repository Consistency

This command applies to a number of objects that have underlying repository volumes. The command produces a report based on analysis of the data in the repository.

Syntax

```
check [snapGroup[snapGroupName] repositoryConsistency |  
snapVolume[snapVolumeName] repositoryConsistency |  
volume[volumeName] repositoryConsistency |  
volumeCopy target[targetName] repositoryConsistency |  
asyncMirrorGroup[asyncMirrorGroupName] repositoryConsistency  
localVolume="localVolumeName" ]  
file="filePath"
```

Parameters

Parameter	Description
repositoryConsistency	Specifies that the snapshot group, snapshot volume, volume, or asynchronous mirror group (depending on which is specified by the corresponding parameter) is checked for consistency.
snapGroup	The alphanumeric identifier (including - and _) of the snapshot group on which to run a consistency check. Enclose the snapshot group identifier in square brackets ([]).
snapVolume	The alphanumeric identifier (including - and _) of the snapshot volume on which to run a consistency check. Enclose the snapshot volume identifier in square brackets ([]).
volume	The alphanumeric identifier (including - and _) of the thin volume on which to run a consistency check. Enclose the volume identifier in square brackets ([]).
volumeCopy	Specifies that the target of a volume copy relationship is checked for consistency.
asyncMirrorGroup	The alphanumeric identifier (including - and _) of the asynchronous mirror group on which to run a consistency check. Enclose the asynchronous mirror group identifier in square brackets ([]).
localVolume	Use only with the <code>asyncMirrorGroup</code> parameter to specify the alphanumeric identifier of a local volume that participates in the mirror relationship. You can specify a local volume that is either the source or the destination on a mirror relationship. Enclose the local volume identifier in double quotation marks (" ").
target	Use only in conjunction with the <code>volumeCopy</code> parameter to specify the alphanumeric identifier of the volume that is the target of the volume copy relationship. Enclose the volume identifier in double quotation marks (" ").

Parameter	Description
file	<p>The file path and the file name to which you want to save the report that results from the consistency check. Enclose the file name in double quotation marks (" "). For example:</p> <pre>file="C:\Program Files\CLI\logs\repoConsistency.txt"</pre> <p>This parameter must appear last, after any of the optional parameters.</p>

Notes

The report generated by this command is intended for use by the Technical Support to analyze problems.

Minimum Firmware Level

7.83

Check Storage Array Connectivity

This command verifies that the local storage array and the remote storage array have a communication path and displays the connection details between the local and remote storage array.

Before creating an asynchronous mirror group, you should check whether the local storage array and the remote storage array can communicate with each other. When you execute this command, the system queries for the remote storage array on all eligible host ports to determine what ports are connected to the remote storage array. The result of the test is a list of all ports on the local storage array along with a list of the remote storage array port addresses accessible through that port.

NOTE Connectivity is tested using all possible channels, and if it is a dual controller configuration, connectivity is checked from each controller. It might take up to 20 minutes to check connectivity between two storage arrays.

Syntax

```
check storageArray connectivity
(remoteStorageArrayName="storageArrayName" |
remoteStorageArrayWwid=<wwID>)
```

Parameters

Parameter	Description
remoteStorageArrayName	The name for the remote storage array for which you are checking connectivity. Enclose the storage array name in double quotation marks (" ").

Parameter	Description
remoteStorageArrayNameWwid	The World Wide Identifier (WWID) of the storage array for which you are checking connectivity. You can use the WWID instead of the storage array name to identify the storage array. Enclose the WWID in angle brackets (<>).

Notes

Before a mirror can be established between two storage arrays, they must be connected through a Fibre Channel connection or an iSCSI connection.

Fibre Channel asynchronous mirroring must be activated on both storage arrays before they can communicate with each other for mirroring or connectivity checking.

If the local storage array supports iSCSI, connectivity over iSCSI is checked. If the local storage array supports Fibre Channel and mirroring over Fibre Channel has been activated, connectivity over Fibre Channel is checked. If the local storage array does not support iSCSI or Fibre Channel, an error message is displayed and the operation aborted.

Minimum Firmware Level

7.84

Check Synchronous Mirroring Status

This command returns the status of a remote-mirror volume. Use this command to determine when the status of the remote-mirror volume becomes Optimal.

NOTE In previous versions of this command the feature identifier was `remoteMirror`. This feature identifier is no longer valid and is replaced by `syncMirror`.

Syntax

```
check syncMirror localVolume [volumeName] optimalStatus
timeout=timeoutValue
```

Parameters

Parameter	Description
localVolume	The name of any remote-mirror volume. The remote-mirror volume can be the primary volume or the secondary volume of a remote mirrored pair. Enclose the volume name in square brackets ([]). If the volume name has special characters, you also must enclose the volume name in double quotation marks (" ").
timeout	The time interval within which the software can return the remote-mirror volume status. The timeout value is in minutes.

Notes

This command waits until the status becomes Optimal or the timeout interval expires. Use this command when you run the Asynchronous Synchronous Mirroring utility.

For more information, see the topic "Asynchronous Synchronous Mirroring Utility."

Minimum Firmware Level

6.10

Check Volume Parity

This command checks a volume for parity and media errors and writes the results of the check to a file.

Syntax

```
check volume [volumeName] parity  
[parityErrorFile=filename |  
mediaErrorFile=filename |  
priority=(highest | high | medium | low | lowest) |  
startingLBA=LBAvalue |  
endingLBA=LBAvalue |  
verbose=(TRUE | FALSE)]
```

Parameters

Parameter	Description
volume	The name of the specific volume for which you want to check parity. Enclose the volume name in square brackets ([]). If the volume name has special characters, you also must enclose the volume name in double quotation marks (" ") inside square brackets. If the volume name consists only of numbers, such as "1002," you also must enclose the volume name in double quotation marks (" ") inside square brackets.
parityErrorFile	The file path and the file name to which you want to save the parity error information. Enclose the file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\logs\parerr.txt" This command does not automatically append a file extension to the saved file. You must specify a file extension when entering the file name.

Parameter	Description
mediaErrorFile	The file path and the file name to which you want to save the media error information. Enclose the file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\logs\mederr.txt" This command does not automatically append a file extension to the saved file. You must specify a file extension when entering the file name.
priority	The priority that the parity check has relative to host I/O activity. Valid values are highest, high, medium, low, or lowest.
startingLBA	The starting logical block address. Use integer values.
endingLBA	The ending logical block address. Use integer values.
verbose	The setting to capture progress details, such as percent complete, and to show the information as the volume parity is being repaired. To capture progress details, set this parameter to TRUE. To prevent capturing progress details, set this parameter to FALSE.

Notes

The starting logical block address and the ending logical block address are useful for very large single-volume LUNs. Running a volume parity check on a very large single volume LUN can take a long time. By defining the beginning address and ending address of the data blocks, you can reduce the time that a volume parity check takes to complete.

Minimum Firmware Level

6.10

Clear Asynchronous Mirroring Fault

This command clears an asynchronous mirroring "sticky" fault from one or more asynchronous mirror groups and one or more asynchronous mirror group member volumes.

An asynchronous mirror group and its member volumes can encounter asynchronous mirroring "sticky" faults, which occur at a single point-in-time but do not impact the functionality of the mirrors. These type of faults must be reviewed, but might or can? not require any changes to the configuration.

An asynchronous mirror group and its member volumes might or can have more than one associated sticky fault. This command clears all of the faults associated with the asynchronous mirror group and its member volume. However, if an asynchronous mirror group has a fault and one of its member volumes has a fault, clearing the fault on the asynchronous mirror group does not clear the fault on its member volume.

Syntax

```
clear asyncMirrorFault(all | asyncMirrorGroup
["asyncMirrorGroupName" ] |
asyncMirrorGroups ["asyncMirrorGroupName1"
... "asyncMirrorGroupNameN" ] |
volume ["volumeName" ] |
volumes ["volumeName1" ... "volumeNameN" ])
```

Parameter

Parameter	Description
all	Use this parameter if you want to clear all asynchronous mirroring faults from all asynchronous mirror groups and on all asynchronous mirror group member volumes.
asyncMirrorGroup	The name of the asynchronous mirror group from which you want to clear the asynchronous mirroring fault. Enclose the asynchronous mirror group name in square brackets ([]). If the asynchronous mirror group name has special characters, you also must enclose the asynchronous mirror group name in double quotation marks (" ").
asyncMirrorGroups	The names of several asynchronous mirror groups from which you want to clear the asynchronous mirroring fault. Enter the names of the asynchronous mirror groups using these rules: <ul style="list-style-type: none"> ■ Enclose all of the names in square brackets ([]). ■ Enclose each of the names in double quotation marks (" "). ■ Separate each of the names with a space.
volume	The name of the specific member volume (in an asynchronous mirror group) from which you want to clear the asynchronous mirroring fault. Enclose the member volume name in square brackets ([]). If the volume name has special characters, you also must enclose the volume name in double quotation marks (" ") inside square brackets. If the volume name consists only of numbers, such as "1002," you also must enclose the volume name in double quotation marks (" ") inside square brackets.
volumes	The names of several member volumes (in an asynchronous mirror group) from which you want to clear the asynchronous mirroring fault. Enter the names of the member volumes using these rules: <ul style="list-style-type: none"> ■ Enclose all of the names in square brackets ([]). ■ Enclose each of the names in double quotation marks (" "). ■ Separate each of the names with a space.

Minimum Firmware Level

7.84

Clear Drive Channel Statistics

This command resets the statistics for all of the drive channels.

Syntax

```
clear all DriveChannels stats
```

Parameters

None.

Minimum Firmware Level

6.10

Clear Storage Array Configuration

Use this command to perform one of these operations:

- Clear the entire storage array configuration, and return it back to the initial installation state
- Clear the configuration except for security information and identification information
- Clear volume group configuration information and volume configuration information only

ATTENTION Possible damage to the storage array configuration – As soon as you run this command, the existing storage array configuration is deleted.

Syntax

```
clear storageArray configuration [all | volumeGroups]
```

Parameters

Parameter	Description
None	If you do not enter a parameter, this command removes all configuration information for the storage array, except for information related to security and identification.
all	The setting to remove the entire configuration of the storage array, including security information and identification information. Removing all configuration information returns the storage array to its initial state.
volumeGroups	The setting to remove the volume configuration and the volume group configuration. The rest of the configuration stays intact.

Notes

With this command you can clear the configuration of the storage array in a number of different ways. When you clear the storage array for volume groups, Recovery mode automatically starts. In recovery mode, onboard (cache) backups (if they exist on the platform) are retained. When the storage array is in recovery mode, you can, optionally, restore the storage array configuration from a host-based backup file or from one of the onboard (cache) backups.

When you run this command, the storage array becomes unresponsive, and all script processing is canceled. You must remove and re-add the storage array to resume communication with the host. To remove an unresponsive storage array, use this SMcli wrapper command:

```
SMcli -X -n storageArrayName
```

The X is a unique SMcli terminal and must be uppercase.

To re-add the storage array, use this SMcli wrapper command:

```
SMcli -A -n storageArrayName
```

The A is a unique SMcli terminal and must be uppercase.

Minimum Firmware Level

6.10

7.10 adds these parameters:

- all
- volumeGroups

7.83 add the recovery mode capability.

Clear Storage Array Core Dump

This command sets a flag on a controller to allow a new core dump to overwrite an existing core dump.

Syntax

```
set storageArray coreDumpAllowOverWrite
```

Parameters

None.

Notes

When you retrieve a core dump from the controller cache to a host, a flag is set on the controller to indicate that the core dump does not need to be retrieved. This setting persists for 48 hours. If a new core dump occurs during that period the new core dump is saved to the controller cache and overwrites any previous core dump data in cache.

You can use the `set storageArray coreDumpAllowOverWrite` command to set the controller flag so that a new core dump will overwrite any previous core dump. Without retrieving a core dump, this command sets the flag as if you had.

Minimum Firmware Level

7.83

Clear Storage Array Event Log

This command clears the Event Log in the storage array by deleting the data in the Event Log buffer.

ATTENTION Possible damage to the storage array configuration – As soon as you run this command, the existing Event Log in the storage array is deleted.

Syntax

```
clear storageArray eventLog
```

Parameters

None.

Minimum Firmware Level

6.10

Clear Storage Array Firmware Pending Area

This command deletes a firmware image or NVSRAM values that you have previously downloaded from the pending area buffer.

ATTENTION Possible damage to the storage array configuration – As soon as you run this command, the contents of the existing pending area in the storage array are deleted.

Syntax

```
clear storageArray firmwarePendingArea
```

Parameters

None.

Minimum Firmware Level

6.10

Clear Storage Array Recovery Mode

This command forces a storage array to exit recovery mode.

Syntax

```
clear storageArray recoveryMode
```

Parameters

None.

Notes

Recovery mode is entered during start-of-day operations when the system configuration is cleared and a valid on board backup is available. This mode is exited by either restoring the system configuration from a backup location, or by clearing the existing on board backups. While recovery mode is in force, a **needs attention** condition is raised and the Recovery Guru is available from the user interface. However, the system configuration is empty during recovery mode.

Minimum Firmware Level

7.83

Clear Volume Reservations

This command clears persistent volume reservations.

Syntax

```
clear (allVolumes | volume [volumeName] |  
volumes ["volumeName1" ... "volumeNameN"]) reservations
```

Parameters

Parameter	Description
allVolumes	The setting to clear persistent volume reservations on all of the volumes in the storage array.
volume or volumes	The name of the specific volume for which you are retrieving information. Enclose the volume name in square brackets ([]). If the volume name has special characters, you also must enclose the volume name in double quotation marks (" ") inside square brackets. If the volume name consists only of numbers, such as "1002," you also must enclose the volume name in double quotation marks (" ") inside square brackets. You can enter more than one volume name. Enclose all of the volume names in one set of square brackets ([]). Enclose each volume name in double quotation marks (" "). Separate each volume name with a white space.

Notes

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

Minimum Firmware Level

5.40

Clear Volume Unreadable Sectors

This command clears unreadable sector information from one or more volumes.

Syntax

```
clear (allVolumes | volume [volumeName] |  
volumes ["volumeName1" ... "volumeNameN"]) unreadableSectors
```

Parameters

Parameter	Description
allVolumes	The setting to clear unreadable sector information from all of the volumes in the storage array.
volume or volumes	The name of the specific volume for which you are retrieving information. Enclose the volume name in square brackets ([]). If the volume name has special characters, you also must enclose the volume name in double quotation marks (" ") inside square brackets. If the volume name consists only of numbers, such as "1002," you also must enclose the volume name in double quotation marks (" ") inside square brackets. You can enter more than one volume name. Enclose all of the volume names in one set of square brackets ([]). Enclose each volume name in double quotation marks (" "). Separate each volume name with a white space.

Notes

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

Minimum Firmware Level

6.10

Configure Automatic Support Bundle Collection

NOTE This command is an SMcli command, not a script command. You must run this command from a command line. You cannot run this command from the script editor in the storage management software

This command enables or disables the automatic collection of support bundles on one or more storage arrays when a critical MEL event is detected.

Syntax

```
SMcli -supportBundle auto (enable|disable)  
(all | storageArrayName)  
[data=fileName]
```

Parameters

Parameter	Description
enable	Use this parameter to enable automatic collection of support bundles when a critical MEL event is detected.
disable	Use this parameter to disable automatic collection of support bundles where the feature was previously enabled.
all	Use this parameter to apply the command to all storage arrays detected from the host.
storageArrayName	Apply the command to the named storage array.
data	The file path and the file name to which you want to save the support bundle information.

Notes

If automatic support bundle collection is enabled on all storage arrays, storage arrays that are discovered later will have their support bundles collected. If automatic support bundle is enabled for all storage arrays and later automatic support bundle collection is disabled on one or more storage arrays, then storage arrays that are discovered later will not have their support bundles collected.

Minimum Firmware Level

7.83

Convert Snapshot (Legacy) Volume to Snapshot Group

This command migrates from a snapshot (legacy) volume to a snapshot group, and converts from an existing snapshot (legacy) repository volume to a snapshot image repository volume. A maximum of four snapshots (legacy) can be converted to a snapshot group. The snapshots (legacy) must be in a Disabled state and must not be involved in a volume copy operation. The base volume can have only four snapshots (legacy). If the base volume has more than four snapshots (legacy), you must delete the extra snapshots (legacy) before running this command.

NOTE You must disable any snapshots (legacy) before attempting to convert from a snapshot (legacy) volume to a snapshot group. Use the `disableSnapshot` command to disable a snapshot (legacy).

Syntax

```
convert snapshotVolume baseVolume="baseVolumeName"
```

Parameter

Parameter	Description
baseVolume	The name of the base volume that has the snapshots (legacy) that you want to convert into a snapshot group. Enclose the base volume name in double quotation marks (" ").

Notes

You can use any combination of alphanumeric characters, underscore (_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

Minimum Firmware Level

7.83

Create Asynchronous Mirror Group

This command creates a new, empty asynchronous mirror group on both the local storage array and the remote storage array. An asynchronous mirror group is a container that can house several mirrored pairs so that they can be managed as one entity. You create an asynchronous mirror group to define the synchronization settings for all mirrored pairs within the group. Each mirrored pair in an asynchronous mirror group share the same synchronization settings, primary and secondary role, and write mode.

The asynchronous mirror group is associated with the local storage array and remote storage array that is used for mirroring. The local storage array is the primary side of the asynchronous mirror group, while the remote storage array is the secondary side of the asynchronous mirror group. All volumes added to the asynchronous mirror group on the local storage array hold the primary role in the mirror relationship. Subsequently, all volumes added to the asynchronous mirror group on the remote storage array hold the secondary role in the mirror relationship.

Make sure that you execute the `Create Asynchronous Mirror Group` command on the local storage array. Asynchronous mirror group creation is initiated from the storage array that contains the volumes that hold the primary role in the mirror relationship. You use the `Create Asynchronous Mirror Group` command to specify the identity of the remote storage array that contains the volumes that hold the secondary role in the mirror relationship.

Syntax

```
create asyncMirrorGroup userLabel="asyncMirrorGroupName"  
(remoteStorageArrayName="storageArrayName" |  
remoteStorageArrayWwn="wwID")  
interfaceType=(FC | iSCSI)  
[remotePassword="password"  
syncInterval=integer (minutes | hours | days)  
warningSyncThreshold=integer (minutes | hours | days)  
warningRecoveryThreshold=integer (minutes | hours | days)  
warningThresholdPercent=percentValue  
autoResync=(TRUE | FALSE)]
```

Parameters

Parameter	Description
userLabel	<p>The name of the new asynchronous mirror group that you want to create. Enclose the new asynchronous mirror group name in double quotation marks (" ").</p> <p>The name must be unique on the local and remote storage arrays.</p>
remoteStorageArrayName	<p>The name for the remote storage array on which you are mirroring the asynchronous mirror group. Enclose the storage array name in double quotation marks (" ").</p> <p>The remote storage array must have the same connection type as the local storage array.</p>
remoteStorageArrayNameWwn	<p>The World Wide Identifier (WWID) of the remote storage array on which you are mirroring the asynchronous mirror group. You can use the WWID instead of the storage array name to identify the storage array. Enclose the WWID in angle brackets (<>).</p> <p>The remote storage array must have the same connection type as the local storage array.</p>
interfaceType	<p>Specify which connection type to use: either Fibre Channel fabric or iSCSI interface. (The default is Fibre Channel.) The local storage array and the remote storage array must be connected through a proper Fibre Channel fabric or iSCSI interface.</p>
remotePassword	<p>The password for the remote storage array. Use this parameter when the remote storage array is password protected. Enclose the password in double quotation marks (" ").</p>

Parameter	Description
syncInterval	Specify the length of time between automatically sending updates of modified data from the local storage array to the remote storage array. You can specify the length of time in minutes, hours, or days.
warningSyncThreshold	Specify the length of time to wait until a warning is triggered for cases in which the synchronization of all of the volumes within the asynchronous mirror group takes longer than the defined time. You can specify the length of time in minutes, hours, or days.
warningRecoveryThreshold	<p>Specify the length of time to wait until a warning is triggered when the automatic data update for the point-in-time image on the remote storage array is older than the defined time. Define the threshold from the end of the previous update. You can specify the length of time in minutes, hours, or days.</p> <hr/> <p>NOTE You must set the Recovery Point Threshold to be twice as large as the synchronization interval threshold.</p> <hr/>
warningThresholdPercent	Specify the length of time to wait until a warning is triggered when the capacity of a mirror repository volume reaches the defined percentage. Define the threshold by percentage (%) of the capacity remaining.
autoResync	<p>The settings for automatic resynchronization between the primary volumes and the secondary volumes of an asynchronous mirrored pair within an asynchronous mirror group. This parameter has these values:</p> <ul style="list-style-type: none"> ■ enabled– Automatic resynchronization is turned on. You do not need to do anything further to resynchronize the primary volume and the secondary volume. ■ disabled– Automatic resynchronization is turned off. To resynchronize the primary volume and the secondary volume, you must run the <code>resume asyncMirrorGroup</code> command.

Notes

- The Asynchronous Mirroring premium feature must be enabled and activated on the local and remote storage arrays that will be used for mirror activities.
- You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

- The local and remote storage arrays must be connected through a Fibre Channel fabric or iSCSI interface.
- Passwords are stored on each storage array in a management domain. If a password was not previously set, you do not need a password. The password can be any combination of alphanumeric characters with a maximum of 30 characters. (You can define a storage array password by using the `set storageArray` command.)
- Depending on your configuration, there is a maximum number of asynchronous mirror groups you can create on a storage array.
- Asynchronous mirror groups are created empty and mirrored pairs are added to them later on. Only mirrored pairs can be added to an asynchronous mirror group. Each mirrored pair is associated with exactly one asynchronous mirror group.
- The Asynchronous Mirroring process is initiated at a defined synchronization interval. Periodic point-in-time images are replicated as only the changed data is copied and not the entire volume.

Minimum Firmware Level

7.84

Create Consistency Group

NOTE This command does not apply to the snapshot (legacy) commands.

This command creates a new, empty consistency group that can contain snapshot groups. You must add the snapshot groups using the `set consistencyGroup addCGMember` command.

Syntax

```
create consistencyGroup userLabel="consistencyGroupName"
[repositoryFullPolicy=(failBaseWrites | purgeSnapImages) |
repositoryFullLimit=percentValue |
autoDeleteLimit=numberOfSnapImages |
enableSchedule=(TRUE | FALSE) |
schedule (immediate | snapSchedule) |
rollbackPriority=(lowest | low | medium | high | highest)]
```

Parameters

Parameter	Description
userLabel	The name of the new consistency group that you want to create. Enclose the new consistency group name in double quotation marks (" ").

Parameter	Description
repositoryFullPolicy	How you want snapshot processing to continue if the snapshot repository volumes are full. You can choose to fail writes to the base volume (<code>failBaseWrites</code>) or delete (purge) the snapshot images (<code>purgeSnapImages</code>). The default action is <code>purgeSnapImages</code> .
repositoryFullLimit	The percentage of repository capacity at which you receive a warning that the snapshot repository volume is nearing full. Use integer values. For example, a value of 70 means 70 percent. The default value is 75.
autoDeleteLimit	Each snapshot group can be configured to perform automatic deletion of its snapshot images to keep the total number of snapshot images in the snapshot group at or below a designated level. When this option is enabled, then any time a new snapshot image is created in the snapshot group, the system automatically deletes the oldest snapshot image in the group to comply with the limit value. This action frees repository capacity so it can be used to satisfy ongoing copy-on-write requirements for the remaining snapshot images.
enableSchedule	Whether the ability to schedule a snapshot operation is turned on or turned off. To turn on snapshot scheduling, set this parameter to <code>TRUE</code> . To turn off snapshot scheduling, set this parameter to <code>FALSE</code> .
schedule	<p>Use this parameter to schedule a snapshot operation.</p> <p>You can use one of these options for setting a schedule for a snapshot operation:</p> <ul style="list-style-type: none"> ■ <code>immediate</code> ■ <code>startDate</code> ■ <code>scheduleDay</code> ■ <code>startTime</code> ■ <code>scheduleInterval</code> ■ <code>endDate</code> ■ <code>timesPerDay</code> ■ <code>timeZone</code> ■ <code>scheduleDate</code> ■ <code>month</code> <p>See the "Notes" section for information explaining how to use these options.</p>

Parameter	Description
rollBackPriority	Determines whether system resources should be allocated to the rollback operation at the expense of system performance. A value of <code>high</code> indicates that the rollback operation is prioritized over all other host I/O. A value of <code>low</code> indicates that the rollback operation should be performed with minimal impact to host I/O.

Notes

A consistency group is a logical entity that enables you to manage in batch form all of the snapshot images that you add to the collection. The consistency group is a collection of snapshot groups that have mutual consistency requirements or dependencies for their snapshot images. Any snapshot images that you create and use for this collection must be managed in accordance with the consistency dependencies.

You can use any combination of alphanumeric characters, underscore (`_`), hyphen (`-`), and pound (`#`) for the names. Names can have a maximum of 30 characters.

The snapshot images in a consistency group can be deduced based on the existence of a snapshot image within a consistency group. All snapshot images that reside within a consistency group share a common time stamp and sequence number.

An operation on a snapshot image consistency group is treated as a single request, and causes all pending I/O operations to the associated base volume of each member to be drained and suspended before creating the snapshot images. If creation of the snapshot images cannot be completed successfully for all of the consistency group members, the operation fails and has no effect (that is, new snapshot images are not created).

Based on this behavior all members for a consistency group usually have the same number of snapshot images. However, when a new member is added to a consistency group, that new member lacks the snapshot images that were previously created on the established members of the consistency group. The lack of snapshot images is not considered an error condition. Ensuing requests for deletion or rollback of snapshot images that only exist on a subset of the consistency group members will only affect the members for which the specified snapshot images actually exist.

Auto Delete

You can configure each snapshot group to automatically delete its snapshot images to keep the total number of snapshot images in the snapshot group at or below a maximum number of images. When the number of snapshot images in the snapshot group is at the maximum limit, the `autoDeleteLimit` parameter automatically deletes snapshot images whenever a new snapshot image is created in the snapshot group. The `autoDeleteLimit` parameter deletes the oldest snapshot images in the snapshot group until the maximum number of images defined with the parameter is met. This has the effect of freeing repository capacity so it can be used to satisfy ongoing copy-on-write requirements for the remaining snapshot images.

Scheduling Snapshot Images in a Consistency Group

The `enableSchedule` parameter and the `schedule` parameter provide a way for you to schedule automatic snapshots (legacy). Using these parameters, you can schedule snapshots daily, weekly, or monthly (by day or by date). The `enableSchedule` parameter turns on or turns off the ability to schedule snapshots. When you enable scheduling, you use the `schedule` parameter to define when you want the snapshots to occur.

This list explains how to use the options for the `schedule` parameter:

- `immediate` – As soon as you enter the command, a snapshot volume is created and a copy-on-write operation begins.
- `startDate` – A specific date on which you want to create a snapshot volume and perform a copy-on-write operation. The format for entering the date is `MM:DD:YY`. If you do not provide a start date, the current date is used. An example of this option is `startDate=06:27:11`.
- `scheduleDay` – A day of the week on which you want to create a snapshot volume and perform a copy-on-write operation. You can enter these values: `monday`, `tuesday`, `wednesday`, `thursday`, `friday`, `saturday`, `sunday`, and `all`. An example of this option is `scheduleDay=wednesday`.
- `startTime` – The time of a day that you want to create a snapshot volume and start performing a copy-on-write operation. The format for entering the time is `HH:MM`, where `HH` is the hour and `MM` is the minute past the hour. Use a 24-hour clock. For example, 2:00 in the afternoon is 14:00. An example of this option is `startTime=14:27`.
- `scheduleInterval` – An amount of time, in minutes, that you want to have as a minimum between copy-on-write operation. You can create a schedule in which you have overlapping copy-on-write operations because of the duration of a copy operation. You can make sure that you have time between copy-on-write operations by using this option. The maximum value for the `scheduleInterval` option is 1440 minutes. An example of this option is `scheduleInterval=180`.
- `endDate` – A specific date on which you want to stop creating a snapshot volume and end the copy-on-write operations. The format for entering the date is `MM:DD:YY`. An example of this option is `endDate=11:26:11`.
- `timesPerDay` – The number of times that you want the schedule to run in a day. An example of this option is `timesPerDay=4`.

- `timeZone` – Use this parameter to define the time zone in which the storage array is operating. You can define the time zone in one of two ways:
 - `GMT±HH:MM` – The time zone offset from GMT. Enter the offset in hours and minutes. For example `GMT-06:00` is the central time zone in the United States.
 - `Text string` – Standard time zone text strings. For example: `"USA/Chicago"` or `"Germany/Berlin"`. Time zone text strings are case sensitive. If you enter an incorrect text string, GMT time is used. Enclose the text string in double quotation marks.
- `scheduleDate` – A day of the month on which you want to create a snapshot volume and perform a copy-on-write operation. The values for the days are numerical and in the range of 1–31. Enclose the value for the day in double quotation marks inside parentheses. If you want to create a snapshot volume on a specific date of a specific month use the `scheduleDate` option with the `month` option. If you have set up a weekly schedule, you cannot use the `scheduleDate` option. An example of the `scheduleDate` option is `scheduleDate=("15")`.
- `month` – A specific month on which you want to create a snapshot volume and perform a copy-on-write operation. The values for the months are: `jan`, `feb`, `mar`, `apr`, `may`, `jun`, `jul`, `aug`, `sep`, `oct`, `nov`, and `dec`. Enclose the value in parentheses. You can enter more than one month by enclosing the months in a single set of parentheses and separating each month with a white space. If you want to create a snapshot volume on a specific date of a specific month use the `scheduleDate` option with the `month` option. These are examples of the `month` option:
 - `month=(mar)`
 - `month=(mar apr may)`

The code string for defining a schedule is similar to these examples:

```
enableSchedule=true schedule startTime=14:27
enableSchedule=true schedule scheduleInterval=180
enableSchedule=true schedule timeZone=GMT-06:00
enableSchedule=true schedule timeZone="USA/Chicago"
enableSchedule=true schedule month=(mar) scheduleDate=("15")
```

If you also use the `scheduleInterval` option, the firmware chooses between the `timesPerDay` option and the `scheduleInterval` option by selecting the lowest value of the two options. The firmware calculates an integer value for the `scheduleInterval` option by dividing 1440 by the `scheduleInterval` option value that you set. For example, $1440/180 = 8$. The firmware then compares the `timesPerDay` integer value with the calculated `scheduleInterval` integer value and uses the smaller value.

To remove a schedule, use the `delete volume` command with the `schedule` parameter. The `delete volume` command with the `schedule` parameter deletes only the schedule, not the snapshot volume.

When performing a rollback in a consistency group, the default operation is to rollback all members of the consistency group. If a rollback cannot be started successfully for all of the members in the consistency group, the rollback fails and has no effect. The snapshot image is not rolled back.

Minimum Firmware Level

7.83

7.86 adds the `scheduleDate` option and the `month` option.

Create Consistency Group Snapshot Image

This command creates a new snapshot image for each base volume that is a member of a snapshot consistency group.

Syntax

```
create cgSnapImage consistencyGroup="consistencyGroupName"
```

Parameter

Parameter	Description
<code>consistencyGroup</code>	The name of the consistency group for which you are creating snapshot images. Enclose the consistency group name in double quotation marks (" ").

Notes

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

The command causes all pending I/O operations to each base volume that is a member of the consistency group to be drained and suspended before creating the snapshot images. If the creation of all of the snapshot images cannot be completed successfully for all of the consistency group members, the operation fails and new snapshot images are not created.

Normally, all members of a snapshot consistency group have the same number of snapshot images. When you add a new member to a snapshot consistency group, that new member lacks the snapshot images that were previously created on the established members of the snapshot consistency group. This is not an error condition. Requests for deletion or rollback of snapshot images that exist on only a subset of the snapshot consistency group members affects only those members for which the specified snapshot image actually exists.

Minimum Firmware Level

7.83

Create Consistency Group Snapshot Volume

This command creates a snapshot volume of specific images in the base volumes in a consistency group. You can select one base volume or more than one base volumes from the consistency group to include in the snapshot volume. When you create a snapshot volume of a consistency group you are creating a volume with contents that you can view.

Syntax With User Specified Base Volumes

```
create cgSnapVolume userLabel="cgVolumeName"  
cgSnapImageID="snapCGID:imageID"  
members=(baseVolume1:repos_XXXX ... baseVolumen:repos_YYYY)
```

Syntax When Setting the Consistency Group Snapshot Volume to Read Only

```
create cgSnapVolume userLabel="cgVolumeName"  
cgSnapImageID="snapCGID:imageID"  
readOnly
```

Syntax When Setting the Repository Full Limit

```
create cgSnapVolume userLabel="cgVolumeName"  
cgSnapImageID="snapCGID:imageID"  
members=(baseVolume1:repos_XXXX | baseVolume1:  
(volumeGroupName [capacity=capacityValue]) | (  
baseVolume1:diskPoolName  
[capacity=capacityValue]) ... baseVolumen:repos_YYYY |  
baseVolumen:  
(volumeGroupName [capacity=capacityValue]) | baseVolumen:  
(diskPoolName  
[capacity=capacityValue])) repositoryFullLimit=percentValue
```

Parameters

Parameter	Description
userLabel	The name that you want to give the consistency group snapshot volume that you are creating. Enclose the consistency group snapshot volume name in double quotation marks (" ").

Parameter	Description
cgSnapImageID	<p>The name of the snapshot image in a consistency group. The name of a snapshot image is comprised of two parts:.</p> <ul style="list-style-type: none"> ■ The name of the consistency group ■ An identifier for the snapshot image in the consistency group. <p>The identifier for the snapshot image can be one of these:</p> <ul style="list-style-type: none"> ■ NEWEST - Use this option when you want to show the latest snapshot image created in the consistency group. ■ OLDEST - Use this option when you want to show the earliest snapshot image created in the consistency group. <p>Enclose the snapshot image name in double quotation marks (" ") inside square brackets ([]).</p> <p>You can enter more than one snapshot image name or sequence number. Enclose all of the snapshot image names in one set of double quotation marks ("") inside square brackets ([]). Separate each snapshot image name with a white space.</p>
members	<p>The identifier of one base volume or more than one base volumes that you want to add. The members identifier is comprised of the base volume name concatenated with the repository volume name. You must use a colon (:) between the two names. Enclose all of the member identifiers in parentheses. If you enter more than one member separated the members with a space.</p> <p>When you do not use the members parameter, all of the members are automatically added to the new consistency group snapshot volume.</p>

Parameter	Description
repositoryVolume	<p>The name of the repository volume that will contain the consistency group member volumes.</p> <p>You have two options for defining the name of a repository volume:</p> <ul style="list-style-type: none"> ■ Use an existing repository volume: name ■ Create a new repository volume when you run this command <p>The name of an existing repository volume is comprised of two parts:</p> <ul style="list-style-type: none"> ■ The term <i>repos</i> ■ A four digit numerical identifier that the storage management software assigns to the repository volume name <p>Enclose the name of the existing repository volume in double quotation marks (" ").</p> <p>If you want to create a new repository volume when you run this command you must enter the name of either a a volume group or a disk pool in which you want the repository volume. Optionally, you can also define the capacity of the repository volume. If you want to define the capacity you can use these values:</p> <ul style="list-style-type: none"> ■ An integer value that represents a percentage of the base volume capacity ■ A decimal fraction value that represents a percentage of the base volume capacity ■ A specific size for the repository volume. Size is defined in units of <code>bytes</code>, KB, MB, GB, or TB. <p>If you do not use the capacity option, the storage management software sets the capacity to 20 percent of the base volume capacity.</p> <p>When you run this command the storage management software creates the repository volume for the snapshot volume.</p>
repositoryFullLimit	<p>The percentage of repository capacity at which the consistency group snapshot repository volume is nearing full. Use integer values. For example, a value of 70 means 70 percent.</p>
readOnly	<p>The setting to determine whether you can write to the snapshot volume or only read from the snapshot volume. To write to the snapshot volume, do not include this parameter. To prevent writing to the snapshot volume, include this parameter.</p>

Notes

You can use any combination of alphanumeric characters, underscore (`_`), hyphen (`-`), and pound (`#`) for the names. Names can have a maximum of 30 characters.

The name of a snapshot image has two parts separated by a colon (`:`):

- The identifier of the snapshot group
- The identifier of the snapshot image

If you do not specify the `repositoryVolumeType` or `readOnly` parameters, the storage management software selects the repositories for the consistency group snapshot volume. If the volume group or disk pool where the base volume resides does not have enough space, this command fails.

The `create cgSnapVolume` command has unique forms that are explained by these examples:

- Creating a read/write consistency group snapshot volume on a snapshot consistency group named "snapCG1" that has three members `cgm1`, `cgm2`, and `cgm3`. The repository volumes already exist and selected by the user in this command.

```
create cgSnapVolume userLabel="cgSnapVolume1"
cgSnapImageID="snapCG1:oldest"
members=(cgm1:repos_0010 cgm2:repos_0011
cgm3:repos_0007);
```

Note the use of the colon (`:`) in the name of the snapshot image to be included in the consistency group snapshot volume. The colon is a delimiter that separates the name of the snapshot volume from a particular snapshot image that you might want to use. You can use one of these options following the colon:

- An integer value that is the actual sequence number of the snapshot image.
- `newest` - Use this option when you want to show the latest consistency group snapshot image.
- `oldest` - Use this option when you want to show the earliest snapshot image created.

The use of the colon following the names of the members of the snapshot consistency group acts define the mapping between the member and a repository volume. For example, in `cgm1:repos_10`, member `cgm1` maps to repository volume `repos_0010`.

- Creating a read/write consistency group snapshot volume on a snapshot consistency group named "snapCG1" of only members `cgm1` and `cgm2`:

```
create cgSnapVolume userLabel="cgSnapVolume2"
cgSnapImageID="snapCG1:14214"
members=(cgm1:repos_1000 cgm2:repos_1001);
```

- Creating a read-only consistency group snapshot volume on a snapshot consistency group named snapCG1 that has three members cgm1, cgm2, and cgm3:

```
create cgSnapVolume userLabel="cgSnapVolume3"
cgSnapImageID="snapCG1:oldest" readOnly;
```

- Creating a consistency group snapshot volume that has a repository full limit set to 60 percent on a snapshot consistency group named snapCG1 that has three members cgm1, cgm2, and cgm3:

```
create cgSnapVolume userLabel="cgSnapVolume3"
cgSnapImageID="snapCG1:oldest"
repositoryFullLimit=60;
```

- Creating a read/write consistency group snapshot volume with automatic repository selection on a snapshot consistency group named snapCG1 that has three members cgm1, cgm2, and cgm3:

```
create cgSnapVolume userLabel="cgSnapVolume4"
cgSnapImageID="snapCG1:oldest" ;
```

Minimum Firmware Level

7.83

Create Consistency Group Snapshot Volume Mapping

This command creates a logical mapping from a consistency group snapshot volume to a host or a host group.

Syntax

```
create mapping cgSnapVolume="snapVolumeName"
(host="hostName" | hostGroup=("hostGroupName" |
defaultGroup)
```

Parameters

Parameter	Description
cgSnapVolume	The name of the consistency group snapshot volume for which you want to create a logical mapping. Enclose the consistency group snapshot volume name in double quotation marks (" ").
host	The name of a host to which you want to create a logical mapping. Enclose the host name in double quotation marks (" ").
hostGroup	The name of a host group to which you want to create a logical mapping. Enclose the host group name in double quotation marks (" "). If you use the defaultGroup keyword, do not enclose it in quotation marks.

Notes

A host is a computer that is attached to the storage array and accesses the volumes on the storage array through the host ports. You can define specific mappings to an individual host. You also can assign the host to a host group that shares access to one or more volumes.

A host group is an optional topological element that you can define if you want to designate a collection of hosts that share access to the same volumes. The host group is a logical entity. Define a host group only if you have two or more hosts that share access to the same volumes.

Minimum Firmware Level

7.83

Create Disk Pool

This command creates a new disk pool based on the specified parameters. You can create the disk pool by entering either a list of drives or a type of drive that you want to use for the disk pool.

NOTE If you enter a list of drives, make sure that all of the drives have the same capacity. If the drives do not have the same capacity each drive in the disk pool reports capacity equal to the smallest drive.

Syntax

```
create diskPool
(drives=(trayID1,drawerID1,slotID1 ...
trayIDN,drawerIDN,slotIDN) |
driveType=(fibre|sas|sata))
userLabel="diskPoolName"
[driveCount=driveCountValue |
warningThreshold=(warningThresholdValue|default) |
criticalThreshold=(criticalThresholdValue|default) |
criticalPriority=(highest|high|medium|low|lowest) |
backgroundPriority=(highest|high|medium|low|lowest) |
degradedPriority=(highest|high|medium|low|lowest) |
securityType=(none|capable|enabled) |
driveMediaType=(hdd | ssd | allMedia | unknown) |
dataAssurance=(none|enabled)]
```

Parameters

Parameter	Description
<code>drives</code>	The drives that you want to assign to the disk pool that you want to create. For high capacity drive trays, specify the tray ID value, the drawer ID value, and the slot ID value for each drive that you assign to the disk pool. For low-capacity drive trays, specify the tray ID value and the slot ID value for each drive that you assign to the disk pool. Tray ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the tray ID values, the drawer ID values, and the slot ID values in parentheses.
<code>driveType</code>	<p>The type of drive that you want to use in the disk pool. You cannot mix drive types. When you use the <code>driveType</code> parameter all of the unassigned drives of that type that are compatible with each other are used for the creation of the disk pool.</p> <p>You must use this parameter when you have more than one type of drive in your storage array.</p> <p>Valid drive types are:</p> <ul style="list-style-type: none"> ■ fibre ■ SATA ■ SAS <p>If you do not specify a drive type, the command defaults to <code>fibre</code>.</p>
<code>userLabel</code>	The name that you want to give the new disk pool. Enclose the disk pool name in double quotation marks (" ").
<code>driveCount</code>	The <code>driveCount</code> parameter limits the disk pool candidates to the given number. When using this parameter, the minimum value that you can enter is 11. Use this parameter only when the <code>driveType</code> parameter is specified.
<code>warningThreshold</code>	<p>The percentage of storage capacity at which you receive a warning alert that the disk pool is nearing full. Use integer values. For example, a value of 60 means 60 percent. For best operation, the value for this parameter must be less than the value for the <code>criticalThreshold</code> parameter.</p> <p>Valid values are from 0 to 100.</p> <p>The default value is 50.</p> <p>Setting this parameter to 0 disables warning alerts.</p> <p>If you set this to <code>default</code>, the warning alert threshold value is determined by the controller firmware.</p>

Parameter	Description
criticalThreshold	<p>The percentage of storage capacity at which you receive a critical alert that the disk pool is nearing full. Use integer values. For example, a value of 70 means 70 percent. For best operation, the value for this parameter must be greater than the value for the warningThreshold parameter.</p> <p>Valid values are from 0 to 100.</p> <p>The default value is 85 percent.</p> <p>Setting this parameter to 0 disables both warning alerts and critical alerts.</p> <p>If you set this to default, the critical alert threshold value is determined by the controller firmware.</p>
criticalPriority	<p>The priority for reconstruction operations for critical events on the disk pool. For example, disk pool reconstruction after at least two drive failures.</p> <p>Valid values are highest, high, medium, low, and lowest. The default value is highest.</p>
backgroundPriority	<p>The priority for background operations on the disk pool.</p> <p>Valid values are highest, high, medium, low, and lowest. The default value is low.</p>
degradedPriority	<p>The priority for degraded activities on the disk pool. For example, disk pool reconstruction after one drive failures.</p> <p>Valid values are highest, high, medium, low, and lowest. The default value is high.</p>
securityType	<p>The setting to specify the security level when creating the disk pool. All volume candidates for the disk pool will have the specified security type. These settings are valid:</p> <ul style="list-style-type: none"> ■ none – The volume candidates are not secure. ■ capable – The volume candidates are capable of having security set, but security has not been enabled. ■ enabled – The volume candidates have security enabled. <p>The default value is none</p>

Parameter	Description
driveMediaType	<p>The type of drive media that you want to use for the disk pool</p> <p>You must use this parameter when you have more than one type of drive media in your storage array.</p> <p>Valid drive media are:</p> <ul style="list-style-type: none"> ■ hdd – Use this option when you have hard drives. ■ ssd – Use this option when you have solid state disks. ■ unknown – Use if you are not sure what types of drive media are in the drive tray. ■ allMedia – Use this option when you want to use all types of drive media that are in the drive tray. <p>The default value is hdd</p>
dataAssurance	<p>The setting to specify that a disk pool has data assurance protection to make sure that the data maintains its integrity. When you use this parameter, only protected drives can be used for the disk pool. These settings are valid:</p> <ul style="list-style-type: none"> ■ none – The disk pool does not have data assurance protection. ■ enabled – The disk pool has data assurance protection. The disk pool supports protected information and is formatted with protection information enabled. <p>The default value is none</p>

Notes

Each disk pool name must be unique. You can use any combination of alphanumeric characters, underscore (_), hyphen(-), and pound (#) for the user label. User labels can have a maximum of 30 characters.

If the parameters you specify cannot be satisfied by any of the available candidate drives, the command fails. Normally, all drives that match the quality of service attributes are returned as the top candidates. However, if you specifying a drive list, some of the available drives returned as candidates might not match the quality of service attributes.

If you do not specify a value for an optional parameter, a default value is assigned.

Drives

When you use the `driveType` parameter, all of the unassigned drives that are of that drive type are used to create the disk pool. If you want to limit the number of drives found by the `driveType` parameter in the disk pool, you can specify the number of drives using the `driveCount` parameter. You can use the `driveCount` parameter only when you use the `driveType` parameter.

The `drives` parameter supports both high-capacity drive trays and low-capacity drive trays. A high-capacity drive tray has drawers that hold the drives. The drawers slide out of the drive tray to provide access to the drives. A low-capacity drive tray does not have drawers. For a high-capacity drive tray, you must specify the identifier (ID) of the drive tray, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive tray, you need only specify the ID of the drive tray and the ID of the slot in which a drive resides. For a low-capacity drive tray, an alternative method for identifying a location for a drive is to specify the ID of the drive tray, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

If you enter specifications for a high-capacity drive tray, but a drive tray is not available, the storage management software returns an error message.

Disk Pool Alert Thresholds

Each disk pool has two progressively severe levels of alerts to inform users when the storage capacity of the disk pool is approaching full. The threshold for an alert is a percent of the used capacity to the total usable capacity in the disk pool. The alerts are:

- Warning – This is the first level of alert and indicates that the used capacity in a disk pool is approaching full. When the threshold for the warning alert is reached, a Needs Attention condition is generated and an event is posted to the storage management software. The warning threshold is superseded by the critical threshold. The default warning threshold is 50 percent.
- Critical – This is the most severe level of alert and indicates that the used capacity in a disk pool is approaching full. When the threshold for the critical alert is reached, a Needs Attention condition is generated and an event is posted to the storage management software. The warning threshold is superseded by the critical threshold. The default threshold for the critical alert is 85 percent.

To be effective, the value for a warning alert must always be less than the value for a critical alert. If the value for the warning alert is the same as the value for a critical alert, only the critical alert is sent.

Disk Pool Background Operations

Disk pools support these background operations:

- Reconstruction
- Instant Availability Format (IAF)
- Format

- Dynamic Capacity Expansion (DCE)
- Dynamic Volume Expansion (DVE) (For disk pools, DVE is actually not a background operation, but DVE is supported as a synchronous operation.)

Disk pools do not queue background commands. You can start several background commands sequentially, but starting more than one background operation at a time delays the completion of commands that you started previously. The relative priority levels for the supported background operations are:

1. Reconstruction
2. Format
3. IAF
4. DCE

Security Type

The `securityType` parameter is valid for drives that are capable of full disk encryption (FDE). With FDE, the controller firmware can create a key and activate the Drive Security feature. The Drive Security feature encrypts data as the data is written to the drive and decrypts the data as the data is read from the drive. Without the key created by the controller, the data written to the drive is inaccessible.

Before you can set the `securityType` parameter to `capable` or `enabled`, you must create a storage array security key. Use the `create storageArray securityKey` command to create a storage array security key. These commands are related to the security key:

- `create storageArray securityKey`
- `enable diskPool [diskPoolName] security`
- `export storageArray securityKey`
- `import storageArray securityKey`
- `set storageArray securityKey`
- `start secureErase (drive | drives)`

Minimum Firmware Level

7.83

Create Host

This command creates a new host. If you do not specify a host group in which to create the new host, the new host is created in the Default Group.

Syntax

```
create host userLabel="hostName"
[hostGroup=( "hostGroupName" | defaultGroup )]
[hostType=(hostTypeIndexLabel | hostTypeIndexNumber )]
```

Parameters

Parameter	Description
userLabel	The name that you want to give the host that you are creating. Enclose the host name in double quotation marks (" ").
hostGroup	The name of the host group in which you want to create a new host. Enclose the host group name in double quotation marks (" "). (If a host group does not exist, you can create a new host group by using the <code>create hostGroup</code> command.) The <code>defaultGroup</code> option is the host group that contains the host to which the volume is mapped.
hostType	The index label or the index number that identifies the host type. Use the <code>show storageArray hostTypeTable</code> command to generate a list of available host type identifiers. If the host type has special characters, enclose the host type in double quotation marks (" ").

Notes

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

A host is a computer that is attached to the storage array and accesses the volumes on the storage array through the host ports. You can define specific mappings to an individual host. You also can assign the host to a host group that shares access to one or more volumes.

A host group is an optional topological element that you can define if you want to designate a collection of hosts that share access to the same volumes. The host group is a logical entity. Define a host group only if you have two or more hosts that share access to the same volumes.

If you do not specify a host group in which to place the host that you are creating, the newly defined host belongs to the default host group.

Minimum Firmware Level

5.20

7.10 adds the `hostType` parameter.

Create Host Group

This command creates a new host group.

Syntax

```
create hostGroup userLabel="hostGroupName"
```

Parameter

Parameter	Description
userLabel	The name that you want to give the host group that you are creating. Enclose the host name in double quotation marks (" ").

Notes

A host group is an optional topological element that you can define if you want to designate a collection of hosts that share access to the same volumes. The host group is a logical entity. Define a host group only if you have two or more hosts that can share access to the same volumes.

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

Minimum Firmware Level

5.20

Create Host Port

This command creates a new host port identification on a host bus adapter (HBA) or on a host channel adapter (HCA). The identification is a software value that represents the physical HBA or HCA host port to the controller. Without the correct host port identification, the controller cannot receive instructions or data from the host port.

Syntax

```
create hostPort identifier=("wwID" | "gid")
userLabel="portLabel"
host="hostName"
interfaceType=(FC | SAS | IB)
```

Parameters

Parameter	Description
identifier	The 8-byte World Wide Identifier (WWID) or the 16-byte group identifier (GID) of the HBA or HCA host port. Enclose the WWID or the GID in double quotation marks (" ").
userLabel	The name that you want to give to the new HBA or HCA host port. Enclose the host port label in double quotation marks (" ").
host	The name of the host for which you are defining an HBA or HCA host port. Enclose the host name in double quotation marks (" ").

Parameter	Description
interfaceType	<p>The identifier of the type of interface for the host port.</p> <p>The choices for the types of host port interfaces are:</p> <ul style="list-style-type: none"> ■ FC – Fibre Channel ■ SAS – Serial-Attached SCSI ■ IB – Infiniband <p>An FC or a SAS selection requires an 8-byte WWID. An IB selection requires a 16-byte group identifier (gid).</p> <p>If you do not specify the type of interface, FC is used as the default interface for the host port.</p>

Notes

An HBA host port or an HCA host port is a physical connection on a host bus adapter or on a host channel adapter that resides in a host computer. An HBA host port or an HCA host port provides host access to the volumes in a storage array. If the HBA or the HCA has only one physical connection (one host port), the terms host port and host bus adapter or host channel adapter are synonymous.

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

Minimum Firmware Level

5.20

7.10 deprecates the `hostType` parameter. The `hostType` parameter has been added to the `create host` command.

7.32 adds the `interfaceType` parameter.

Create iSCSI Initiator

This command creates a new iSCSI initiator object.

Syntax

```
create iscsiInitiator iscsiName="iscsiID"
userLabel="name"
host="hostName"
[chapSecret="securityKey"]
```

Parameters

Parameters	Description
iscsiName	The default identifier of the iSCSI initiator. Enclose the identifier in double quotation marks (" ").

Parameters	Description
userLabel	The name that you want to use for the iSCSI initiator. Enclose the name in double quotation marks (" ").
host	The name of the host in which the iSCSI initiator is installed. Enclose the name in double quotation marks (" ").
chapSecret	The security key that you want to use to authenticate a peer connection. Enclose the security key in double quotation marks (" ").

Notes

Challenge Handshake Authentication Protocol (CHAP) is a protocol that authenticates the peer of a connection. CHAP is based upon the peers sharing a *secret*. A secret is a security key that is similar to a password.

Use the `chapSecret` parameter to set up the security keys for initiators that require a mutual authentication.

Minimum Firmware Level

7.10

Create RAID Volume (Automatic Drive Select)

This command creates a volume group across the drives in the storage array and a new volume in the volume group. The storage array controllers choose the drives to be included in the volume.

NOTE If you have drives with different capacities, you cannot automatically create volumes by specifying the `driveCount` parameter. If you want to create volumes with drives of different capacities, see "[Create RAID Volume \(Manual Drive Select\)](#)."

Syntax

```
create volume driveCount=numberOfDrives
volumeGroupUserLabel="volumeGroupName"
raidLevel=(0 | 1 | 3 | 5 | 6)
userLabel="volumeName"
driveMediaType=(HDD | SSD | unknown | allMedia)
[driveType=(fibre | SATA | SAS)
capacity=volumeCapacity
owner=(a | b)
cacheReadPrefetch=(TRUE | FALSE)
segmentSize=segmentSizeValue
usageHint=(fileSystem | dataBase | multiMedia)
trayLossProtect=(TRUE | FALSE)
drawerLossProtect=(TRUE | FALSE)
dssPreAllocate=(TRUE | FALSE)
securityType=(none | capable | enabled)
dataAssurance=(none | enabled)]
```

Parameters

Parameter	Description
driveCount	The number of unassigned drives that you want to use in the volume group.
volumeGroupUserLabel	The alphanumeric identifier (including – and _) that you want to give the new volume group. Enclose the new volume group name in double quotation marks (" ").
raidLevel	The RAID level of the volume group that contains the volume. Valid values are 0, 1, 3, 5, or 6.
userLabel	The name that you want to give to the new volume. Enclose the new volume name in double quotation marks (" ").
driveMediaType	The type of drive media that you want to use for the volume group. Valid drive media are these: <ul style="list-style-type: none">■ HDD – Use this option when you have hard drives in the drive tray.■ SSD – Use this option when you have solid state drives in the drive tray.■ unknown – Use if you are not sure what types of drive media are in the drive tray.■ allMedia – Use this option when you want to use all types of drive media that are in the drive tray.

Parameter	Description
driveType	<p>The type of drive that you want to use in the volume. You cannot mix drive types.</p> <p>You must use this parameter when you have more than one type of drive in your storage array.</p> <p>Valid drive types are :</p> <ul style="list-style-type: none"> ■ fibre ■ SATA ■ SAS <p>If you do not specify a drive type, the command defaults to fibre.</p>
capacity	<p>The size of the volume that you are adding to the storage array. Size is defined in units of bytes, KB, MB, GB, or TB.</p>
owner	<p>The controller that owns the volume. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. If you do not specify an owner, the controller firmware determines the owner.</p>
cacheReadPrefetch	<p>The setting to turn on or turn off cache read prefetch. To turn off cache read prefetch, set this parameter to FALSE. To turn on cache read prefetch, set this parameter to TRUE.</p>
segmentSize	<p>The amount of data (in KB) that the controller writes on a single drive in a volume before writing data on the next drive. Valid values are 8, 16, 32, 64, 128, 256, or 512.</p>
usageHint	<p>The setting for both cacheReadPrefetch parameter and the segmentSize parameter to be default values. The default values are based on the typical I/O usage pattern of the application that is using the volume. Valid values are fileSystem, dataBase, or multiMedia.</p>
trayLossProtect	<p>The setting to enforce tray loss protection when you create the volume group. To enforce tray loss protection, set this parameter to TRUE. The default value is FALSE.</p>
drawerLossProtect	<p>The setting to enforce drawer loss protection when you create the mirror repository volume group. To enforce drawer loss protection, set this parameter to TRUE. The default value is FALSE.</p>
dssPreAllocate	<p>The setting to make sure that reserve capacity is allocated for future segment size increases. The default value is TRUE.</p>

Parameter	Description
securityType	<p>The setting to specify the security level when creating the volume groups and all associated volumes. These settings are valid:</p> <ul style="list-style-type: none"> ■ none – The volume group and volumes are not secure. ■ capable – The volume group and volumes are capable of having security set, but security has not been enabled. ■ enabled – The volume group and volumes have security enabled.
dataAssurance	<p>The setting to specify that a volume group, and the volumes within the volume group, has data assurance protection to make sure that the data maintains its integrity. When you use this parameter, only protected drives can be used for the volume group. These settings are valid:</p> <ul style="list-style-type: none"> ■ none – The volume group does not have data assurance protection. ■ enabled – The volume group has data assurance protection. The volume group supports protected information and is formatted with protection information enabled.

Notes

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

The `driveCount` parameter lets you choose the number of drives that you want to use in the volume group. You do not need to specify the drives by tray ID and slot ID. The controllers choose the specific drives to use for the volume group.

The `owner` parameter defines which controller owns the volume.

If you do not specify a capacity using the `capacity` parameter, all of the drive capacity that is available in the volume group is used. If you do not specify capacity units, `bytes` is used as the default value.

Cache Read Prefetch

Cache read prefetch lets the controller copy additional data blocks into cache while the controller reads and copies data blocks that are requested by the host from the drives into cache. This action increases the chance that a future request for data can be fulfilled from cache. Cache read prefetch is important for multimedia applications that use sequential data transfers. The configuration settings for the storage array that you use determine the number of additional data blocks that the controller reads into cache. Valid values for the `cacheReadPrefetch` parameter are `TRUE` or `FALSE`.

Segment Size

The size of a segment determines how many data blocks that the controller writes on a single drive in a volume before writing data on the next drive. Each data block stores 512 bytes of data. A data block is the smallest unit of storage. The size of a segment determines how many data blocks that it contains. For example, an 8-KB segment holds 16 data blocks. A 64-KB segment holds 128 data blocks.

When you enter a value for the segment size, the value is checked against the supported values that are provided by the controller at run time. If the value that you entered is not valid, the controller returns a list of valid values. Using a single drive for a single request leaves other drives available to simultaneously service other requests.

If the volume is in an environment where a single user is transferring large units of data (such as multimedia), performance is maximized when a single data transfer request is serviced with a single data stripe. A data stripe is the segment size that is multiplied by the number of drives in the volume group that are used for data transfers. In this case, multiple drives are used for the same request, but each drive is accessed only once.

For optimal performance in a multiuser database or file system storage environment, set your segment size to minimize the number of drives that are required to satisfy a data transfer request.

You do not need to enter a value for the `cacheReadPrefetch` parameter or the `segmentSize` parameter. If you do not enter a value, the controller firmware uses the `usageHint` parameter with `fileSystem` as the default value. Entering a value for the `usageHint` parameter and a value for the `cacheReadPrefetch` parameter or a value for the `segmentSize` parameter does not cause an error. The value that you enter for the `cacheReadPrefetch` parameter or the `segmentSize` parameter takes priority over the value for the `usageHint` parameter.

Tray Loss Protection and Drawer Loss Protection

For tray loss protection to work, each drive in a volume group must be on a separate tray. If you set the `trayLossProtect` parameter to `TRUE` and have selected more than one drive from any one tray, the storage array returns an error. If you set the `trayLossProtect` parameter to `FALSE`, the storage array performs operations, but the volume group that you create might not have tray loss protection.

Tray loss protection is not valid when you create volumes on existing volume groups.

The `drawerLossProtect` parameter defines if data on a volume is accessible if a drawer fails. When you assign the drives, if you set the `drawerLossProtect` parameter to `TRUE` and select more than one drive from any one drawer, the storage array returns an error. If you set the `drawerLossProtect` parameter to `FALSE`, the storage array performs operations, but the volume group that you create might not have drawer loss protection.

You must set the `trayLossProtect` parameter and the `drawerLossProtect` parameter to the same value. Both of the parameters must be either `TRUE` or `FALSE`. If the `trayLossProtect` parameter and the `drawerLossProtect` parameter are set to different values, the storage array returns an error.

Security Type

The `securityType` parameter is valid for drives that are capable of full disk encryption (FDE). With FDE, the controller firmware can create a key and activate the Drive Security feature. The Drive Security feature encrypts data as the data is written to the drive and decrypts the data as the data is read from the drive. Without the key created by the controller, the data written to the drive is inaccessible.

Before you can set the `securityType` parameter to `capable` or `enabled`, you must create a storage array security key. Use the `create storageArray securityKey` command to create a storage array security key. These commands are related to the security key:

- `create storageArray securityKey`
- `set storageArray securityKey`
- `import storageArray securityKey`
- `export storageArray securityKey`
- `start secureErase (drive | drives)`
- `enable volumeGroup [volumeGroupName] security`

Minimum Firmware Level

5.20

7.10 adds RAID Level 6 capability and the `dssPreAllocate` parameter.

7.50 adds the `securityType` parameter.

7.60 adds the `drawerLossProtect` parameter.

7.75 adds the `dataAssurance` parameter.

Create RAID Volume (Free Extent Based Select)

This command creates a volume in the free space of a volume group.

Syntax

```
create volume volumeGroup="volumeGroupName"
userLabel="volumeName"
[freeCapacityArea=freeCapacityIndexNumber
capacity=volumeCapacity
owner=(a | b)
cacheReadPrefetch=(TRUE | FALSE)
segmentSize=segmentSizeValue
usageHint=(fileSystem | dataBase | multiMedia)]
[dssPreAllocate=(TRUE | FALSE)
securityType=(none | capable | enabled)
dataAssurance=(none | enabled)]
```

Parameters

Parameter	Description
volumeGroup	The alphanumeric identifier (including - and _) for a specific volume group in your storage array. Enclose the volume group name in double quotation marks (" ").
userLabel	The name that you want to give the new volume. Enclose the new volume name in double quotation marks (" ").
freeCapacityArea	The index number of the free space in an existing volume group that you want to use to create the new volume. Free capacity is defined as the free capacity between existing volumes in a volume group. For example, a volume group might have the following areas: volume 1, free capacity, volume 2, free capacity, volume 3, free capacity. To use the free capacity following volume 2, you would enter this index number: freeCapacityArea=2 Run the show volumeGroup command to determine whether the free capacity area exists.
capacity	The size of the volume that you are adding to the storage array. Size is defined in units of bytes, KB, MB, GB, or TB.
owner	The controller that owns the volume. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. If you do not specify an owner, the controller firmware determines the owner.
cacheReadPrefetch	The setting to turn on or turn off cache read prefetch. To turn on cache read prefetch, set this parameter to TRUE. To turn off cache read prefetch, set this parameter to FALSE.

Parameter	Description
<code>segmentSize</code>	The amount of data (in KB) that the controller writes on a single drive in a volume before writing data on the next drive. Valid values are 8, 16, 32, 64, 128, 256, or 512.
<code>usageHint</code>	The settings for both the <code>cacheReadPrefetch</code> parameter and the <code>segmentSize</code> parameter to be default values. The default values are based on the typical I/O usage pattern of the application that is using the volume. Valid values are <code>fileSystem</code> , <code>dataBase</code> , or <code>multiMedia</code> .
<code>dssPreAllocate</code>	The setting to make sure that reserve capacity is allocated for future segment size increases. The default value is <code>TRUE</code> .
<code>securityType</code>	The setting to specify the security level when creating the volume groups and all associated volumes. These settings are valid: <code>none</code> – The volume group and volumes are not secure. <code>capable</code> – The volume group and volumes are capable of having security set, but security has not been enabled. <code>enabled</code> – The volume group and volumes have security enabled.
<code>dataAssurance</code>	The setting to specify that a volume group, and the volumes within the volume group, has data assurance protection to make sure that the data maintains its integrity. When you use this parameter, only protected drives can be used for the volume group. These settings are valid: <ul style="list-style-type: none"> ■ <code>none</code> – The volume group does not have data assurance protection. ■ <code>enabled</code> – The volume group has data assurance protection. The volume group supports protected information and is formatted with protection information enabled.

Notes

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

The `owner` parameter defines which controller owns the volume. The preferred controller ownership of a volume is the controller that currently owns the volume group.

If you do not specify a capacity using the `capacity` parameter, all of the available capacity in the free capacity area of the volume group is used. If you do not specify capacity units, `bytes` is used as the default value.

Segment Size

The size of a segment determines how many data blocks that the controller writes on a single drive in a volume before writing data on the next drive. Each data block stores 512 bytes of data. A data block is the smallest unit of storage. The size of a segment determines how many data blocks that it contains. For example, an 8-KB segment holds 16 data blocks. A 64-KB segment holds 128 data blocks.

When you enter a value for the segment size, the value is checked against the supported values that are provided by the controller at run time. If the value that you entered is not valid, the controller returns a list of valid values. Using a single drive for a single request leaves other drives available to simultaneously service other requests.

If the volume is in an environment where a single user is transferring large units of data (such as multimedia), performance is maximized when a single data transfer request is serviced with a single data stripe. A data stripe is the segment size that is multiplied by the number of drives in the volume group that are used for data transfers. In this case, multiple drives are used for the same request, but each drive is accessed only once.

For optimal performance in a multiuser database or file system storage environment, set your segment size to minimize the number of drives that are required to satisfy a data transfer request.

Cache Read Prefetch

Cache read prefetch lets the controller copy additional data blocks into cache while the controller reads and copies data blocks that are requested by the host from the drives into cache. This action increases the chance that a future request for data can be fulfilled from cache. Cache read prefetch is important for multimedia applications that use sequential data transfers. The configuration settings for the storage array that you use determine the number of additional data blocks that the controller reads into cache. Valid values for the `cacheReadPrefetch` parameter are `TRUE` or `FALSE`. You do not need to enter a value for the `cacheReadPrefetch` parameter or the `segmentSize` parameter. If you do not enter a value, the controller firmware uses the `usageHint` parameter with `fileSystem` as the default value.

Entering a value for the `usageHint` parameter and a value for the `cacheReadPrefetch` parameter or a value for the `segmentSize` parameter does not cause an error. The value that you enter for the `cacheReadPrefetch` parameter or the `segmentSize` parameter takes priority over the value for the `usageHint` parameter.

Security Type

The `securityType` parameter is valid for drives that are capable of full disk encryption (FDE). With FDE, the controller firmware can create a key and activate the Drive Security feature. The Drive Security feature encrypts data as the data is written to the drive and decrypts the data as the data is read from the drive. Without the key created by the controller, the data written to the drive is inaccessible.

Before you can set the `securityType` parameter to `capable` or `enabled`, you must create a storage array security key. Use the `create storageArray securityKey` command to create a storage array security key. These commands are related to the security key:

- `create storageArray securityKey`
- `set storageArray securityKey`
- `import storageArray securityKey`
- `export storageArray securityKey`
- `start secureErase (drive | drives)`
- `enable volumeGroup [volumeGroupName] security`
- `create hostPort identifier`

Minimum Firmware Level

5.20

7.10 adds the `dssPreAllocate` parameter.

7.50 adds the `securityType` parameter.

7.75 adds the `dataAssurance` parameter.

Create RAID Volume (Manual Drive Select)

This command creates a new volume group and volume and lets you specify the drives for the volume.

NOTE You cannot use mixed drive types in the same volume group and volume. This command fails if you specify different types of drives for the RAID volume.

Syntax

```
create volume drives=(trayID1,drawerID1,slotID1 ...
trayIDn,drawerIDn,slotIDn)
volumeGroupUserLabel="volumeGroupName"
raidLevel=(0 | 1 | 3 | 5 | 6)
userLabel="volumeName"
[capacity=volumeCapacity
owner=(a | b)
cacheReadPrefetch=(TRUE | FALSE)
segmentSize=segmentSizeValue
usageHint=(fileSystem | dataBase | multiMedia)
trayLossProtect=(TRUE | FALSE)
drawerLossProtect=(TRUE | FALSE)
dssPreAllocate=(TRUE | FALSE)
securityType=(none | capable | enabled)
dataAssurance=(none | enabled)]
```

Parameters

Parameter	Description
<code>drives</code>	The drives that you want to assign to the volume that you want to create. For high-capacity drive trays, specify the tray ID value, the drawer ID value, and the slot ID value for each drive that you assign to the volume. For low-capacity drive trays, specify the tray ID value and the slot ID value for each drive that you assign to the volume. Tray ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the tray ID values, the drawer ID values, and the slot ID values in parentheses.
<code>volumeGroupUserLabel</code>	The alphanumeric identifier (including <code>-</code> and <code>_</code>) that you want to give the new volume group. Enclose the volume group identifier in double quotation marks (<code>" "</code>).
<code>raidLevel</code>	The RAID level of the volume group that contains the volume. Valid values are 0, 1, 3, 5, or 6.
<code>userLabel</code>	The name that you want to give the new volume. Enclose the new volume name in double quotation marks (<code>" "</code>).
<code>capacity</code>	The size of the volume that you are adding to the storage array. Size is defined in units of <code>bytes</code> , <code>KB</code> , <code>MB</code> , <code>GB</code> , or <code>TB</code> .
<code>owner</code>	The controller that owns the volume. Valid controller identifiers are <code>a</code> or <code>b</code> , where <code>a</code> is the controller in slot A, and <code>b</code> is the controller in slot B. If you do not specify an owner, the controller firmware determines the owner.
<code>cacheReadPrefetch</code>	The setting to turn on or turn off cache read prefetch. To turn off cache read prefetch, set this parameter to <code>FALSE</code> . To turn on cache read prefetch, set this parameter to <code>TRUE</code> .
<code>segmentSize</code>	The amount of data (in <code>KB</code>) that the controller writes on a single drive in a volume before writing data on the next drive. Valid values are 8, 16, 32, 64, 128, 256, or 512.
<code>usageHint</code>	The settings for both the <code>cacheReadPrefetch</code> parameter and the <code>segmentSize</code> parameter to be default values. The default values are based on the typical I/O usage pattern of the application that is using the volume. Valid values are <code>fileSystem</code> , <code>dataBase</code> , or <code>multiMedia</code> .
<code>trayLossProtect</code>	The setting to enforce tray loss protection when you create the repository. To enforce tray loss protection, set this parameter to <code>TRUE</code> . The default value is <code>FALSE</code> .
<code>drawerLossProtect</code>	The setting to enforce drawer loss protection when you create the mirrored repository volume. To enforce drawer loss protection, set this parameter to <code>TRUE</code> . The default value is <code>FALSE</code> .

Parameter	Description
dssPreAllocate	The setting to make sure that reserve capacity is allocated for future segment size increases. This default value is TRUE.
securityType	The setting to specify the security level when creating the volume groups and all associated volumes. These settings are valid: <ul style="list-style-type: none"> ■ none – The volume group and volumes are not secure. ■ capable – The volume group and volumes are capable of having security set, but security has not been enabled. ■ enabled – The volume group and volumes have security enabled.
dataAssurance	The setting to specify that a volume group, and the volumes within the volume group, has data assurance protection to make sure that the data maintains its integrity. When you use this parameter, only protected drives can be used for the volume group. These settings are valid: <ul style="list-style-type: none"> ■ none – The volume group does not have data assurance protection. ■ enabled – The volume group has data assurance protection. The volume group supports protected information and is formatted with protection information enabled.

Notes

The `drives` parameter supports both high-capacity drive trays and low-capacity drive trays. A high-capacity drive tray has drawers that hold the drives. The drawers slide out of the drive tray to provide access to the drives. A low-capacity drive tray does not have drawers. For a high-capacity drive tray, you must specify the identifier (ID) of the drive tray, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive tray, you need only specify the ID of the drive tray and the ID of the slot in which a drive resides. For a low-capacity drive tray, an alternative method for identifying a location for a drive is to specify the ID of the drive tray, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

If you set the `raidLevel` parameter to RAID 1, the controller firmware takes the list of drives and pairs them by using this algorithm:

Data drive	=	X
Parity drive	=	N/2 + X

In this algorithm X is 1 to N/2, and N is the number of drives in the list. For example, if you have six drives, the mirrored pairs are as follows:

Data	Parity
1	$N/2 + 1 = 4$
2	$N/2 + 2 = 5$
3	$N/2 + 3 = 6$

You can use any combination of alphanumeric characters, underscore (`_`), hyphen (`-`), and pound (`#`) for the names. Names can have a maximum of 30 characters.

The `owner` parameter defines which controller owns the volume. The preferred controller ownership of a volume is the controller that currently owns the volume group.

If you do not specify a capacity using the `capacity` parameter, all of the drive capacity that is available in the volume group is used. If you do not specify capacity units, `bytes` is used as the default value.

Segment Size

The size of a segment determines how many data blocks that the controller writes on a single drive in a volume before writing data on the next drive. Each data block stores 512 bytes of data. A data block is the smallest unit of storage. The size of a segment determines how many data blocks that it contains. For example, an 8-KB segment holds 16 data blocks. A 64-KB segment holds 128 data blocks.

When you enter a value for the segment size, the value is checked against the supported values that are provided by the controller at run time. If the value that you entered is not valid, the controller returns a list of valid values. Using a single drive for a single request leaves other drives available to simultaneously service other requests.

If the volume is in an environment where a single user is transferring large units of data (such as multimedia), performance is maximized when a single data transfer request is serviced with a single data stripe. A data stripe is the segment size that is multiplied by the number of drives in the volume group that are used for data transfers. In this case, multiple drives are used for the same request, but each drive is accessed only once.

For optimal performance in a multiuser database or file system storage environment, set your segment size to minimize the number of drives that are required to satisfy a data transfer request.

Cache Read Prefetch

Cache read prefetch lets the controller copy additional data blocks into cache while the controller reads and copies data blocks that are requested by the host from the drive into cache. This action increases the chance that a future request for data can be fulfilled from cache. Cache read prefetch is important for multimedia applications

that use sequential data transfers. The configuration settings for the storage array that you use determine the number of additional data blocks that the controller reads into cache. Valid values for the `cacheReadPrefetch` parameter are `TRUE` or `FALSE`.

You do not need to enter a value for the `cacheReadPrefetch` parameter or the `segmentSize` parameter. If you do not enter a value, the controller firmware uses the `usageHint` parameter with `fileSystem` as the default value. Entering a value for the `usageHint` parameter and a value for the `cacheReadPrefetch` parameter or a value for the `segmentSize` parameter does not cause an error. The value that you enter for the `cacheReadPrefetch` parameter or the `segmentSize` parameter takes priority over the value for the `usageHint` parameter.

Tray Loss Protection and Drawer Loss Protection

For tray loss protection to work, each drive in a volume group must be on a separate tray. If you set the `trayLossProtect` parameter to `TRUE` and have selected more than one drive from any one tray, the storage array returns an error. If you set the `trayLossProtect` parameter to `FALSE`, the storage array performs operations, but the volume group that you create might not have tray loss protection.

Tray loss protection is not valid when you create volumes on existing volume groups.

The `drawerLossProtect` parameter defines if data on a volume is accessible if a drawer fails. When you assign the drives, if you set the `drawerLossProtect` parameter to `TRUE` and select more than one drive from any one drawer, the storage array returns an error. If you set the `drawerLossProtect` parameter to `FALSE`, the storage array performs operations, but the volume group that you create might not have drawer loss protection.

You must set the `trayLossProtect` parameter and the `drawerLossProtect` parameter to the same value. Both of the parameters must be either `TRUE` or `FALSE`. If the `trayLossProtect` parameter and the `drawerLossProtect` parameter are set to different values, the storage array returns an error.

Security Type

The `securityType` parameter is valid for drives that are capable of full disk encryption (FDE). With FDE, the controller firmware can create a key and activate the Drive Security feature. The Drive Security feature encrypts data as the data is written to the drive and decrypts the data as the data is read from the drive. Without the key created by the controller, the data written to the drive is inaccessible.

Before you can set the `securityType` parameter to `capable` or `enabled`, you must create a storage array security key. Use the `create storageArray securityKey` command to create a storage array security key. These commands are related to the security key:

- `create storageArray securityKey`
- `enable volumeGroup [volumeGroupName] security`
- `export storageArray securityKey`

- `import storageArray securityKey`
- `set storageArray securityKey`
- `start secureErase (drive | drives)`

Minimum Firmware Level

5.20

7.10 adds RAID Level 6 capability and the `dssPreAllocate` parameter.

7.60 adds the `drawerID` user input and the `drawerLossProtect` parameter.

7.75 adds the `dataAssurance` parameter.

Create Read-Only Snapshot Volume

This command creates a read-only snapshot volume for the snapshot images of a base volume. To change a read-only snapshot volume to a read/write volume, use the `set snapVolume convertToReadWrite` command.

NOTE You cannot use this command for a snapshot image that is used in online volume copy.

Syntax

```
create snapVolume userLabel="snapVolumeName"
snapImageID="snapCGID: imageID"
readOnly
```

Parameters

Parameter	Description
<code>userLabel</code>	The name that you want to give to a snapshot volume. Enclose the snapshot volume name in double quotation marks (" ").

Parameter	Description
snapImageID	<p>The alphanumeric identifier of a snapshot image that you want to add to the new read-only snapshot volume. The identifier of a snapshot image is comprised of two parts:</p> <ul style="list-style-type: none"> ■ The name of the snapshot group ■ An identifier for the snapshot image in the snapshot group <p>The identifier for the snapshot image can be one of these:</p> <ul style="list-style-type: none"> ■ An integer value that is the sequence number of the snapshot in the snapshot group. ■ NEWEST - Use this option when you want to show the latest snapshot image created in the snapshot group. ■ OLDEST - Use this option when you want to show the earliest snapshot image created in the snapshot group. <p>Enclose the snapshot image name in double quotation marks (" ").</p>
readOnly	<p>This parameter sets the snapshot volume to read-only. This parameter is actually a Boolean; however, in the context of this command, the Boolean value is always TRUE.</p> <hr/> <p>NOTE With the <code>readOnly</code> parameter, snapshot volumes are not created.</p> <hr/>

Notes

You can use any combination of alphanumeric characters, underscore (`_`), hyphen (`-`), and pound (`#`) for the names. Names can have a maximum of 30 characters.

The identifier of a snapshot image has two parts separated by a colon (`:`):

- The name of the snapshot group
- The identifier of the snapshot image

For example, if you want to create a read-only volume named `engData1` using the most recent snapshot image in a snapshot group that has the name `snapGroup1`, you would use this command:

```
create snapVolume userLabel="engData1"
snapImageID="snapGroup1:newest" readOnly;
```

Minimum Firmware Level

7.83

Create Snapshot (Legacy) Volume

This command creates a snapshot (legacy) volume of a base volume. You can also use this command to create a new repository volume group if one does not already exist, or if you would prefer a different repository volume group. This command defines three ways to create a snapshot (legacy) volume:

- In a new repository volume group created from user-defined drives
- In a new repository volume group created from a user-defined number of drives
- In an existing repository volume group

If you choose to define a number of drives, the controller firmware chooses which drives to use for the snapshot (legacy) volume.

Syntax (User-Defined Drives)

```
create snapshotVolume baseVolume="baseVolumeName"
(repositoryRAIDLevel=(1 | 3 | 5 | 6)
repositoryDrives=(trayID1,drawerID1,slotID1 ...
trayIDn,drawerIDn,slotIDn)
[repositoryVolumeGroupUserLabel="repositoryVolumeGroupName"
trayLossProtect=(TRUE | FALSE)
drawerLossProtect=(TRUE | FALSE)
freeCapacityArea=freeCapacityIndexNumber
userLabel="snapshotVolumeName"
warningThresholdPercent=percentValue
repositoryPercentOfBase=percentValue
repositoryUserLabel="repositoryName"
repositoryFullPolicy=(failBaseWrites | failSnapshot)
enableSchedule=(TRUE | FALSE)
schedule (immediate | snapshotSchedule)]
```

Syntax (User-Defined Number of Drives)

```
create snapshotVolume baseVolume="baseVolumeName"
repositoryRAIDLevel=(1 | 3 | 5 | 6)
repositoryDriveCount=numberOfDrives
[repositoryVolumeGroupUserLabel="repositoryVolumeGroupName"
driveMediaType=(HDD | SSD | unknown | allMedia)]
driveType=(fibre | SATA | SAS)
trayLossProtect=(TRUE | FALSE)
drawerLossProtect=(TRUE | FALSE)
userLabel="snapshotVolumeName"
warningThresholdPercent=percentValue
repositoryPercentOfBase=percentValue
repositoryUserLabel="repositoryName"
repositoryFullPolicy=(failBaseWrites | failSnapshot)
enableSchedule=(TRUE | FALSE)
schedule (immediate | snapshotSchedule)]
```

Syntax (Existing Repository Volume Group)

```
create snapshotVolume baseVolume="baseVolumeName"
[repositoryVolumeGroup="repositoryVolumeGroupName"
repositoryUserLabel="repositoryName"
freeCapacityArea=freeCapacityIndexNumber
userLabel="snapshotVolumeName"
warningThresholdPercent=percentValue
repositoryPercentOfBase=percentValue
repositoryFullPolicy=(failBaseWrites | failSnapshot)
enableSchedule=(TRUE | FALSE)
schedule (immediate | snapshotSchedule)]
```

Parameters

Parameter	Description
<code>baseVolume</code>	The name of the base volume from which you want to take a snapshot (legacy). Enclose the base volume name in double quotation marks (" ").
<code>repositoryRAIDLevel</code>	Use this parameter when you create a new volume group. The RAID level for the snapshot (legacy) repository volume group. Valid values are 1, 3, 5, or 6.
<code>repositoryDrives</code>	Use this parameter when you create a new volume group. The drives that you want to assign to the snapshot (legacy) repository volume group. For high-capacity drive trays, specify the tray ID value, the drawer ID value, and the slot ID value for each drive that you assign to the snapshot (legacy) repository volume. For low-capacity drive trays, specify the tray ID value and the slot ID value for each drive that you assign to the snapshot (legacy) repository volume. Tray ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the tray ID values, the drawer ID values, and the slot ID values in parentheses.
<code>repositoryDriveCount</code>	Use this parameter when you create a new volume group. The number of unassigned drives that you want to use for the snapshot (legacy) repository volume group.

Parameter	Description
repositoryVolumeGroupUserLabel	<p>Use this parameter when you create a new volume group.</p> <p>The name of a new volume group to be used for the repository volume. Enclose the repository volume group name in double quotation marks (" ").</p>
repositoryVolumeGroup	<p>The name of an existing volume group where you want to place the repository volume. Use this parameter if you do not want to put the repository volume in the same volume group as the base volume. The default is to use the same volume group for both the base volume and the repository volume. Enclose the name of the repository volume group in double quotation marks (" ").</p>
userLabel	<p>The name that you want to give to the snapshots (legacy) volume. If you do not want to provide a name, the CLI creates a name using the base volume user label that you provide.</p>
trayLossProtect	<p>The setting to enforce tray loss protection when you create the snapshot (legacy) repository volume. To enforce tray loss protection, set this parameter to TRUE. The default value is FALSE.</p>
drawerLossProtect	<p>The setting to enforce drawer loss protection when you create the mirror repository volume. To enforce drawer loss protection, set this parameter to TRUE. The default value is FALSE.</p>

Parameter	Description
driveMediaType	<p>The type of drive medium that you want to use for the mirror repository volume. Valid drive media are these:</p> <ul style="list-style-type: none"> ■ HDD – Use this option when you have hard drives in the drive tray. ■ SSD – Use this option when you have solid state drives in the drive tray. ■ unknown – Use if you are not sure what types of drive media are in the drive tray. ■ allMedia – Use this option when you want to use all types of drive media that are in the drive tray. <p>Use this parameter when you use the repositoryDriveCount parameter.</p> <p>You must use this parameter when you have more than one type of drive media in your storage array.</p>
driveType	<p>The type of drive that you want to use in the volume. You cannot mix drive types.</p> <p>You must use this parameter when you have more than one type of drive in your storage array.</p> <p>Valid drive types are:</p> <ul style="list-style-type: none"> ■ fibre ■ SATA ■ SAS <p>If you do not specify a drive type, the command defaults to fibre.</p> <p>Use this parameter when you use the repositoryDriveCount parameter.</p>

Parameter	Description
freeCapacityArea	<p>The index number of the free space in an existing volume group that you want to use to create the snapshot (legacy) repository volume. Free capacity is defined as the free capacity between existing volumes in a volume group. For example, a volume group might have these areas: volume 1, free capacity, volume 2, free capacity, volume 3, free capacity. To use the free capacity following volume 2, you would specify:</p> <pre>freeCapacityArea=2</pre> <p>Run the <code>show volumeGroup</code> command to determine if a free capacity area exists.</p>
warningThresholdPercent	<p>The percentage of repository capacity at which you receive a warning that the snapshot (legacy) repository volume is nearing full. Use integer values. For example, a value of 70 means 70 percent. The default value is 50.</p>
repositoryPercentOfBase	<p>The size of the snapshot (legacy) repository volume as a percentage of the base volume. Use integer values. For example, a value of 40 means 40 percent. The default value is 20.</p>
repositoryUserLabel	<p>The name that you want to give to the snapshot (legacy) repository volume. Enclose the snapshot (legacy) repository volume name in double quotation marks (" ").</p>
repositoryFullPolicy	<p>How you want snapshot (legacy) processing to continue if the snapshot (legacy) repository volume is full. You can choose to fail writes to the base volume (<code>failBaseWrites</code>) or fail the snapshot (legacy) volume (<code>failSnapshot</code>). The default value is <code>failSnapshot</code>.</p>
enableSchedule	<p>Use this parameter to turn on or to turn off the ability to schedule a snapshot (legacy) operation. To turn on snapshot (legacy) scheduling, set this parameter to <code>TRUE</code>. To turn off snapshot (legacy) scheduling, set this parameter to <code>FALSE</code>.</p>

Parameter	Description
schedule	<p>Use this parameter to schedule a snapshot operation.</p> <p>You can use one of these options for setting a schedule for a snapshot operation:</p> <ul style="list-style-type: none"> ■ immediate ■ startDate ■ scheduleDay ■ startTime ■ scheduleInterval ■ endDate ■ timesPerDay ■ timeZone ■ scheduleDate ■ month <p>See the "Notes" section for information explaining how to use these options.</p>

Notes

The volume that you are taking a snapshot (legacy) of must be a standard volume in the storage array. The maximum number of snapshot (legacy) volumes that you can create is one-half of the total number of volumes that are supported by a controller.

You can use any combination of alphanumeric characters, underscore (`_`), hyphen (`-`), and pound (`#`) for the names. Names can have a maximum of 30 characters.

One technique for naming the snapshot (legacy) volume and the snapshot (legacy) repository volume is to add a hyphenated suffix to the original base volume name. The suffix distinguishes between the snapshot (legacy) volume and the snapshot (legacy) repository volume. For example, if you have a base volume with a name of Engineering Data, the snapshot (legacy) volume can have a name of Engineering Data-S1, and the snapshot (legacy) repository volume can have a name of EngineeringData-R1.

If you do not choose a name for either the snapshot (legacy) volume or the snapshot (legacy) repository volume, the storage management software creates a default name by using the base volume name. An example of the snapshot (legacy) volume name that the controllers might create is, if the base volume name is aaa and does not have a snapshot (legacy) volume, the default snapshot (legacy) volume name is aaa-1. If the base volume already has $n-1$ number of snapshot (legacy) volumes, the default name

is `aaa-n`. An example of the snapshot (legacy) repository volume name that the controller might create is, if the base volume name is `aaa` and does not have a snapshot (legacy) repository volume, the default snapshot (legacy) repository volume name is `aaa-R1`. If the base volume already has $n-1$ number of snapshot (legacy) repository volumes, the default name is `aaa-Rn`.

If you do not specify the unconfigured space or free space, the snapshot (legacy) repository volume is placed in the same volume group as the base volume. If the volume group where the base volume resides does not have enough space, this command fails.

The `repositoryDrives` parameter supports both high-capacity drive trays and low-capacity drive trays. A high-capacity drive tray has drawers that hold the drives. The drawers slide out of the drive tray to provide access to the drives. A low-capacity drive tray does not have drawers. For a high-capacity drive tray, you must specify the identifier (ID) of the drive tray, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive tray, you need only specify the ID of the drive tray and the ID of the slot in which a drive resides. For a low-capacity drive tray, an alternative method for identifying a location for a drive is to specify the ID of the drive tray, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

Tray Loss Protection and Drawer Loss Protection

When you assign the drives, if you set the `trayLossProtect` parameter to `TRUE` and have selected more than one drive from any one tray, the storage array returns an error. If you set the `trayLossProtect` parameter to `FALSE`, the storage array performs operations, but the volume group that you create might not have tray loss protection.

When the controller firmware assigns the drives, if you set the `trayLossProtect` parameter to `TRUE`, the storage array returns an error if the controller firmware cannot provide drives that result in the new volume group having tray loss protection. If you set the `trayLossProtect` parameter to `FALSE`, the storage array performs the operation even if it means the volume group might not have tray loss protection.

The `drawerLossProtect` parameter defines if data on a volume is accessible if a drawer fails. When you assign the drives, if you set the `drawerLossProtect` parameter to `TRUE` and select more than one drive from any one drawer, the storage array returns an error. If you set the `drawerLossProtect` parameter to `FALSE`, the storage array performs operations, but the volume group that you create might not have drawer loss protection.

If you have a storage configuration that includes a drive tray that has drawers to hold the drives, follow these guidelines when configuring tray loss protection:

- If you set `trayLossProtect` to `TRUE`, then you must set `drawerLossProtect` to `TRUE`.
- If you set `trayLossProtect` to `FALSE`, then you can set `drawerLossProtect` to either `TRUE` or `FALSE`.

If you set `trayLossProtect` to `TRUE` and `drawerLossProtect` to `FALSE`, the storage array returns an error.

Scheduling Snapshots (Legacy)

The `enableSchedule` parameter and the `schedule` parameter provide a way for you to schedule automatic snapshots (legacy). Using these parameters, you can schedule snapshots (legacy) daily, weekly, or monthly (by day or by date). The `enableSchedule` parameter turns on or turns off the ability to schedule snapshots (legacy). When you enable scheduling, you use the `schedule` parameter to define when you want the snapshots (legacy) to occur.

This list explains how to use the options for the `schedule` parameter:

- `immediate` – As soon as you enter the command, a snapshot (legacy) volume is created and a copy-on-write operation begins.
- `startDate` – A specific date on which you want to create a snapshot (legacy) volume and perform a copy-on-write operation. The format for entering the date is `MM:DD:YY`. If you do not provide a start date, the current date is used. An example of this option is `startDate=06:27:11`.
- `scheduleDay` – A day of the week on which you want to create a snapshot (legacy) volume and perform a copy-on-write operation. The values that you can enter are: `monday`, `tuesday`, `wednesday`, `thursday`, `friday`, `saturday`, `sunday`, and `all`. An example of this option is `scheduleDay=wednesday`.
- `startTime` – The time of a day that you want to create a snapshot (legacy) volume and start performing a copy-on-write operation. The format for entering the time is `HH:MM`, where `HH` is the hour and `MM` is the minute past the hour. Use a 24-hour clock. For example, 2:00 in the afternoon is 14:00. An example of this option is `startTime=14:27`.
- `scheduleInterval` – An amount of time, in minutes, that you want to have as a minimum between copy-on-write operation. You can create a schedule in which you have overlapping copy-on-write operations because of the duration of a copy operation. You can make sure that you have time between copy-on-write operations by using this option. The maximum value for the `scheduleInterval` option is 1440 minutes. An example of this option is `scheduleInterval=180`.
- `endDate` – A specific date on which you want to stop creating a snapshot (legacy) volume and end the copy-on-write operations. The format for entering the date is `MM:DD:YY`. An example of this option is `endDate=11:26:11`.
- `timesPerDay` – The number of times that you want the schedule to run in a day. An example of this option is `timesPerDay=4`.

- `timeZone` – Use this parameter to define the time zone in which the storage array is operating. You can define the time zone in one of two ways:
 - `GMT±HH:MM` – The time zone offset from GMT. Enter the offset in hours and minutes. For example `GMT-06:00` is the central time zone in the United States.
 - `Text string` – Standard time zone text strings. For example: `"USA/Chicago"` or `"Germany/Berlin"`. Time zone text strings are case sensitive. If you enter an incorrect text string, GMT time is used. Enclose the text string in double quotation marks.
- `scheduleDate` – A day of the month on which you want to create a snapshot volume and perform a copy-on-write operation. The values for the days are numerical and in the range of 1–31. Enclose the value for the day in double quotation marks inside parentheses. If you want to create a snapshot volume on a specific date of a specific month use the `scheduleDate` option with the `month` option. If you have set up a weekly schedule, you cannot use the `scheduleDate` option. An example of the `scheduleDate` option is `scheduleDate=("15")`.
- `month` – A specific month on which you want to create a snapshot volume and perform a copy-on-write operation. The values for the months are: `jan`, `feb`, `mar`, `apr`, `may`, `jun`, `jul`, `aug`, `sep`, `oct`, `nov`, and `dec`. Enclose the value in parentheses. You can enter more than one month by enclosing the months in a single set of parentheses and separating each month with a white space. If you want to create a snapshot volume on a specific date of a specific month use the `scheduleDate` option with the `month` option. These are examples of the `month` option:
 - `month=(mar)`
 - `month=(mar apr may)`

The code string for defining a schedule is similar to these examples:

```
enableSchedule=true schedule startTime=14:27
enableSchedule=true schedule scheduleInterval=180
enableSchedule=true schedule timeZone=GMT-06:00
enableSchedule=true schedule timeZone="USA/Chicago"
enableSchedule=true schedule month=(mar) scheduleDate=("15")
```

If you also use the `scheduleInterval` option, the firmware will choose between the `timesPerDay` option and the `scheduleInterval` option by selecting the lowest value of the two options. The firmware calculates an integer value for the `scheduleInterval` option by dividing 1440 by the `scheduleInterval` option value that you set. For example, $1440/180 = 8$. The firmware then compares the `timesPerDay` integer value with the calculated `scheduleInterval` integer value and uses the smaller value.

To remove a schedule, use the `delete volume` command with the `schedule` parameter. The `delete volume` command with the `schedule` parameter deletes only the schedule, not the snapshot (legacy) volume.

Minimum Firmware Level

5.00

7.10 adds RAID 6 Level capability.

7.60 adds the `drawerID` user input, the `driveMediaType` parameter, and the `drawerLossProtect` parameter.

7.77 adds scheduling.

7.83 removes the `noEndDate` option. This option is not supported.

7.86 adds the `scheduleDate` option and the `month` option.

Create Snapshot Group

This command creates a new snapshot group and the associated repository volume. A snapshot group contains a sequence of snapshot images of an associated base volume. A snapshot group has a repository volume that is used to save data for all of the snapshot images that are part of the snapshot group.

Syntax

```
create snapGroup userLabel="snapGroupName"
sourceVolume="volumeName"
(repositoryVolume="repos_xxxx" |
repositoryVolume=(volumeGroupName [capacity=capacityValue] |
repositoryVolume=(diskPoolName [capacity=capacityValue]))
[repositoryFullPolicy=(failBaseWrites | purgeSnapImages) |
rollbackPriority=(highest | high | medium | low | lowest) |
repositoryFullLimit=percentValue |
autoDeleteLimit=numberOfSnapImages |
enableSchedule=(TRUE | FALSE)
schedule (immediate | snapshotSchedule)]
```

Parameters

Parameter	Description
<code>userLabel</code>	The name that you want to give the new snapshot group. Enclose the snapshot group identifier in double quotation marks (" ").
<code>sourceVolume</code>	The name of the volume that you want to use as the source for your snapshot images. Enclose the source volume name in double quotation marks (" ").

Parameter	Description
repositoryVolume	<p>The name of the repository volume that will contain the changed data of the snapshot group.</p> <p>You have two options for defining the name of a repository volume:</p> <ul style="list-style-type: none"> ■ Use an existing repository volume: name ■ Create a new repository volume when you run this command <p>The name of an existing repository volume is comprised of two parts:</p> <ul style="list-style-type: none"> ■ The term <i>repos</i> ■ A four digit numerical identifier that the storage management software assigns to the repository volume name <p>Enclose the name of the existing repository volume in double quotation marks (" ").</p> <p>If you want to create a new repository volume when you run this command you must enter the name of either a volume group or a disk pool in which you want the repository volume. Optionally, you also can define the capacity of the repository volume. If you want to define the capacity you can use these values:</p> <ul style="list-style-type: none"> ■ An integer value that represents a percentage of the base volume capacity ■ A decimal fraction value that represents a percentage of the base volume capacity ■ A specific size for the repository volume. Size is defined in units of <code>bytes</code>, KB, MB, GB, or TB. <p>If you do not use the capacity option, the storage management software sets the capacity to 20 percent of the base volume capacity.</p> <p>When you run this command the storage management software creates the repository volume for the snapshot volume.</p>
repositoryFullPolicy	<p>Defines how snapshot image processing continues if the snapshot group repository volume is full. You can choose to fail I/O writes to the base volume (<code>failBaseWrites</code>) or delete (purge) the snapshot images (<code>purgeSnapImages</code>) in the repository volume. The <code>purgeSnapImages</code> option deletes the oldest snapshot images to free up space. The default action is <code>purgeSnapImages</code>.</p>

Parameter	Description
rollBackPriority	Determines whether system resources should be allocated to the rollback operation at the expense of system performance. A value of <code>high</code> indicates that the rollback operation is prioritized over all other host I/O. A value of <code>low</code> indicates that the rollback operation should be performed with minimal impact to host I/O. The default value is <code>medium</code> .
repositoryFullLimit	The percentage of repository capacity at which you receive a warning that the snapshot group repository volume is nearing full. Use integer values. For example, a value of 70 means 70 percent. The default value is 75.
autoDeleteLimit	Each snapshot group can be configured to perform automatic deletion of its snapshot images to keep the total number of snapshot images in the snapshot group at or below a designated level. When this option is enabled, then any time a new snapshot image is created in the snapshot group, the system automatically deletes the oldest snapshot image in the group to comply with the limit value. This action frees repository capacity so it can be used to satisfy ongoing copy-on-write requirements for the remaining snapshot images.
enableSchedule	Use this parameter to turn on or to turn off the ability to schedule a snapshot operation. To turn on snapshot scheduling, set this parameter to <code>TRUE</code> . To turn off snapshot scheduling, set this parameter to <code>FALSE</code> .

Parameter	Description
schedule	<p>Use this parameter to schedule a snapshot operation.</p> <p>You can use one of these options for setting a schedule for a snapshot operation:</p> <ul style="list-style-type: none"> ■ immediate ■ startDate ■ scheduleDay ■ startTime ■ scheduleInterval ■ endDate ■ timesPerDay ■ timeZone ■ scheduleDate ■ month <p>See the "Notes" section for information explaining how to use these options.</p>

Notes

Each snapshot group name must be unique. You can use any combination of alphanumeric characters, underscore (_), hyphen(-), and pound (#) for the user label. User labels can have a maximum of 30 characters.

To create a snapshot group, you must have an associated repository volume in which you store the snapshot images. You can either use an existing repository volume or create a new repository volume. You can create the repository volume when you create the snapshot group. A snapshot group repository volume is an expandable volume that is structured as a concatenated collection of up to 16 standard volume entities. Initially, an expandable repository volume has only a single element. The capacity of the expandable repository volume is exactly that of the single element. You can increase the capacity of an expandable repository volume by attaching additional standard volumes to it. The composite expandable repository volume capacity then becomes the sum of the capacities of all of the concatenated standard volumes.

A snapshot group has a strict ordering of snapshot images based on the time that each snapshot image is created. A snapshot image that is created after another snapshot image is a *successor* relative to that other snapshot image. A snapshot image that is created before another snapshot image is a *predecessor* relative to that other one.

A snapshot group repository volume must satisfy a minimum capacity requirement that is the sum of the following:

- 32 MB to support fixed overhead for the snapshot group and for copy-on-write processing.
- Capacity for rollback processing, which is 1/5000th of the capacity of the base volume.

The minimum capacity is enforcement by the controller firmware and the storage management software.

When you first create a snapshot group, it does not contain any snapshot images. When you create snapshot images, you add the snapshot images to a snapshot group. Use the `create snapImage` command to create snapshot images and add the snapshot images to a snapshot group.

A snapshot group can have one of these states:

- **Optimal** – The snapshot group is operating normally.
- **Full** – The snapshot group repository is full. Additional copy-on-write operations can not be performed. This state is possible only for snapshot groups that have the Repository Full policy set to Fail Base Writes. Any snapshot group in a Full state causes a Needs-Attention condition to be posted for the storage array.
- **Over Threshold** – The snapshot group repository volume usage is at or beyond its alert threshold. Any snapshot group in this state causes a Needs-Attention condition to be posted for the storage array.
- **Failed** – The snapshot group has encountered a problem that has made all snapshot images in the snapshot group unusable. For example, certain types of repository volume failures can cause a Failed state. To recover from a Failed state use the `revive snapGroup` command.

Automatic Snapshot Image Deletion

You can configure each snapshot group to automatically delete the snapshot images by using the `autoDeleteLimit` parameter. Automatically deleting the snapshot images enables you to avoid having to routinely, manually delete the images that you do not want and that might prevent the creation of future snapshot images because the repository volume is full. When you use the `autoDeleteLimit` parameter it causes the storage management software to automatically delete snapshot images, starting with the oldest. The storage management software deletes snapshot images until it reaches a number of snapshot images that is equal to the number that you enter with `autoDeleteLimit` parameter. When new snapshot images are added to the repository volume, the storage management software deletes the oldest snapshot images until the `autoDeleteLimit` parameter number is reached.

Scheduling Snapshots

The `enableSchedule` parameter and the `schedule` parameter provide a way for you to schedule creating snapshot images for a snapshot group. Using these parameters, you can schedule snapshots daily, weekly, or monthly (by day or by date).

The `enableSchedule` parameter turns on or turns off the ability to schedule snapshots. When you enable scheduling, you use the `schedule` parameter to define when you want the snapshots to occur.

This list explains how to use the options for the `schedule` parameter:

- `immediate` – As soon as you enter the command, a snapshot image is created, and a copy-on-write operation begins.
- `startDate` – A specific date on which you want to create a snapshot image and perform a copy-on-write operation. The format for entering the date is `MM:DD:YY`. If you do not provide a start date, the current date is used. An example of this option is `startDate=06:27:11`.
- `scheduleDay` – A day of the week on which you want to create a snapshot image and perform a copy-on-write operation. You can enter these values: `monday`, `tuesday`, `wednesday`, `thursday`, `friday`, `saturday`, `sunday`, and `all`. An example of this option is `scheduleDay=wednesday`.
- `startTime` – The time of a day that you want to create a snapshot image and start performing a copy-on-write operation. The format for entering the time is `HH:MM`, where `HH` is the hour and `MM` is the minute past the hour. Use a 24-hour clock. For example, 2:00 in the afternoon is 14:00. An example of this option is `startTime=14:27`.
- `scheduleInterval` – An amount of time, in minutes, that you want to have as a minimum between copy-on-write operations. You can possibly create a schedule in which you have overlapping copy-on-write operations because of the duration of a copy operation. You can make sure that you have time between copy-on-write operations by using this option. The maximum value for the `scheduleInterval` option is 1440 minutes. An example of this option is `scheduleInterval=180`.
- `endDate` – A specific date on which you want to stop creating a snapshot image and end the copy-on-write operation. The format for entering the date is `MM:DD:YY`. An example of this option is `endDate=11:26:11`.
- `noEndDate` – Use this option if you do not want your scheduled copy-on-write operation to end. If you later decide to end the copy-on-write operations you must re-enter the `set snapGroup` command and specify an end date.
- `timesPerDay` – The number of times that you want the schedule to run in a day. An example of this option is `timesPerDay=4`.
- `timeZone` – Use this parameter to define the time zone in which the storage array is operating. You can define the time zone in one of two ways:
 - `GMT±HH:MM` – The time zone offset from GMT. Enter the offset in hours and minutes. For example `GMT-06:00` is the central time zone in the United States.
 - `Text string` – Standard time zone text strings. For example: `"USA/Chicago"` or `"Germany/Berlin"`. Time zone text strings are case sensitive. If you enter an incorrect text string, GMT time is used. Enclose the text string in double quotation marks.

- `scheduleDate` – A day of the month on which you want to create a snapshot volume and perform a copy-on-write operation. The values for the days are numerical and in the range of 1–31. Enclose the value for the day in double quotation marks inside parentheses. If you want to create a snapshot volume on a specific date of a specific month use the `scheduleDate` option with the `month` option. If you have set up a weekly schedule, you cannot use the `scheduleDate` option. An example of the `scheduleDate` option is `scheduleDate=("15")`.
- `month` – A specific month on which you want to create a snapshot volume and perform a copy-on-write operation. The values for the months are: `jan`, `feb`, `mar`, `apr`, `may`, `jun`, `jul`, `aug`, `sep`, `oct`, `nov`, and `dec`. Enclose the value in parentheses. You can enter more than one month by enclosing the months in a single set of parentheses and separating each month with a white space. If you want to create a snapshot volume on a specific date of a specific month use the `scheduleDate` option with the `month` option. These are examples of the `month` option:
 - `month=(mar)`
 - `month=(mar apr may)`

The code string for defining a schedule is similar to these examples:

```
enableSchedule=true schedule startTime=14:27
enableSchedule=true schedule scheduleInterval=180
enableSchedule=true schedule timeZone=GMT-06:00
enableSchedule=true schedule timeZone="USA/Chicago"
enableSchedule=true schedule month=(mar) scheduleDate=( "15" )
```

If you also use the `scheduleInterval` option, the firmware chooses between the `timesPerDay` option and the `scheduleInterval` option by selecting the lowest value of the two options. The firmware calculates an integer value for the `scheduleInterval` option by dividing 1440 by the `scheduleInterval` option value that you set. For example, $1440/180 = 8$. The firmware then compares the `timesPerDay` integer value with the calculated `scheduleInterval` integer value and uses the smaller value.

To remove a schedule, use the `delete volume` command with the `schedule` parameter. The `delete volume` command with the `schedule` parameter deletes only the schedule, not the snapshot volume.

Minimum Firmware Level

7.83

7.86 adds the `scheduleDate` option and the `month` option.

Create Snapshot Image

This command creates a new snapshot image in one or more existing snapshot groups. Before you can create a snapshot image, you must first have at least one snapshot group into which you can place the snapshot image. To create a snapshot group use the `create snapGroup` command.

Syntax

```
create snapImage (snapGroup="snapGroupName" |
snapGroups=("snapGroupName1" . . . "snapGroupName"))
```

Parameters

Parameter	Description
snapGroup or snapGroups	The name of the snapshot group that will contain the snapshot image. Enclose the name of the snapshot group in double quotation marks (" "). If you enter more than one snapshot group name, enclose each of the snapshot group names in double quotation marks (" "). Enclose all of the snapshot group names in parentheses.

Notes

You can use any combination of alphanumeric characters, underscore (_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

A snapshot image is a logical point-in-time image of the contents of an associated *base* volume. The snapshot image is created instantly and records the state of the base volume at that moment. Every snapshot image is created in the context of exactly one *snapshot group*. A snapshot group is a sequence of snapshot images of the associated base volume. A snapshot group has one *repository volume* used to save all of the data from the snapshot images. The snapshot images in a snapshot group have a specific order. The specific order of the snapshot images enables you to manage the snapshot images, such as restoring a specific snapshot image to the base volume or deleting the snapshot images that you no longer need.

The result of creating a snapshot image of a consistency group is a snapshot image of every member volume of the consistency group.

Minimum Firmware Level

7.83

Create Snapshot Volume

This command creates a snapshot volume with read-write capabilities for snapshot images of a base volume. You can map the snapshot volume to a host and all of the host writes reside in the repository volume associated with the snapshot volume. You can assign the new snapshot volume to an existing repository volume, or you can create a new repository volume in a volume group or disk pool.

NOTE You cannot use this command for a snapshot image that is used in online volume copy.

Syntax

```
create snapVolume userLabel="snapVolumeName"  
snapImageID="snapCGID:imageID"  
[(repositoryVolume="repos_xxxx" |  
repositoryVolume=(volumeGroupName [capacity=capacityValue])  
repositoryVolume=(diskPoolName [capacity=capacityValue]))  
repositoryFullLimit=percentValue]
```

Parameters

Parameter	Description
userLabel	The name that you want to give to a snapshot volume. Enclose the snapshot volume name in double quotation marks (" ").
snapImageID	<p>The alphanumeric identifier of a snapshot image that you want to add to the new snapshot volume. The identifier of a snapshot image is comprised of two parts:</p> <ul style="list-style-type: none">■ The name of the snapshot group■ An identifier for the snapshot image in the snapshot group <p>The identifier for the snapshot image can be one of these:</p> <ul style="list-style-type: none">■ An integer value that is the sequence number of the snapshot in the snapshot group.■ NEWEST - Use this option when you want to show the latest snapshot image created in the snapshot group.■ OLDEST - Use this option when you want to show the earliest snapshot image created in the snapshot group. <p>Enclose the snapshot image name in double quotation marks (" ").</p>

Parameter	Description
repositoryVolume	<p>The name of the repository volume that holds changed data from the snapshot image.</p> <p>You have two options for defining the name of a repository volume:</p> <ul style="list-style-type: none"> ■ Use an existing repository volume name ■ Create a new repository volume when you run this command <p>The name of an existing repository volume is comprised of two parts:</p> <ul style="list-style-type: none"> ■ The term <i>repos</i> ■ A four digit numerical identifier that the storage management software assigns to the repository volume name <p>Enclose the name of the existing repository volume in double quotation marks (" ").</p> <p>If you want to create a new repository volume when you run this command you must enter the name of either a volume group or a disk pool in which you want the repository volume. Optionally, you can also define the capacity of the repository volume. If you want to define the capacity you can use these values:</p> <ul style="list-style-type: none"> ■ An integer value that represents a percentage of the base volume capacity ■ A decimal fraction value that represents a percentage of the base volume capacity ■ A specific size for the repository volume. Size is defined in units of bytes, KB, MB, GB, or TB. <p>If you do not use the capacity option, the storage management software sets the capacity to 20 percent of the base volume capacity.</p> <p>When you run this command the storage management software creates the repository volume for the snapshot volume.</p>
repositoryFullLimit	<p>The percentage of repository capacity at which you receive a warning that the snapshot repository volume is nearing full. Use integer values. For example, a value of 70 means 70 percent. The default value is 75.</p>

Notes

You can use any combination of alphanumeric characters, underscore (`_`), hyphen (`-`), and pound (`#`) for the names. Names can have a maximum of 30 characters.

The identifier of a snapshot image has two parts separated by a colon (`:`):

- The name of the snapshot group
- The identifier of the snapshot image

For example, if you want to create a snapshot volume named `snapData1` using the most recent snapshot image in a snapshot group that has the name `snapGroup1` with a maximum fill limit of 80 percent for the repository volume, you would use this command:

```
create snapVolume userLabel="snapData1"  
snapImageID="snapGroup1:newest"  
repositoryVolume="repos_1234" repositoryFullLimit=80;
```

The repository volume identifier is automatically created by the storage management software and the firmware when you create a new snapshot group. You cannot rename the repository volume because renaming the repository volume breaks the linkage with the snapshot images.

Minimum Firmware Level

7.83

Create SSD Cache

This command creates a read cache for a storage array using Solid State Disks (SSDs). Using high performance SSDs to cache read data improves the application I/O performance and response times, and delivers sustained performance improvement across different workloads, especially for high-IOP workloads. SSD cache maximizes the use of expensive fast SSDs.

SSD cache works in addition to the primary cache in the controller DRAM. With controller cache, the data is stored in DRAM after a host read. With SSD cache, the data is copied from user-specified base volumes, and then cached on SSDs.

Syntax

```
create ssdCache userLabel="ssdCacheName"  
drives=( trayID1,drawerID1,slotID1 ...  
trayIDn,drawerIDn,slotIDn)  
[updateExistingVolumes=( TRUE | FALSE )]
```

Parameters

Parameter	Description
userLabel	The alphanumeric identifier (including - and _) that you want to name the new SSD cache. Enclose the identifier in double quotation marks (" "). You can use any combination of alphanumeric characters, underscore (_), hyphen (-), and pound (#) for the identifier. Identifiers can have a maximum of 30 characters.
drives	The drives that you want to use to create the SSD cache. For high-capacity drive trays, specify the tray ID value, the drawer ID value, and the slot ID value for each drive that you want to assign to the SSD cache. For low-capacity drive trays, specify the tray ID value and the slot ID value for each drive that you assign to the SSD cache. Tray ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the tray ID values, the drawer ID values, and the slot ID values in parentheses.
updateExistingVolumes	This optional parameter specifies whether the SSD cache should be enabled for all existing volumes in the storage array. To enable the SSD cache for all existing volumes, set this parameter to TRUE. If you want to enable the SSD cache for individual volumes later, set this parameter to FALSE. The default value is TRUE.

Notes

A storage array can have only one SSD cache.

Only volumes created using hard disk drives can use the SSD cache. You cannot enable SSD cache on snapshot (legacy) volumes or snapshot images.

SSDs do not currently support full disk encryption (FDE). Therefore, volumes that have FDE enabled cannot use the SSD cache.

If all of the SSDs in the SSD cache are data assurance (DA)-capable and the DA premium feature is enabled, DA is automatically enabled for the SSD cache and cannot be disabled. In addition, you cannot add non-DA capable SSDs to a DA-enabled SSD cache.

Minimum Firmware Level

7.84

Create Storage Array Security Key

This command creates a new security key for a storage array that has full disk encryption (FDE) drives. This command also sets the security definitions and sets the state to Security Enabled. To use this command successfully, you need to have enough FDE drives to create at least one volume group or one disk pool.

NOTE Before you create a storage array security key, you must set the password for the storage array. Use the `set storageArray` command to set the password for the storage array.

Syntax

```
create storageArray securityKey  
[keyIdentifier="keyIdentifierString" |  
passPhrase="passPhraseString" |  
file="fileName" |  
commitSecurityKey=(TRUE | FALSE)
```

Parameters

Parameter	Description
keyIdentifier	A character string that you can read that is a wrapper around a security key. Enclose the key identifier in double quotation marks (" ").
passPhrase	A character string that encrypts the security key so that you can store the security key in an external file. Enclose the pass phrase in double quotation marks (" "). For information about the correct form for creating a valid pass phrase, refer to the Notes in this command description.
file	The file path and the file name to which you want to save the security key. For example: file="C:\Program Files\CLI\sup\seckey.slk" IMPORTANT – If you are creating the security key for storage array running the storage management software GUI, you must add a file extension of .slk to the end of the file name. If you are creating the security key on a storage array that is not running the storage management software GUI, you can use any file extension or no file extension. Enclose the file path and name in double quotation marks (" ").

Parameter	Description
<code>commitSecurityKey</code>	This parameter commits the security key identifier to the storage array for all FDE drives as well as the controllers. After the security key identifier is committed, a key is required to read data or write data. The data can only be read or changed by using a key, and the drive can never be used in a non-secure mode without rendering the data useless or totally erasing the drive.

Notes

Use this command for local key management only.

The controller firmware creates a lock that restricts access to the FDE drives. FDE drives have a state called Security Capable. When you create a security key, the state is set to Security Enabled, which restricts access to all FDE drives that exist within the storage array.

You can have a storage array configuration with more than one set of encrypted volume groups or disk pools. Each volume group or disk pool can have a unique security key. The character string generated by the `keyIdentifier` parameter is a string that you can read and that enables you to identify the security key that you need. You can create a `keyIdentifier` by using one of these methods:

- You can enter up to 189 alphanumeric characters for a key identifier. The key identifier cannot have these characters:
 - White spaces
 - Punctuation
 - Symbols
- If you do not enter the `keyIdentifier` parameter, the controller automatically generates the `keyIdentifier` parameter.

Additional characters are automatically generated and appended to the end of the string that you enter for the key identifier. If you do not enter any string for the `keyIdentifier` parameter, the key identifier consists of only the characters that are automatically generated.

Your pass phrase must meet these criteria:

- The pass phrase must be between eight and 32 characters long.
- The pass phrase must contain at least one uppercase letter.
- The pass phrase must contain at least one lowercase letter.
- The pass phrase must contain at least one number.
- The pass phrase must contain at least one non-alphanumeric character, for example, `<>@+`.

NOTE If your pass phrase does not meet these criteria, you will receive an error message and will be asked to retry the command.

Minimum Firmware Level

7.40

Create Synchronous Mirroring

This command creates both the primary volume and the secondary volume for a synchronous remote mirrored pair. This command also sets the write mode (synchronous write mode or asynchronous write mode) and the synchronization priority.

NOTE In previous versions of this command the feature identifier was `remoteMirror`. This feature identifier is no longer valid and is replaced by `syncMirror`.

Syntax

```
create syncMirror primary="primaryVolumeName"
secondary="secondaryVolumeName"
(remoteStorageArrayName="storageArrayName" |
remoteStorageArrayWwn="wwID" )
[remotePassword="password"
syncPriority=(highest | high | medium | low | lowest)
autoResync=(enabled | disabled)
writeOrder=(preserved | notPreserved)
writeMode=(synchronous | asynchronous)]
```

Parameters

Parameter	Description
<code>primary</code>	The name of an existing volume on the local storage array that you want to use for the primary volume. Enclose the primary volume name in double quotation marks (" ").
<code>secondary</code>	The name of an existing volume on the remote storage array that you want to use for the secondary volume. Enclose the secondary volume name in double quotation marks (" ").
<code>remoteStorageArrayName</code>	The name of the remote storage array. Enclose the remote storage array name in double quotation marks (" ").
<code>remoteStorageArrayWwn</code>	The World Wide Identifier (WWID) of the remote storage array. Enclose the WWID in double quotation marks (" ").
<code>remotePassword</code>	The password for the remote storage array. Use this parameter when the remote storage array is password protected. Enclose the password in double quotation marks (" ").

Parameter	Description
<code>syncPriority</code>	The priority that full synchronization has relative to host I/O activity. Valid values are <code>highest</code> , <code>high</code> , <code>medium</code> , <code>low</code> , or <code>lowest</code> .
<code>autoResync</code>	The settings for automatic resynchronization between the primary volumes and the secondary volumes of a remote-mirror pair. This parameter has these values: <ul style="list-style-type: none"> ■ <code>enabled</code>– Automatic resynchronization is turned on. You do not need to do anything further to resynchronize the primary volume and the secondary volume. ■ <code>disabled</code>– Automatic resynchronization is turned off. To resynchronize the primary volume and the secondary volume, you must run the <code>resume syncMirror</code> command.
<code>writeOrder</code>	The write order for data transmission between the primary volume and the secondary volume. Valid values are <code>preserved</code> or <code>notPreserved</code> .
<code>writeMode</code>	How the primary volume writes to the secondary volume. Valid values are <code>synchronous</code> or <code>asynchronous</code> .

Notes

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

When you choose the primary volume and the secondary volume, the secondary volume must be of equal or greater size than the primary volume. The RAID level of the secondary volume does not have to be the same as the primary volume.

Product shipments using the CE6998 or CE7900 controller define a maximum of 128 remote mirrors. The CDE3992 and CDE3994 controllers can define a maximum of 64 remote mirrors.

Passwords are stored on each storage array in a management domain. If a password was not previously set, you do not need a password. The password can be any combination of a alphanumeric characters with a maximum of 30 characters. (You can define a storage array password by using the `set storageArray` command.)

Synchronization priority defines the amount of system resources that are used to synchronize the data between the primary volume and the secondary volume of a mirror relationship. If you select the highest priority level, the data synchronization uses the most system resources to perform the full synchronization, which decreases performance for host data transfers.

Create Volume in Disk Pool

The `writeOrder` parameter applies only to asynchronous write modes and makes the mirrored pair part of a consistency group. Setting the `writeOrder` parameter to `preserved` causes the remote mirrored pair to transmit data from the primary volume to the secondary volume in the same order as the host writes to the primary volume. In the event of a transmission link failure, the data is buffered until a full synchronization can occur. This action can require additional system overhead to maintain the buffered data, which slows operations. Setting the `writeOrder` parameter to `notPreserved` frees the system from having to maintain data in a buffer, but it requires forcing a full synchronization to make sure that the secondary volume has the same data as the primary volume.

Minimum Firmware Level

6.10

This command creates a new standard RAID volume or a thin volume in an existing disk pool.

NOTE Some parameters for creating a standard RAID volume in a volume group are not compatible for creating volumes of any type in a disk pool. When using older scripts to create volumes in disk pools, make sure that all of the parameters are valid for disk pools. Invalid parameters prevent the scripts from running correctly and cause an error to be posted.

Syntax for Creating a Standard Volume

```
create volume diskPool="diskPoolName"
userLabel="volumeName"
capacity=volumeCapacity
[thinProvisioned=(TRUE | FALSE) |
owner=(a|b) |
mapping=(none|default) |
dataAssurance=(none|enabled) |
cacheReadPrefetch=(TRUE | FALSE)]
```

Syntax for Creating a Thin Provisioned Volume

```

create volume diskPool="diskPoolName"
userLabel="volumeName"
capacity=volumeCapacity
[thinProvisioned=(TRUE | FALSE) |
owner=(a|b) |
mapping=(none|default) |
dataAssurance=(none|enabled) |
(existingRepositoryLabel=existingRepositoryName |
newRepositoryCapacity=newRepositoryCapacityValue [KB | MB |
GB | TB | Bytes]) |
repositoryMaxCapacity=repositoryMaxCapacityValue[KB|MB|GB|TB
|Bytes] |
warningThresholdPercent=warningThresholdPercentValue |
repositoryExpansionPolicy=(automatic|manual) |
cacheReadPrefetch=(TRUE | FALSE)]
  
```

Parameters

Parameter	Description
diskPool	The name of the disk pool in which to create the new volume. Enclose the disk pool name in double quotation marks (" ").
userLabel	The name that you want to give the new volume. Enclose the volume name in double quotation marks (" ").
capacity	<p>The size of the volume that you are creating.</p> <p>For a standard RAID volume, the capacity that will be allocated for the volume.</p> <p>For a thin volume, the virtual capacity value that will be exposed by the thin volume.</p> <p>Size is defined in units of bytes, KB, MB, GB, or TB. The following are examples of the syntax:</p> <ul style="list-style-type: none"> ■ capacity=500MB ■ capacity=2GB
thinProvisioned	This parameter enables thin provisioning for the new volume. To use thin provisioning, set this parameter to TRUE. If you do not want thin provisioning, set this parameter to FALSE.

Parameter	Description
owner	The controller that owns the volume. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. If you do not specify an owner, the controller firmware determines the owner.
mapping	<p>This parameter enables you to map the volume to a host. If you want to map later, set this parameter to <code>none</code>. If you want to map now, set this parameter to <code>default</code>. The volume is mapped to all hosts that have access to the storage pool.</p> <p>The default value is <code>none</code>.</p>
dataAssurance	<p>The setting to specify that a volume group, and the volumes within the volume group, has data assurance protection to make sure that the data maintains its integrity. When you use this parameter, only protected drives can be used for the volume group. These settings are valid:</p> <ul style="list-style-type: none"> ■ <code>none</code> – The volume group does not have any explicit data assurance defined. ■ <code>enabled</code> – The volume group has data assurance protection. The volume group supports protected information and is formatted with protection information enabled. <p>The default value is <code>enabled</code> if the disk pool is data assurance capable.</p>
existingRepositoryLabel	This parameter identifies an existing repository for a thin volume. A repository volume has the physical capacity for a thin volume. This parameter applies only for thin provisioning. If you use the <code>existingRepositoryLabel</code> parameter, you must not use the <code>newRepositoryCapacity</code> parameter.
newRepositoryCapacity	<p>This parameter creates a new repository for a thin volume. A repository volume has the physical capacity for a thin volume. Use this parameter only if you set the value of the <code>thinProvisioned</code> parameter to <code>TRUE</code>.</p> <p>Size is defined in units of MB, GB, or TB. The following are examples of the syntax:</p> <ul style="list-style-type: none"> ■ <code>capacity=500MB</code> ■ <code>capacity=2GB</code> <p>The default value is 50% of the virtual capacity.</p>

Parameter	Description
<code>repositoryMaxCapacity</code>	<p>This parameter defines the maximum capacity for a repository for a thin volume. Use this parameter only if you set the value of the <code>thinProvisioned</code> parameter to <code>TRUE</code>.</p> <p>Size is defined in units of MB, GB, or TB. The following are examples of the syntax:</p> <ul style="list-style-type: none"> ■ <code>capacity=500MB</code> ■ <code>capacity=2GB</code>
<code>warningThresholdPercent</code>	<p>The percentage of thin volume capacity at which you receive a warning alert that the thin volume is nearing full. Use integer values. For example, a value of 70 means 70 percent.</p> <p>Valid values are from 1 to 100.</p> <p>Setting this parameter to 100 disables warning alerts.</p>
<code>repositoryExpansionPolicy</code>	<ul style="list-style-type: none"> ■ <code>automatic</code> ■ <code>manual</code>
<code>cacheReadPrefetch</code>	<ul style="list-style-type: none"> ■ <code>TRUE</code> ■ <code>FALSE</code>

Notes

Each volume name must be unique. You can use any combination of alphanumeric characters, underscore (`_`), hyphen (`-`), and pound (`#`) for the user label. User labels can have a maximum of 30 characters.

For thin volumes, the `capacity` value specifies the virtual capacity of the volume, and the `repositoryCapacity` specifies the capacity of the volume created as the repository volume. Use the `existingRepositoryLabel` to specify an existing unused repository volume instead of creating a new volume.

For best results when creating a thin volume, the repository volume must already exist or must be created in an already existing disk pool. If you do not specify some of the optional parameters when creating thin volumes the storage management software will attempt to create the repository volume. The most desirable candidate volume is a repository volume that already exists and that is within the size requirements. The next most desirable candidate volume is a new repository volume that is created in the disk pool free extent.

Repository volumes for thin volumes cannot be created in volume groups.

Data Assurance Management

If `dataAssurance` parameter is set to `enabled`, only data assurance capable drives will be considered for volume candidates, otherwise both data assurance capable and non-data assurance capable drives will be considered. If only data assurance drives are available the new volume will be created using the enabled `dataAssurance` drives.

If `dataAssurance` parameter is set to `none`, data assurance is not defined. The volume possibly will or will not have data assurance, depending on the availability of data assurance enabled drives. You can enable data assurance only on volumes that support data assurance. If only data assurance drives are available, the volumes will be data assurance enabled.

Minimum Firmware Level

7.83

Create Volume Copy

This command creates a volume copy and starts the volume copy operation. This command is valid for both snapshot (legacy) volume copy pairs and snapshot image volume copy pairs.

ATTENTION Starting a volume copy operation overwrites all existing data on the target volume, makes the target volume read-only to hosts, and fails all snapshot (legacy) volumes or snapshot image volumes associated with the target volume, if any exist. If you have used the target volume as a copy before, be sure you no longer need the data or have it backed up.

This command creates volume copies in two ways:

- Volume copy without snapshot (legacy) also called *offline* volume copy
- Volume copy with either snapshot (legacy) or snapshot image, also called *online* volume copy

If you use volume copy without either snapshot (legacy) or snapshot image, you cannot write to the source volume until the copy operation is complete. If you want to be able to write to the source volume before the copy operation is complete, use volume copy with snapshot (legacy) or snapshot image. You can select volume copy with snapshot (legacy) or snapshot image through the optional parameters in the command syntax.

After completion of the volume copy with snapshot (legacy) operation, the snapshot (legacy) is disabled. After completion of the volume copy with snapshot image operation, the snapshot image is deleted and the snapshot volume is disabled.

NOTE You can have a maximum of eight volume copies in progress at one time. If you try to create more than eight volume copies at one time, the controllers return a status of Pending until one of the volume copies that is in progress finishes and returns a status of Complete.

Syntax

```
create volumeCopy source="sourceName"
target="targetName"
[copyPriority=(highest | high | medium | low | lowest)
targetReadOnlyEnabled=(TRUE | FALSE)
copyType=(offline | online)
repositoryPercentOfBase=(20 | 40 | 60 | 120 | default) |
repositoryGroupPreference=(sameAsSource | otherThanSource |
default)]
```

Parameters

Parameter	Description
source	The name of an existing volume that you want to use as the source volume. Enclose the source volume name in double quotation marks (" ").
target	The name of an existing volume that you want to use as the target volume. Enclose the target volume name in double quotation marks (" ").
copyPriority	The priority that volume copy has relative to host I/O activity. Valid values are highest, high, medium, low, or lowest.
targetReadOnlyEnabled	The setting so that you can write to the target volume or only read from the target volume. To write to the target volume, set this parameter to FALSE. To prevent writing to the target volume, set this parameter to TRUE.
copyType	Use this parameter to create a volume copy with a snapshot (legacy) or snapshot image. Creating a volume copy with a snapshot (legacy) or snapshot image enables you to continue to write to the source volume while creating the volume copy. To create a volume copy with a snapshot (legacy) or snapshot image, set this parameter to online. To create a volume copy without a snapshot (legacy) or snapshot image, set this parameter to offline. If you do not use this parameter, the volume copy is created without a snapshot (legacy) or snapshot image.

Parameter	Description
repositoryPercentOfBase	<p>This parameter determines the size of the repository volume for the snapshot (legacy) or snapshot image when you are creating a volume copy with a snapshot (legacy) or snapshot image. The size of the repository volume is expressed as a percentage of the source volume, which is also called the base volume. Valid values for this parameter are 20, 40, 60, 120, and default. The default value is 20. If you do not use this parameter, the firmware uses a value of 20 percent.</p> <p>You must use the <code>copyType</code> parameter with the <code>repositoryPercentOfBase</code> parameter.</p>
repositoryGroupPreference	<p>This parameter determines to which volume group the snapshot (legacy) repository volume or snapshot image repository volume is written. You have these choices:</p> <ul style="list-style-type: none"> ■ <code>sameAsSource</code> – The snapshot (legacy) repository volume or snapshot image volume is written to the same volume group as the source volume if space is available. ■ <code>otherThanSource</code> – The snapshot (legacy) repository volume is written to a different volume group. Firmware determines which volume group based on available space on the volume groups. ■ <code>default</code> – The snapshot (legacy) repository volume or snapshot image repository volume is written to any volume group that has space. <p>For best performance, use the <code>sameAsSource</code> option.</p> <p>You must use the <code>copyType</code> parameter with the <code>repositoryGroupPreference</code> parameter.</p>

Notes

You can use any combination of alphanumeric characters, happens, and underscores for the names. Names can have a maximum of 30 characters.

Copy priority defines the amount of system resources that are used to copy the data between the source volume and the target volume of a volume copy pair. If you select the highest priority level, the volume copy uses the most system resources to perform volume copy, which decreases performance for host data transfers.

Minimum Firmware Level

5.40

7.77 adds creating a volume copy with snapshot (legacy).

Create Volume Group

This command creates either a free-capacity volume group or a volume group with one volume when you enter a set of unassigned drives.

Syntax

```
create volumeGroup
drives=(trayID1,drawerID1,slotID1 ...
trayIDn,drawerIDn,slotIDn)
raidLevel=(0 | 1 | 3 | 5 | 6)
userLabel="volumeGroupName"
[driveMediaType=(HDD | SSD | unknown | allMedia)
driveType=(fibre | SATA | SAS)
trayLossProtect=(TRUE | FALSE)
drawerLossProtect=(TRUE | FALSE)
securityType=(none | capable | enabled)
dataAssurance=(none | enabled)]
```

Parameters

Parameter	Description
drives	The drives that you want to assign to the volume group that you want to create. For high-capacity drive trays, specify the tray ID value, the drawer ID value, and the slot ID value for each drive that you assign to the volume group. For low-capacity drive trays, specify the tray ID value and the slot ID value for each drive that you assign to the volume group. Tray ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the tray ID values, the drawer ID values, and the slot ID values in parentheses.
raidLevel	The RAID level of the volume group that contains the volume. Valid values are 0, 1, 3, 5, or 6.
userLabel	The alphanumeric identifier (including - and _) that you want to give the new volume group. Enclose the volume group identifier in double quotation marks (" ").

Parameter	Description
driveMediaType	<p>The type of drive media that you want to use for the volume group</p> <p>You must use this parameter when you have more than one type of drive media in your storage array.</p> <p>Valid drive media are:</p> <ul style="list-style-type: none"> ■ HDD – Use this option when you have hard drives in the drive tray. ■ SSD – Use this option when you have solid state drives in the drive tray. ■ unknown – Use if you are not sure what types of drive media are in the drive tray. ■ allMedia – Use this option when you want to use all types of drive media that are in the drive tray.
driveType	<p>The type of drive that you want to use in the volume group. You cannot mix drive types.</p> <p>You must use this parameter when you have more than one type of drive in your storage array.</p> <p>Valid drive types are :</p> <ul style="list-style-type: none"> ■ fibre ■ SATA ■ SAS <p>If you do not specify a drive type, the command defaults to fibre.</p>
trayLossProtect	<p>The setting to enforce tray loss protection when you create the volume group. To enforce tray loss protection, set this parameter to TRUE. The default value is FALSE.</p>
drawerLossProtect	<p>The setting to enforce drawer loss protection when you create the volume group. To enforce drawer loss protection, set this parameter to TRUE. The default value is FALSE.</p>
securityType	<p>The setting to specify the security level when creating the volume groups and all associated volumes. These settings are valid:</p> <ul style="list-style-type: none"> ■ none – The volume group and volumes are not secure. ■ capable – The volume group and volumes are capable of having security set, but security has not been enabled. ■ enabled – The volume group and volumes have security enabled.

Parameter	Description
dataAssurance	<p>The setting to specify that a volume group, and the volumes within the volume group, has data assurance protection to make sure that the data maintains its integrity. When you use this parameter, only protected drives can be used for the volume group. These settings are valid:</p> <ul style="list-style-type: none"> ■ none – The volume group does not have data assurance protection. ■ enabled – The volume group has data assurance protection. The volume group supports protected information and is formatted with protection information enabled.

Notes

The `drives` parameter supports both high-capacity drive trays and low-capacity drive trays. A high-capacity drive tray has drawers that hold the drives. The drawers slide out of the drive tray to provide access to the drives. A low-capacity drive tray does not have drawers. For a high-capacity drive tray, you must specify the identifier (ID) of the drive tray, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive tray, you need only specify the ID of the drive tray and the ID of the slot in which a drive resides. For a low-capacity drive tray, an alternative method for identifying a location for a drive is to specify the ID of the drive tray, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

If you do not specify a capacity by using the `capacity` parameter, all of the drive capacity that is available in the volume group is used. If you do not specify capacity units, `bytes` is used as the default value.

Cache Read Prefetch

The `cacheReadPrefetch` command lets the controller copy additional data blocks into cache while the controller reads and copies data blocks that are requested by the host from the drives into cache. This action increases the chance that a future request for data can be fulfilled from cache. Cache read prefetch is important for multimedia applications that use sequential data transfers. The configuration settings for the storage array that you use determine the number of additional data blocks that the controller reads into cache. Valid values for the `cacheReadPrefetch` parameter are `TRUE` or `FALSE`.

You do not need to enter a value for the `cacheReadPrefetch` parameter or the `segmentSize` parameter. If you do not enter a value, the controller firmware uses the `usageHint` parameter with `fileSystem` as the default value. Entering a value for the `usageHint` parameter and a value for the `cacheReadPrefetch` parameter or a value for the `segmentSize` parameter does not cause an error. The

value that you enter for the `cacheReadPrefetch` parameter or the `segmentSize` parameter takes priority over the value for the `usageHint` parameter.

Segment Size

The size of a segment determines how many data blocks that the controller writes on a single drive in a volume before writing data on the next drive. Each data block stores 512 bytes of data. A data block is the smallest unit of storage. The size of a segment determines how many data blocks that it contains. For example, an 8-KB segment holds 16 data blocks. A 64-KB segment holds 128 data blocks.

When you enter a value for the segment size, the value is checked against the supported values that are provided by the controller at run time. If the value that you entered is not valid, the controller returns a list of valid values. Using a single drive for a single request leaves other drives available to simultaneously service other requests.

If the volume is in an environment where a single user is transferring large units of data (such as multimedia), performance is maximized when a single data transfer request is serviced with a single data stripe. A data stripe is the segment size that is multiplied by the number of drives in the volume group that are used for data transfers. In this case, multiple drives are used for the same request, but each drive is accessed only once.

For optimal performance in a multiuser database or file system storage environment, set your segment size to minimize the number of drives that are required to satisfy a data transfer request.

Security Type

The `securityType` parameter is valid for drives that are capable of full disk encryption (FDE). With FDE, the controller firmware can create a key and activate the Drive Security feature. The Drive Security feature encrypts data as the data is written to the drive and decrypts the data as the data is read from the drive. Without the key created by the controller, the data written to the drive is inaccessible.

Before you can set the `securityType` parameter to `capable` or `enabled`, you must create a storage array security key. Use the `create storageArray securityKey` command to create a storage array security key. These commands are related to the security key:

- `create storageArray securityKey`
- `enable volumeGroup [volumeGroupName] security`
- `export storageArray securityKey`
- `import storageArray securityKey`
- `set storageArray securityKey`
- `start secureErase (drive | drives)`

Tray Loss Protection and Drawer Loss Protection

For tray loss protection to work, each drive in a volume group must be on a separate tray. If you set the `trayLossProtect` parameter to `TRUE` and have selected more than one drive from any one tray, the storage array returns an error. If you set the `trayLossProtect` parameter to `FALSE`, the storage array performs operations, but the volume group that you create might not have tray loss protection.

Tray loss protection is not valid when you create volumes on existing volume groups.

The `drawerLossProtect` parameter defines if data on a volume is accessible if a drawer fails. When you assign the drives, if you set the `drawerLossProtect` parameter to `TRUE` and select more than one drive from any one drawer, the storage array returns an error. If you set the `drawerLossProtect` parameter to `FALSE`, the storage array performs operations, but the volume group that you create might not have drawer loss protection.

You must set the `trayLossProtect` parameter and the `drawerLossProtect` parameter to the same value. Both of the parameters must be either `TRUE` or `FALSE`. If the `trayLossProtect` parameter and the `drawerLossProtect` parameter are set to different values, the storage array returns an error.

Data Assurance Management

If `dataAssurance` parameter is set to `enabled`, only data assurance capable drives will be considered for volume candidates, otherwise both data assurance capable and non-data assurance capable drives will be considered. If only data assurance drives are available the new volume group will be created using the enabled data assurance drives.

If `dataAssurance` parameter is set to `none`, data assurance is not defined. The volume group will be comprised of volumes that possibly will or will not have data assurance, depending on the availability of data assurance enabled drives. You can enable data assurance only on volumes created on volume groups that support data assurance. If only data assurance drives are available, the volume groups will be data assurance enabled.

Minimum Firmware Level

7.10

7.50 adds the `securityType` parameter.

7.60 adds the `drawerID` user input, the `driveMediaType` parameter, and the `drawerLossProtect` parameter.

7.75 adds the `dataAssurance` parameter.

Deactivate Asynchronous Mirroring

This command deactivates the Asynchronous Mirroring premium feature.

NOTE All existing asynchronous mirror groups or asynchronous mirrored pairs must be deleted from the local storage array and the remote storage array before the Asynchronous Mirroring feature can be deactivated.

Syntax

```
deactivate storageArray feature=asyncRemoteMirror
```

Parameters

None.

Minimum Firmware Level

7.84

Deactivate Synchronous Mirroring

This command deactivates the Synchronous Mirroring premium feature, disassembles the mirror repository volume, and releases the controller owner of the secondary volume. The controller host port that is dedicated to the secondary volume is available for host data transfers.

NOTE In previous versions of this command the feature identifier was `remoteMirror`. This feature identifier is no longer valid and is replaced by `syncMirror`.

Syntax

```
deactivate storageArray feature=syncMirror
```

Parameters

None.

Minimum Firmware Level

6.10

Delete Asynchronous Mirror Group

This command deletes one or more asynchronous mirror groups from the local storage array and the remote storage array.

NOTE The asynchronous mirror group must be empty before it can be successfully deleted. You must remove all asynchronous mirrored pairs from the asynchronous mirror group before using this command.

Syntax

```
delete asyncMirrorGroup
(allAsyncMirrorGroups |
asyncMirrorGroup["asyncMirrorGroupName"] |
asyncMirrorGroups
["asyncMirrorGroupName_01" "asyncMirrorGroupName_02"])
```

Parameter

Parameter	Description
allAsyncMirrorGroups	Use this parameter if you want to remove all asynchronous mirror groups from the local storage array and the remote storage array.
asyncMirrorGroup	The name of the asynchronous mirror group that you want to delete. Enclose the asynchronous mirror group name in square brackets ([]). If the asynchronous mirror group name has special characters, you also must enclose the asynchronous mirror group name in double quotation marks (" ").
asyncMirrorGroups	The names of several asynchronous mirror groups that you want to delete. Enter the names of the asynchronous mirror groups using these rules: <ul style="list-style-type: none">■ Enclose all of the names in square brackets ([]).■ Enclose each of the names in double quotation marks (" ").■ Separate each of the names with a space.

Minimum Firmware Level

7.84

Delete Consistency Group

This command deletes a snapshot consistency group. This command works in two ways:

- You can delete both the consistency group and the repository volumes contained by the consistency group.
- You can delete only the consistency group and leave the repository volumes that are contained by the consistency group intact.

Syntax

```
delete consistencyGroup [consistencyGroupName
[deleteRepositoryMembers=(TRUE | FALSE)]]
```

Parameter

Parameter	Description
consistencyGroup	The name of the snapshot consistency group that you want to delete. Enclose the snapshot consistency group name in square brackets ([]). If the snapshot consistency group name has special characters, you also must enclose the snapshot consistency group name in double quotation marks (" ").
deleteRepositoryMembers	The setting to delete on or retain the repository volumes. To delete the repository volumes, set this parameter to TRUE. To retain the repository volumes, set this parameter to FALSE. The default setting is FALSE.

Minimum Firmware Level

7.83

Delete Consistency Group Snapshot Image

This command deletes the snapshot images in a consistency group.

Syntax

```
delete cgSnapImage consistencyGroup="consistencyGroupName"  
[(deleteCount=numberOfSnapImages |  
retainCount=numberOfSnapImages) |  
ignoreSnapVolume=(TRUE | FALSE)
```

Parameters

Parameter	Description
consistencyGroup	The name of the consistency group from which you want to delete the snapshot images. Enclose the consistency group name in double quotation marks (" ").
deleteCount	The number of snapshot images that you want to delete from the consistency group. Use integer values. This parameter deletes the oldest snapshot image first and continues to delete the oldest snapshot images until reaching the number that you enter.
retainCount	The number of snapshot images that you want to keep in the consistency group. Use integer values. This parameter keeps the most recent snapshot images in the consistency group.

Parameter	Description
ignoreSnapVolume	Determines whether the associated consistency group snapshot volume is kept or deleted. This parameter applies only if the consistency group snapshot image is associated with a consistency group snapshot volume. To keep the snapshot volume, set this parameter to TRUE. To delete the snapshot volume, set this parameter to FALSE. The default value is FALSE.

Notes

If the snapshot images cannot be deleted for all of the relevant member volumes of the consistency group, the operation fails and none of the snapshot images are deleted.

When you delete a consistency group snapshot image that is associated with a consistency group snapshot volume, the corresponding snapshot volume member in the consistency group snapshot volume are transitioned to the Stopped state. A snapshot volume member in the Stopped state no longer has a relationship to the snapshot group of the deleted snapshot image. However, a snapshot volume member in the Stopped state keeps its relationship to its consistency group snapshot volume.

Minimum Firmware Level

7.83

Delete Consistency Group Snapshot Volume

This command deletes the snapshot volume of a consistency group. Optionally, you can also delete the repository members.

Syntax

```
delete cgSnapVolume ["snapVolumeName"]
[deleteRepositoryMembers=(TRUE | FALSE)]
```

Parameters

Parameter	Description
cgSnapVolume	The name of the consistency group snapshot volume that you want to delete. Enclose the name of the consistency group snapshot volume in double quotation marks (") inside square brackets ([]).
deleteRepositoryMembers	The parameter to save or delete the member volumes. To save the member volumes, set this parameter to TRUE. To delete the member volumes, set this parameter to FALSE. The default value is TRUE.

Minimum Firmware Level

7.83

Delete Disk Pool

ATTENTION Possible damage to the storage array configuration – All of the data in the disk pool is lost as soon as you run this command.

This command deletes a disk pool. Depending on your version of the storage management software, this command also deletes all of the volumes in the disk pool. If your version of the storage management software does not support automatic deletion of the volumes, you can force the deletion of the disk pool and volumes.

Syntax

```
delete diskPool [diskPoolName]  
[force=(TRUE | FALSE)]
```

Parameter

Parameter	Description
diskPoolName	The name of the disk pool that you want to delete. Enclose the disk pool name in square brackets ([]). If the disk pool name has special characters, you also must enclose the host group name in double quotation marks (" ").
force	This parameter forces the deletion of volumes if your storage management software does not support automatic deletion of the volumes in the disk pool. To force the deletion of a disk pool and the volumes that it contains, set this parameter to TRUE. The default is FALSE.

Notes

Each disk pool name must be unique. You can use any combination of alphanumeric characters, underscore (_), hyphen(-), and pound (#) for the user label. User labels can have a maximum of 30 characters.

Minimum Firmware Level

7.83

Delete Host

This command deletes one or more hosts.

Syntax

```
delete (host [hostName] |  
hosts ["hostName1" ... "hostNameN"])
```

Parameters

Parameter	Description
host	The name of the host that you want to delete. Enclose the host name in square brackets ([]). If the host name has special characters, you also must enclose the host name in double quotation marks (" ").
hosts	The names of several hosts that you want to delete. Enter the names of the hosts using these rules: <ul style="list-style-type: none">■ Enclose all of the names in square brackets ([]).■ Enclose each of the names in double quotation marks (" ").■ Separate each of the names with a space.

Notes

A host is a computer that is attached to the storage array and accesses the volumes on the storage array through the host ports on the host.

Minimum Firmware Level

5.20

Delete Host Group

This command deletes a host group.

ATTENTION Possible damage to the storage array configuration – This command deletes all of the host definitions in the host group.

Syntax

```
delete hostGroup [hostGroupName]
```

Parameter

Parameter	Description
hostGroup	The name of the host group that you want to delete. Enclose the host group name in square brackets ([]). If the host group name has special characters, you also must enclose the host group name in double quotation marks (" ").

Notes

A host group is an optional topological element that is a collection of hosts that share access to the same volumes. The host group is a logical entity.

Minimum Firmware Level

5.20

Delete Host Port

This command deletes a host port identification. The identification is a software value that represents the physical host port to the controller. By deleting the identification, the controller no longer recognizes instructions and data from the host port.

Syntax

```
delete hostPort [hostPortName]
```

Parameter

Parameter	Description
hostPort	The name of the host port that you want to delete. Enclose the name of the host port in square brackets ([]).

Notes

A host port is a physical connection on a host adapter that resides within a host computer. A host port provides a host access to the volumes in a storage array.

Minimum Firmware Level

5.20

Delete iSCSI Initiator

This command deletes a specific iSCSI initiator object.

Syntax

```
delete iscsiInitiator (["iscsiID"] | ["name"])
```

Parameters

Parameter	Description
iscsiInitiator	The identifier of the iSCSI initiator that you want to delete. The identifier of the iSCSI initiator can be either an iSCSI ID or a unique name. Enclose the identifier in double quotation marks (" "). You must also enclose the iscsiID in either square brackets ([]) or angle brackets (<>).

Minimum Firmware Level

7.10

Delete Snapshot (Legacy) Volume

This command deletes one or more snapshot (legacy) volumes or snapshot (legacy) repository volumes. You can also use this command to remove schedules for creating snapshots (legacy).

ATTENTION Possible damage to the storage array configuration – All of the data in the volume is lost as soon as you run this command.

Syntax

```
delete (volume [volumeName] |  
volumes ["volumeName1"a ... "volumeNameN"] )  
[schedule]
```

Parameters

Parameter	Description
volume or volumes	The name of the snapshot (legacy) volume that you want to delete. Enclose the snapshot (legacy) volume name in square brackets ([]). If the snapshot (legacy) volume name has special characters, you also must enclose the snapshot (legacy) volume name in double quotation marks (" "). You can enter more than one volume name. Enclose all of the volume names in one set of square brackets ([]). Enclose each volume name in double quotation marks (" "). Separate each volume name with a white space.
schedule	This parameter deletes the schedule for a specific snapshot (legacy) volume. Only the schedule is deleted, the snapshot (legacy) volume remains.

Minimum Firmware Level

7.77

Delete Snapshot Group

ATTENTION Possible damage to the storage array configuration – All of the data in the snapshot group is lost as soon as you run this command.

This command deletes an entire snapshot group and optionally the associated repository volumes.

Syntax

```
delete snapGroup ["snapGroupName"]  
[deleteRepositoryMembers=(TRUE | FALSE)]
```

Parameters

Parameter	Description
snapGroup	The name of the snapshot group that you want to delete. Enclose the snapshot group name in double quotation marks (" ") inside square brackets ([]).
deleteRepositoryMembers	The parameter to delete or save the repository volumes. To delete the repository volumes, set this parameter to TRUE. To save the repository volumes, set this parameter to FALSE. The default value is FALSE .

Notes

You can use any combination of alphanumeric characters, underscore (`_`), hyphen (`-`), and pound (`#`) for the user label. User labels can have a maximum of 30 characters.

You can delete a snapshot group if it is empty or if it contains snapshot images. All of the snapshot images in the snapshot group are deleted along with the snapshot group. If any existing snapshot image within the snapshot group has an associated snapshot volume, each snapshot volume is stopped and detached from the snapshot image. When you delete a snapshot group the associated repository volume is also deleted. By default, all member volumes in the repository volume are retained as unused, unmapped standard volumes. To remove the member volumes set the `deleteRepositoryMembers` parameter to `TRUE`, or do not use this parameter. To keep the member volumes, set the `deleteRepositoryMembers` parameter to `FALSE`.

Minimum Firmware Level

7.83

Delete Snapshot Image

This command deletes one or more snapshot images from a snapshot group.

Syntax

```
delete snapImage (snapGroup="snapGroupName" |
snapGroups=("snapGroupName1"... "snapGroupName"))
(deleteCount=numberOfSnapImages |
retainCount=numberOfSnapImages |
ignoreSnapVolume=(TRUE | FALSE) |
snapImageID=OLDEST)
```

Parameters

Parameter	Description
<code>snapGroup</code> or <code>snapGroups</code>	<p>The name of the snapshot group that has the snapshot image that you want to delete. Enclose the name of the snapshot group in double quotation marks (" ").</p> <p>If you enter more than one snapshot group name, enclose each of the snapshot group names in double quotation marks (" "). Enclose all of the snapshot group names in parentheses.</p> <p>If you do not use any other parameters with the <code>snapGroup</code> parameter or <code>snapGroups</code> parameter, by default the oldest snapshot image is deleted.</p>

Parameter	Description
deleteCount	<p>The number of snapshot images that you want to delete from the snapshot group. Use integer values.</p> <p>This parameter deletes the oldest snapshot image first and continues to delete the oldest snapshot images until reaching the number that you enter.</p> <p>If the number that you enter is greater than the actual number of all of the snapshot images in the snapshot group, all of the snapshot images will be deleted. The snapshot group is left empty.</p>
retainCount	<p>The number of snapshot images that you want to keep in the snapshot group. Use integer values.</p> <p>This parameter keeps the most recent snapshot images in the snapshot group and deletes older snapshot images.</p> <p>If the number of existing snapshot images in the snapshot group is less than the number that you enter, none of the snapshot images are deleted.</p>
ignoreSnapVolume	<p>Use this parameter to make sure that you do not delete a snapshot image that has a snapshot volume associated with the snapshot image. You can use one of these values:</p> <ul style="list-style-type: none"> ■ TRUE - Use this value to delete the snapshot image even though the snapshot image has an associated snapshot volume. ■ FALSE - Use this value to keep the snapshot image when the snapshot image has an associated snapshot volume. <p>The default value is TRUE.</p>
snapImageID	<p>The snapImageID parameter accepts only the OLDEST option. This parameter deletes the earliest snapshot image created.</p>

Notes

You can delete the oldest snapshot image from a snapshot group repository volume. The definition of a snapshot image that you delete is removed from the system. The space occupied by the snapshot image that you delete from the snapshot group repository volume is released and made available for reuse within the snapshot group.

Any snapshot volumes that exist for a snapshot image transition to the Stopped state when the snapshot image is deleted.

This command will not run when the controller is in Lockdown mode.

Minimum Firmware Level

7.83

Delete Snapshot Volume

This command deletes a snapshot volume and optionally the associated snapshot repository members.

NOTE You cannot use this command for snapshot images involved in online volume copy.

Syntax

```
delete snapVolume ["snapVolumeName"]  
[deleteRepositoryMembers=(TRUE | FALSE)]
```

Parameters

Parameter	Description
snapVolume	The name of the snapshot volume that you want to delete. Enclose the snapshot volume identifier in double quotation marks (" ") inside square brackets ([]).
deleteRepositoryMembers	The parameter to save or delete the repository members. To save the repository members, set this parameter to FALSE. To delete the repository members, set this parameter to TRUE. The default value is TRUE. If you do not use this parameter the repository members are automatically deleted.

Minimum Firmware Level

7.83

Delete SSD Cache

This command deletes the SSD cache. All data in the SSD cache is purged.

Syntax

```
delete ssdCache [ssdCacheName]
```

Parameter

Parameter	Description
ssdCache	The alphanumeric identifier (including - and _) of the SSD cache that you want to delete. Enclose the identifier in square brackets ([]). If the SSD cache name contains special characters or consists only of numbers, you also must enclose the identifier in double quotation marks (" ") inside square brackets.

Minimum Firmware Level

7.84

Delete Volume

This command deletes one or more standard volumes, snapshot (legacy) volumes, or snapshot (legacy) repository volumes.

ATTENTION Possible damage to the storage array configuration – All of the data in the volume is lost as soon as you run this command.

Syntax

```
delete (allVolumes |  
volume [volumeName] |  
volumes [volumeName1 ... volumeNameN])  
[removeVolumeGroup=(TRUE | FALSE)]  
force=(TRUE | FALSE)]
```

Parameters

Parameter	Description
allVolumes	This parameter deletes all of the volumes in a storage array.
volume or volumes	The name of the specific volume for which you are retrieving information. Enclose the volume name in square brackets ([]). If the volume name has special characters, you also must enclose the volume name in double quotation marks (" ") inside square brackets. If the volume name consists only of numbers, such as "1002," you also must enclose the volume name in double quotation marks (" ") inside square brackets. You can enter more than one volume name. Enclose all of the volume names in one set of square brackets ([]). Enclose each volume name in double quotation marks (" "). Separate each volume name with a white space.
removeVolumeGroup	Deleting the last volume in a volume group does not delete the volume group. You can have a standalone volume group (minus any volumes). To remove the standalone volume group, set this parameter to TRUE. To keep standalone volume groups intact, set this parameter to FALSE.
force	Use this parameter to force the immediate deletion of a volume even if the controllers are performing other operations. To immediately force the deletion of a volume, set this parameter to TRUE. To wait until the controllers have finished performing other operations, do not use this parameter or set this parameter to FALSE.

Notes

When you use the `allVolumes` parameter, this command deletes volumes until all of the volumes are removed or until an error is encountered. If an error is encountered, this command does not try to delete the remaining volumes. Deleting volumes from different volume groups is possible. All of the volume groups that become empty are deleted if you set the `removeVolumeGroup` parameter to `TRUE`.

If you want to delete an entire volume group, you can also use the `delete volumeGroup` command.

Minimum Firmware Level

6.10

7.10 adds the `removeVolumeGroup` parameter.

Delete Volume from Disk Pool

This command deletes either normal or thin volumes from a disk pool. You also can use this command to delete any schedules related to the volume. When you delete the schedule the volume is not deleted.

ATTENTION Possible damage to the storage array configuration – All of the data in the volume is lost as soon as you run this command.

Syntax

```
delete (allVolumes |
volume [volumeName] |
volumes [volumeName1 ... volumeNameN] | allVolumes)
[removeVolumeGroup=(TRUE | FALSE)
force=(TRUE | FALSE)
schedule
retainRepositoryMembers=(TRUE | FALSE)]
```

Parameters

Parameter	Description
<code>allVolumes</code>	This parameter deletes all of the volumes on a disk pool.

Parameter	Description
volume or volumes	<p>The name of the volume that you want to delete. Enclose the volume name in square brackets ([]). If the volume name has special characters, you also must enclose the volume name in double quotation marks (" ") inside square brackets. If the volume name consists only of numbers, such as "1002," you also must enclose the volume name in double quotation marks (" ") inside square brackets.</p> <p>You can enter more than one volume name. Enclose all of the volume names in one set of double quotation marks (" ") inside square brackets ([]). Separate each volume name with a white space.</p>
removeVolumeGroup	<p>Deleting the last volume in a volume group does not delete the volume group. You can have a standalone volume group (minus any volumes). To remove the standalone volume group, set this parameter to TRUE. To keep standalone volume groups intact, set this parameter to FALSE.</p>
force	<p>Use this parameter to force the immediate deletion of a volume even if the controllers are performing other operations. To immediately force the deletion of a volume, set this parameter to TRUE. To wait until the controllers have finished performing other operations, do not use this parameter or set this parameter to FALSE.</p>
schedule	<p>This parameter deletes any schedule for related to a specific disk pool volume. Only the schedule is deleted, the disk pool volume remains.</p>
retainRepositoryMembers	<p>When you delete a thin volume, the associated repository volume is deleted by default. However, when the <code>retainRepositoryMembers</code> is set to TRUE, the repository volume is retained. For normal volumes, this parameter has no effect.</p>

Notes

When you use the `allVolumes` parameter, this command deletes volumes until all of the volumes are removed or until an error is encountered. If an error is encountered, this command does not try to delete the remaining volumes. Deleting volumes from different volume groups is possible. All of the volume groups that become empty are deleted if you set the `removeVolumeGroup` parameter to TRUE.

Minimum Firmware Level

7.83

Delete Volume Group

ATTENTION Possible damage to the storage array configuration – All of the data in the volume group is lost as soon as you run this command.

This command deletes an entire volume group and its associated volumes.

Syntax

```
delete volumeGroup [volumeGroupName]
```

Parameter

Parameter	Description
volumeGroup	The alphanumeric identifier (including - and _) of the volume group that you want to delete. Enclose the volume group identifier in square brackets ([]).

Minimum Firmware Level

6.10

Diagnose Controller

This command runs diagnostic tests on the controller. The diagnostic tests consist of loopback tests in which data is written to the drives and read from the drives.

Syntax

```
diagnose controller [(a | b)]  
loopbackDriveChannel=(allchannels | (1 | 2 | 3 | 4 | 5 | 6 |  
7 | 8))  
testID=(1 | 2 | 3 | discreteLines)  
[patternFile="filename"]
```

Parameters

Parameter	Description
controller	The controller on which you want to run the diagnostic tests. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the storage management software returns a syntax error.
loopbackDriveChannel	The drive channels on which you want to run the diagnostic tests. You can either choose to run the diagnostics on all channels or select a specific channel on which to run diagnostics. If you select a specific channel, valid values for the drive channels are 1, 2, 3, 4, 5, 6, 7, or 8.

Parameter	Description
testID	The identifier for the diagnostic test you want to run. The identifier and corresponding tests are as follows: <ul style="list-style-type: none"> ■ 1 – Read test ■ 2 – Write test ■ 3 – Data loop-back test ■ discreteLines – Discrete lines diagnostic test
patternFile	The file path and the file name that contains a data pattern that you want to use as test data. Enclose the file name of the data pattern in double quotation marks (" "). For example: file="C:\Program Files\CLI\sup\patfile.txt"

Notes

When you run a data loop-back test, you can optionally specify a file that contains a data pattern. If you do not specify a file, the controller firmware provides a default pattern.

Discrete lines are control lines and status lines that are connected between two controllers in a controller tray. The discrete lines diagnostic test lets each controller check that control signal transitions can be observed at the control inputs of the alternate controller. The discrete lines diagnostic test automatically runs after each power-cycle or each controller-reset. You can run the discrete lines diagnostic test after you have replaced a component that failed the initial discrete lines diagnostic test. This test applies only to the CE6998 controller tray and the CE7900 controller tray. The discrete lines diagnostic test returns one of these messages:

- When the discrete lines diagnostic test runs successfully, this message appears:

```
The controller discrete lines successfully
passed the diagnostic
test. No failures were detected.
```
- If the discrete lines diagnostic test fails, this message appears:

```
One or more controller discrete lines failed
the diagnostic test.
```
- If the CLI cannot run the discrete lines diagnostic test, the CLI returns Error 270, which means that the discrete lines diagnostic test could not start nor complete.

Minimum Firmware Level

6.10 adds the read test, the write test, and the data loop-back test.

6.14 adds the discrete lines diagnostic test.

7.30 adds the updated drive channel identifier.

Diagnose Synchronous Mirroring

This command tests the connection between the specified primary volumes and the mirror volumes on a storage array with the Synchronous Mirroring premium feature enabled.

NOTE In previous versions of this command the feature identifier was `remoteMirror`. This feature identifier is no longer valid and is replaced by `syncMirror`.

Syntax

```
diagnose syncMirror (primary [primaryVolumeName] |
primaries ["primaryVolumeName1" ... "primaryVolumeNameN"])
testID=connectivity
```

Parameter

Parameter	Description
primary or primaries	The name of the primary volume of the remote mirror pair that you want to test. Enclose the primary volume name in square brackets ([]). If the primary volume name has special characters, you also must enclose the primary volume name in double quotation marks (" "). You can enter more than one volume name. Enclose all of the volume names in one set of square brackets ([]). Enclose each volume name in double quotation marks (" "). Separate each volume name with a white space.

Minimum Firmware Level

6.10

Disable AutoSupport at the EMW Level SMcli Version

NOTE This command is an SMcli command, not a script command. You must run this command from a command line. You cannot run this command from the script editor in the storage management software.

This command turns off the AutoSupport (ASUP) bundle collection feature for all managed storage arrays.

Syntax

```
SMcli disable autoSupportFeature
```

Parameters

None.

Minimum Firmware Level

7.86

Disable External Security Key Management

This command disables external security key management for a storage array that has full disk encryption drives.

Syntax

```
disable storageArray externalKeyManagement  
file="fileName"  
passPhrase="passPhraseString"
```

Parameters

Parameter	Description
file	The file path and the file name that has the security key. For example: file="C:\Program Files\CLI\sup\seckey.slk" IMPORTANT – The file name must have an extension of .slk.
passPhrase	A character string that encrypts the security key so that you can store the security key in an external file.

Notes

Your pass phrase must meet these criteria:

- The pass phrase must be between eight and 32 characters long.
- The pass phrase must contain at least one uppercase letter.
- The pass phrase must contain at least one lowercase letter.
- The pass phrase must contain at least one number.
- The pass phrase must contain at least one non-alphanumeric character, for example, < > @ +.

NOTE If your pass phrase does not meet these criteria, you will receive an error message.

Minimum Firmware Level

7.70

Disable Snapshot (Legacy)

This command stops a copy-on-write operation. This command performs the same action as the legacy `stop snapshot` command.

Syntax

```
disableSnapshot (volume [volumeName] |  
volumes [volumeName1 ... volumeNameN])
```

Parameter

Parameter	Description
volume or volumes	<p>The name of the specific volume for which you want to stop a copy-on-write operation. You can enter more than one volume name.</p> <p>Enclose the volume names using one of these forms:</p> <ul style="list-style-type: none">■ On a Windows command line: <code>\"<i>volumeName</i>\"</code>■ In a Windows script engine window: <code>["<i>volumeName</i>"]</code>■ On a Linux command line: <code>\"<i>volumeName</i>\"</code>■ In a Linux script engine window: <code>["<i>volumeName</i>"]</code>

Notes

Names can be any combination of alphanumeric characters, underscore (`_`), hyphen (`-`), and pound (`#`). Names can have a maximum of 30 characters.

One technique for naming the snapshot (legacy) volume and the snapshot (legacy) repository volume is to add a hyphenated suffix to the original base volume name. The suffix distinguishes between the snapshot (legacy) volume and the snapshot (legacy) repository volume. For example, if you have a base volume with a name of Engineering Data, the snapshot (legacy) volume can have a name of Engineering Data-S1, and the snapshot (legacy) repository volume can have a name of EngineeringData-R1.

If you do not choose a name for either the snapshot (legacy) volume or the snapshot (legacy) repository volume, the storage management software creates a default name by using the base volume name. An example of the snapshot (legacy) volume name that the controllers might create is, if the base volume name is `aaa` and does not have a snapshot (legacy) volume, the default snapshot (legacy) volume name is `aaa-1`. If the base volume already has $n-1$ number of snapshot (legacy) volumes, the default name is `aaa- n` . An example of the snapshot (legacy) repository volume name that the controller might create is, if the base volume name is `aaa` and does not have a snapshot (legacy) repository volume, the default snapshot (legacy) repository volume name is `aaa-R1`. If the base volume already has $n-1$ number of snapshot (legacy) repository volumes, the default name is `aaa-R n` .

Minimum Firmware Level

6.10

Disable Storage Array Feature

This command disables a storage array premium feature. Run the `show storageArray` command to show a list of the feature identifiers for all enabled premium features in the storage array.

Syntax

```
disable storageArray (featurePack |  
feature=featureAttributeList)
```

The *feaatureAttributeList* can be one or more of these attribute values. If you enter more than one attribute value, separate the values with a white space.

- volumeCopy
- snapshot
- asyncMirror
- syncMirror
- mixedDriveTypes
- goldKey
- driveSecurity
- enterpriseSecurityKeyMgr
- thinProvisioning
- storagePartition[2|4|8|16|32|64|96|128|192|256|512|Max]

The form for this attribute value is a combination of the alphabetical term merged with the numerical value, as shown by this example: `storagePartition256`

- driveSlotLimit[16|24|32|48|60|64|72|96|112|10|128|136|144|180|192|256|272|300|360|384|448|480|Max]

The form for this attribute value is a combination of the alphabetical term merged with the numerical value, as shown by this example: `driveSlotLimit360`

- ssdCache
- dataAssurance
- SSDSupport
- highPerformanceTier
- raid6

NOTE The following attribute is for enabling the snapshot (legacy) feature only.

- snapshot[2|4|8|16]

The form for this attribute value is a combination of the alphabetical term merged with the numerical value, as shown by this example: `snapshot16`

NOTE The `remoteMirror` attribute is valid only for firmware versions before 7.84. Starting with firmware 7.84 the `remoteMirror` attribute is replaced by `syncMirror`.

- `remoteMirror[8|16|32|64|128]`

The form for this attribute value is a combination of the alphabetical term merged with the numerical value, as shown by this example: `remoteMirror128`

Parameters

None.

Notes

If you specify the `syncMirror` parameter, this command disables the Synchronous Mirroring premium feature and takes away the structure of the mirror repository volume.

If you disable the High Performance Tier feature, all of the included features are disabled.

Minimum Firmware Level

5.00

6.50 adds these attributes:

- `goldKey`
- `mixedDriveTypes`

7.60 adds the `SSDSupport` attribute.

7.70 adds the `syncMirror8` attribute. Firmware version 7.70 supports a maximum of eight remote mirrored pairs.

7.83 adds these attributes:

- `raid6`
- `snapshot`
- `thinProvisioning`

7.84 adds these attributes:

- `asyncMirror`
- `syncMirror`
- `ssdCache`

Display Automatic Support Bundle Collection Configuration

NOTE This command is an SMcli command, not a script command. You must run this command from a command line. You cannot run this command from the script editor in the storage management software

This command displays the automatic support bundle collection settings.

Syntax

```
SMcli -supportBundle auto show
```

Parameters

None.

Minimum Firmware Level

7.83

Display Automatic Support Bundle Collection Schedule

NOTE This command is an SMcli command, not a script command. You must run this command from a command line. You cannot run this command from the script editor in the storage management software

This command displays the schedule for collecting support bundles for all storage arrays.

Syntax

```
SMcli -supportBundle schedule show
```

Parameters

None.

Minimum Firmware Level

7.83

Download Drive Firmware

This command downloads a firmware image to a drive.

ATTENTION Possible damage to the storage array configuration – Downloading drive firmware incorrectly can result in damage to the drives or a loss of data access.

This command is intended for downloading a firmware image to only one drive at a time. If you use this command in a script, make sure that you use this command only once. If you use this command more than once, the operation can fail. You can download firmware images to all of the drives in a storage array at one time by using the `download storageArray driveFirmware` command.

Syntax

```
download drive [trayID,drawerID,slotID] firmware  
file="filename"
```

Parameters

Parameter	Description
<code>drive</code>	The drive to which you want to download the firmware image. For high-capacity drive trays, specify the tray ID value, the drawer ID value, and the slot ID value for each drive to which you want to download firmware. For low-capacity drive trays, specify the tray ID value and the slot ID value for each drive to which you want to download firmware. Tray ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the tray ID values, the drawer ID values, and the slot ID values in brackets ([]).
<code>file</code>	The file path and the file name of the file that contains the firmware image. Enclose the file path and the file name of the firmware image in double quotation marks (" "). For example: <code>file="C:\Program Files\CLI\dnld\drvfrm.dlp"</code> Valid file names have a .dlp extension.

Notes

Before trying to download drive firmware, take these precautions:

- Stop all I/O activity to the storage array before you download the firmware image. The `download drive` command blocks all I/O activity until the download finishes or fails; however, as a precaution, make sure that all I/O activity that might affect the drive is stopped.
- Make sure that the firmware image file is compatible with the drive tray. If you download a firmware image file that is not compatible with the drive tray that you have selected, the drive tray might become unusable.
- Do not make any configuration changes to the storage array while you download drive firmware. Trying to make a configuration change can cause the firmware download to fail and make the selected drives unusable.

When you download the firmware to the drives, you must provide the full path and file name to the firmware image that is stored on your system.

You can use `download drive` command to test the firmware on one drive before you install the firmware on all of the drives in a storage array. The download returns one of these statuses:

- Successful
- Unsuccessful With Reason
- Never Attempted With Reason

The `drive` parameter supports both high-capacity drive trays and low-capacity drive trays. A high-capacity drive tray has drawers that hold the drives. The drawers slide out of the drive tray to provide access to the drives. A low-capacity drive tray does not

have drawers. For a high-capacity drive tray, you must specify the identifier (ID) of the drive tray, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive tray, you need only specify the ID of the drive tray and the ID of the slot in which a drive resides. For a low-capacity drive tray, an alternative method for identifying a location for a drive is to specify the ID of the drive tray, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

Minimum Firmware Level

6.10

7.60 adds the `drawerID` user input.

Download Environmental Card Firmware

This command downloads environmental services module (ESM) firmware.

Syntax

```
download (allTrays | tray [trayID])
firmware file="filename"
```

Parameters

Parameter	Description
<code>allTray</code>	This parameter downloads new firmware to all of the trays in the storage array.
<code>tray</code>	The drive tray that contains the ESM card to which you want to load new firmware. Tray ID values are 0 to 99. Enclose the tray ID value in square brackets ([]).
<code>file</code>	The file path and the file name of the file that contains the firmware image. Enclose the file path and the file name of the firmware image in double quotation marks (" "). For example: file="C:\Program Files\CLI\dnld\esmfrm.esm" Valid file names have an .esm extension.

Notes

The `tray` parameter downloads new firmware to a specific drive tray. If you need to download new firmware to more than one drive tray, but not all drive trays, you must enter this command for each drive tray.

Minimum Firmware Level

5.20

Download Storage Array Drive Firmware

This command downloads firmware images to all of the drives in the storage array.

Syntax

```
download storageArray driveFirmware file="filename"  
[file="filename2"... file="filenameN"]
```

Parameter

Parameter	Description
file	The file path and the file name of the file that contains the firmware image. Enclose the file path and the file name of the firmware image in double quotation marks (" "). For example: file="C:\Program Files\CLI\dnld\sdrvfrm.dlp" Valid file names have a .dlp extension.

Notes

When you run this command, you can download more than one firmware image file to the drives in a storage array. The number of firmware image files that you can download depends on the storage array. The storage management software returns an error if you try to download more firmware image files than the storage array can accept.

You can schedule downloads for multiple drives at the same time, including multiple drives in a redundant volume group. Each firmware image file contains information about the drive types on which the firmware image runs. The specified firmware images can be downloaded only to a compatible drive. Use the `download drive firmware` command to download a firmware image to a specific drive.

The `download storageArray driveFirmware` command blocks all I/O activity until either download try has been made for each candidate drive or you run the `stop storageArray downloadDriveFirmware` command. When the `download storageArray driveFirmware` command finishes downloading the firmware image, each candidate drive shows the download status for each drive. One of these statuses is returned:

- Successful
- Unsuccessful With Reason
- Never Attempted With Reason

Minimum Firmware Level

5.20

Download Storage Array Firmware/NVSRAM

This command downloads firmware and, optionally, NVSRAM values for the storage array controller. If you want to download only NVSRAM values, use the `download storageArray NVSRAM` command.

Syntax

```
download storageArray firmware [, NVSRAM ]
file="filename" [, "NVSRAM-filename" ]
[downgrade=(TRUE | FALSE)]
[activateNow=(TRUE | FALSE)]
```

Parameters

Parameter	Description
NVSRAM	The setting to download a file with NVSRAM values when you download a firmware file. Do not include square brackets with this parameter. Include a comma after the <code>firmware</code> parameter.
file	The file path and the file name that contains the firmware. Enclose the file path and the file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\dnld\safrm.dlp" Valid file names have a .dlp extension.
NVSRAM-filename	The file path and the file name that contains the NVSRAM values. Enclose the NVSRAM file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\dnld\safrm.dlp" Valid file names have a .dlp extension. Include a comma before the file name downloading both firmware and NVSRAM.
downgrade	<p>ATTENTION Possible damage to the storage array configuration – Downloading a previous version of the controller firmware or NVSRAM incorrectly can result in damage to the controllers or a loss of data access. Contact your Technical Support Representative before using this parameter.</p> <hr/> <p>The setting to load firmware that is a previous version. The default value is FALSE. Set the <code>downgrade</code> parameter to TRUE if you want to download an earlier version of firmware.</p>

Parameter	Description
activateNow	The setting to activate the firmware image and the NVSRAM image. The default value is TRUE. If you set the activateNow parameter to FALSE, you must run the activate storageArray firmware command to activate the firmware values and the NVSRAM values at a later time.

Minimum Firmware Level

5.00

Download Storage Array NVSRAM

This command downloads the NVSRAM values for the storage array controller.

Syntax

```
download storageArray NVSRAM file="filename"
```

Parameter

Parameter	Description
file	The file path and the file name that contains the NVSRAM values. Enclose the NVSRAM file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\dnld\afm.dlp" Valid file names have a .dlp extension.

Minimum Firmware Level

6.10

Download Tray Configuration Settings

This command downloads the factory default settings to all of the drive trays in a storage array or to a specific drive tray in a storage array.

Syntax

```
download (allTrays | tray [trayID]) configurationSettings  
file="filename"
```

Parameters

Parameter	Description
allTray	This parameter downloads new firmware to all of the trays in the storage array.
tray	The drive tray that contains the ESM card to which you want to load new firmware. Tray ID values are 0 to 99. Enclose the tray ID value in square brackets ([]).

Parameter	Description
file	<p>The file path and the file name of the file that contains the firmware image. Enclose the file path and the file name of the firmware image in double quotation marks (" "). For example:</p> <pre>file="C:\Program Files\CLI\dnld\trayset.dlp"</pre> <p>Valid file names have a .dlp extension.</p>

Notes

The `tray` parameter downloads the factory default configuration settings to a specific drive tray. If you need to download the factory default configuration settings to more than one drive tray, but not all drive trays, you must enter this command for each drive tray.

Minimum Firmware Level

7.75

Enable AutoSupport at the EMW Level SMcli Version

NOTE This command is an SMcli command, not a script command. You must run this command from a command line. You cannot run this command from the script editor in the storage management software

This command turns on the AutoSupport (ASUP) bundle collection feature for all managed storage array and makes it possible to transmit the bundle to a predesignated technical support site. After you enable the ASUP feature, any ASUP capable storage array is automatically prepared to collect and send support related data to Technical Support. The data can then be used for remote troubleshooting and problem analysis.

Syntax

```
SMcli enable autoSupportFeature
```

Parameters

None.

Minimum Firmware Level

7.86

Enable Controller Data Transfer

This command revives a controller that has become quiesced while running diagnostics.

Syntax

```
enable controller [(a | b)] dataTransfer
```

Parameter

Parameter	Description
controller	The controller that you want to revive. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the storage management software returns a syntax error.

Minimum Firmware Level

6.10

Enable Disk Pool Security

This command converts a non-secure disk pool to a secure disk pool.

NOTE All of the drives that comprise the disk pool must be security capable.

Syntax

```
enable diskPool [diskPoolName] security
```

Parameter

Parameter	Description
diskPool	The name of the disk pool that you want to place in the Security Enabled state. Enclose the disk pool identifier in square brackets ([]).

Notes

Each disk pool name must be unique. You can use any combination of alphanumeric characters, underscore (_), hyphen(-), and pound (#) for the user label. User labels can have a maximum of 30 characters.

Minimum Firmware Level

7.83

Enable External Security Key Management

This command enables external security key management for a storage array that has full disk encryption drives.

Syntax

```
enable storageArray externalKeyManagement  
file="fileName" |  
passPhrase="passPhraseString"
```

Parameters

Parameter	Description
file	The file path and the file name that has the security key. Enclose the file path and the file name that has the security key in double quotation marks (" "). For example: file="C:\Program Files\CLI\sup\seckey.slk" IMPORTANT – The file name must have an extension of .slk.
passPhrase	A character string that encrypts the security key so that you can store the security key in an external file. Enclose the pass phrase character string in double quotation marks (" ").

Notes

Your pass phrase must meet these criteria:

- The pass phrase must be between eight and 32 characters long.
- The pass phrase must contain at least one uppercase letter.
- The pass phrase must contain at least one lowercase letter.
- The pass phrase must contain at least one number.
- The pass phrase must contain at least one non-alphanumeric character, for example, < > @ +.

NOTE If your pass phrase does not meet these criteria, you will receive an error message.

Minimum Firmware Level

7.70

Enable or Disable SSD Cache for a Volume

This command turns on or off caching using the SSD cache feature for a specific volume. The volume can be either a standard volume, a snapshot volume, or a consistency group snapshot volume.

Syntax Applicable to a Standard Volume

```
set volume ["volumeName"] ssdCacheEnabled=(TRUE | FALSE)
```

Syntax Applicable to a Snapshot Volume

```
set snapVolume ["snapVolumeName"] ssdCacheEnabled=(TRUE | FALSE)
```

Syntax Applicable to a Consistency Group Snapshot Volume

```
set cgSnapVolume ["cgSnapVolumeName"] ssdCacheEnabled=(TRUE | FALSE)
```

Parameters

Parameter	Description
volume or snapVolume or cgSnapVolume	The name of the specific volume for which you want to turn on or off the SSD cache. Enclose the volume name in double quotation marks (" ") inside of square brackets ([]).
ssdCacheEnabled	To turn on SSD cache, set this parameter to TRUE. To turn off SSD cache, set this parameter to FALSE.

Notes

You can turn the SSD cache on or off for only one volume at a time.

When you turn off SSD cache for a volume, the SSD cache for that volume is purged.

Minimum Firmware Level

7.84

Enable Storage Array Feature

This command enables a premium feature for either a permanent upgrade to the storage array or a trial period. This command performs one of these actions:

- Enables a feature key for a permanent upgrade of a feature
- Enables a feature key for a permanent upgrade of a feature pack
- Enables a feature for a trial period

A feature pack is a predefined set of several premium features, such as Storage Partitioning and Synchronous Mirroring. These premium features are combined for the convenience of the users. When users install a feature pack, all of the premium features in the feature pack are installed at one time.

Each premium feature is managed by a license key that is generated for a specific feature or feature pack and a specific storage array. The license key is delivered as a file that you run to apply the license for the feature.

To determine which features are loaded on the storage array run the `show storageArray features` command. The `show storageArray features` command lists all of the premium features installed on the storage array, which premium features can be evaluated for a trial period, which premium features are enabled, and which premium features are disabled.

Syntax to Enable a Feature Key

```
enable storageArray feature file="filename"
```

The `file` parameter identifies the file path and the file name of a valid feature key file. Enclose the file path and the file name in double quotation marks (" "). For example:

```
file="C:\Program Files\CLI\dnld\ftrkey.key"
```

Valid file names for feature key files end with a `.key` extension.

You will need a feature key file for each premium feature that you want to enable.

Syntax to Enable a Feature Pack

```
enable storageArray featurePack file="filename"
```

The `file` parameter identifies the file path and the file name of a valid feature pack file. Enclose the file path and the file name in double quotation marks (" "). For example:

```
file="C:\Program Files\CLI\dnld\ftrpk.key"
```

Valid file names for feature key files end with a `.key` extension.

Syntax to Enable a Feature for a Trial Period

```
enable storageArray feature=featureAttributeList
```

To evaluate a feature for a trial period, you can enter one or more of the following attribute values for the `featureAttributeList`. If you enter more than one attribute value, separate the values with a white space.

- `asyncMirror`
- `syncMirror`
- `snapshot`
- `ssdCache`
- `volumeCopy`
- `thinProvisioning`

ATTENTION Before you enable the High Performance Tier premium feature, stop all host I/O operations to the storage array. When you enable the High Performance Tier premium feature, both controllers in the storage array will immediately reboot.

- `highPerformanceTier`
- `SSDSupport`

Notes

A premium feature is an additional application to enhance the capabilities of a storage array. Following is a list of the available premium feature attributes:

- `volumeCopy`
- `snapshot`
- `asyncMirror`
- `syncMirror`
- `mixedDriveTypes`
- `goldKey`
- `driveSecurity`

- enterpriseSecurityKeyMgr
- thinProvisioning
- storagePartition
- driveSlotLimit
- ssdCache
- dataAssurance
- SSDSupport

ATTENTION Before you enable the High Performance Tier premium feature, stop all host I/O operations to the storage array. When you enable the High Performance Tier premium feature, both controllers in the storage array will immediately reboot.

- highPerformanceTier
- raid6

NOTE The following attribute is for enabling the snapshot (legacy) feature only.

- snapshot[2|4|8|16]

The form for this attribute value is a combination of the alphabetical term merged with the numerical value, as shown by this example: snapshot16

NOTE The remoteMirror attribute is valid only for firmware versions before 7.84. Starting with firmware 7.84 the remoteMirror attribute is replaced by syncMirror.

- remoteMirror[8|16|32|64|128]

The form for this attribute value is a combination of the alphabetical term merged with the numerical value, as shown by this example: remoteMirror128

Minimum Firmware Level

6.10

6.50 adds these attributes:

- goldKey
- mixedDriveTypes

7.50 adds the highPerformanceTier attribute.

7.70 adds the syncMirror8 attribute. Firmware version 7.70 supports a maximum of eight remote mirrored pairs.

7.83 adds these attributes:

- raid6

- snapImage
- thinProvisioning

7.84 adds these attributes:

- asyncMirror
- ssdCache

Enable Volume Group Security

This command converts a non-secure volume group to a secure volume group.

Syntax

```
enable volumeGroup [volumeGroupName] security
```

Parameter

Parameter	Description
volumeGroup	The alphanumeric identifier (including - and _) of the volume group that you want to place in the Security Enabled state. Enclose the volume group identifier in square brackets ([]).

Notes

These conditions must be met to successfully run this command.

- All drives in the volume group must be full disk encryption drives.
- The Drive Security premium feature must be enabled.
- The storage array security key has to be set.
- The volume group is Optimal, and it does not have snapshot (legacy) volumes or repository volumes.

The controller firmware creates a lock that restricts access to the FDE drives. FDE drives have a state called Security Capable. When you create a security key, the state is set to Security Enabled, which restricts access to all FDE drives that exist within the storage array.

Minimum Firmware Level

7.40

Establish Asynchronous Mirrored Pair

Use this command to complete an asynchronous mirrored pair on the remote storage array by adding a secondary volume to an existing asynchronous mirror group.

Before you run this command, the asynchronous mirror group must exist and the primary volume must exist in the asynchronous mirror group. After this command successfully completes, asynchronous mirroring starts between the primary volume and the secondary volume.

The two volumes comprising an asynchronous mirrored pair function as a single entity. Establishing an asynchronous mirrored pair allows you to perform actions on the entire mirrored pair versus the two individual volumes.

Syntax

```
establish asyncMirror volume="secondaryVolumeName"  
asyncMirrorGroup="asyncMirrorGroupName"  
primaryVolume="primaryVolumeName"
```

Parameters

Parameter	Description
volume	The name of an existing volume on the remote storage array that you want to use for the secondary volume. Enclose the volume name in double quotation marks (" ").
asyncMirrorGroup	The name of an existing asynchronous mirror group that you want to use to contain the asynchronous mirrored pair. Enclose the asynchronous mirror group name in double quotation marks (" ").
primaryVolume	The name of an existing volume on the local storage array that you want to use for the primary volume. Enclose the volume name in double quotation marks (" ").

Notes

An asynchronous mirrored pair is comprised of two volumes, a primary volume and a secondary volume, that contain identical copies of the same data. The mirrored pair is a part of an asynchronous mirror group, which allows the mirrored pair to synchronize at the same time as any other mirrored pairs within the asynchronous mirror group .

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

When you choose the primary volume and the secondary volume, the secondary volume must be of equal or greater size than the primary volume. The RAID level of the secondary volume does not have to be the same as the primary volume.

Minimum Firmware Level

7.84

Export Storage Array Security Key

This command saves a full disk encryption (FDE) security key to a file. You can transfer the file from one storage array to another storage array. The file enables you to move FDE drives between storage arrays.

Syntax

```
export storageArray securityKey  
passPhrase="passPhraseString"  
file="fileName"
```

Parameters

Parameter	Description
passPhrase	A character string that encrypts the security key so that you can store the security key in an external file.
file	The file path and the file name to which you want to save the security key. For example: <code>file="C:\Program Files\CLI\sup\seckey.slk"</code> IMPORTANT – You must add a file extension of <code>.slk</code> to the end of the file name.

Notes

The storage array to which you will be moving drives must have drives with a capacity that is equal to or greater than the drives that you are importing.

The controller firmware creates a lock that restricts access to the full disk encryption (FDE) drives. FDE drives have a state called Security Capable. When you create a security key, the state is set to Security Enabled, which restricts access to all FDE drives that exist within the storage array.

Your pass phrase must meet these criteria:

- The pass phrase must be between eight and 32 characters long.
- The pass phrase must contain at least one uppercase letter.
- The pass phrase must contain at least one lowercase letter.
- The pass phrase must contain at least one number.
- The pass phrase must contain at least one non-alphanumeric character, for example, `< > @ +`.

NOTE If your pass phrase does not meet these criteria, you will receive an error message and will be asked to retry the command.

Minimum Firmware Level

7.40

Import Storage Array Security Key

This command unlocks one or more full disk encryption (FDE) drives that you have imported from one storage array to another storage array. Only the FDE drives with the matching security key from the imported storage array are unlocked. After they are unlocked, the security key for the new storage array is applied.

Syntax

```
import storageArray securityKey file="fileName"  
passPhrase="passPhraseString"
```

Parameters

Parameter	Description
file	The file path and the file name that has the original security key of the imported FDE drives. For example: file="C:\Program Files\CLI\sup\seckey.slk" IMPORTANT – The file that has the security key must have a file extension of .slk.
passPhrase	The character string that provides authentication for the security key. The pass phrase is 8 to 32 characters in length. You must use at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character in the pass phrase. A space is not permitted.

Notes

The controller firmware creates a lock that restricts access to the FDE drives. FDE drives have a state called Security Capable. When you create a security key, the state is set to Security Enabled, which restricts access to all FDE drives that exist within the storage array.

Your pass phrase must meet these criteria:

- The pass phrase must be between eight and 32 characters long.
- The pass phrase must contain at least one uppercase letter.
- The pass phrase must contain at least one lowercase letter.
- The pass phrase must contain at least one number.
- The pass phrase must contain at least one non-alphanumeric character, for example, < > @ +.

NOTE If your pass phrase does not meet these criteria, you will receive an error message and will be asked to retry the command.

Minimum Firmware Level

7.40

Increase Capacity of Volume in Disk Pool

This command increases the capacity of either a standard volume or a repository volume in a disk pool. In this command, a standard volume is also called a thick volume.

NOTE You cannot use this command to increase the capacity of a thin volume.

Syntax

```
start increaseVolumeCapacity volume="volumeName"  
incrementalCapacity=volumeCapacity  
[addDrives=(trayID1,drawerID1,slotID1 ...  
trayIDn,drawerIDn,slotIDn)]
```

Parameters

Parameter	Description
volume	The name of the disk pool volume for which you want to increase capacity. Enclose the volume name in square brackets ([]). If the volume name has special characters, you also must enclose the volume name in double quotation marks (" ") inside square brackets. If the volume name consists only of numbers, such as "1002," you also must enclose the volume name in double quotation marks (" ") inside square brackets.
incrementalCapacity	The setting to increase the storage size (capacity) for the volume. Size is defined in units of bytes, KB, MB, GB, or TB. The default value is bytes.
addDrives	The setting to add new drives to the volume. For high-capacity drive trays, specify the tray ID value, the drawer ID value, and the slot ID value for the drive. For low-capacity drive trays, specify the tray ID value and the slot ID value for the drive. Tray ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the tray ID value, drawer ID value, and the slot ID value in parentheses.

Notes

In some cases a drive parameter might appear as valid input for the command syntax. However, you cannot use the drive parameter with this command.

Setting the incrementalCapacity parameter, starts a long-running operation that you cannot stop. Long-running operations are performed in the background and do not prevent you from running other commands. To show the progress of long-running operations, use the show volume actionProgress command.

The addDrives parameter supports both high-capacity drive trays and low-capacity drive trays. A high-capacity drive tray has drawers that hold the drives. The drawers slide out of the drive tray to provide access to the drives. A low-capacity drive tray does not have drawers. For a high-capacity drive tray, you must specify the identifier (ID) of the drive tray, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive tray, you need only specify the ID of the drive tray and the ID of the slot in which a drive resides. For a low-capacity drive tray, an

alternative method for identifying a location for a drive is to specify the ID of the drive tray, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

Minimum Firmware Level

7.83

Initialize Thin Volume

This command initializes or re-initializes a thin volume.

- Used without any of the optional parameters, this command deletes the data on the thin volume. The repository volume capacity is not affected.
- Used with any of the optional parameters, this command cause reinitialization and repository volume actions.

NOTE Initializing a thin volume starts a long-running operation that you cannot stop.

Syntax

```
start volume [volumeName] initialize
[existingRepositoryLabel=existingRepositoryName |
diskPool="diskPoolName" capacity=capacityValue |
retainRepositoryMembers=[TRUE|FALSE]]
```

Parameter

Parameter	Description
volume	The name of the volume for which you are starting the formatting. Enclose the volume name in square brackets ([]). If the volume name has special characters, you also must enclose the volume name in double quotation marks (" ") inside square brackets. If the volume name consists only of numbers, such as "1002," you also must enclose the volume name in double quotation marks (" ") inside square brackets.
existingRepositoryLabel	This parameter replaces the repository volume by the candidate volume specified. <ul style="list-style-type: none"> ■ The value specified is an existing repository volume user label. The volume specified must be an unused repository volume with the name in the proper form. ■ If the newly specified repository volume is on a different disk pool, the thin volume will change ownership to that pool. ■ The old repository volume will be deleted by default.

Parameter	Description
diskPool	This parameter identifies the disk pool in which you want to create a new repository volume with the specified capacity. You must use this parameter with the capacity parameter to create a new repository volume with the specified capacity.
capacity	The size that you want to set for the repository volume that you are creating. Size is defined in units of bytes, KB, MB, GB, or TB. The minimum physical capacity is 4 GB. The maximum physical capacity is 64 TB. You must use this parameter with the diskPool parameter to create a new repository volume with the specified capacity.
retainRepositoryMembers	If this parameter is set to TRUE, the old repository is retained. By default, the old repository is deleted. This parameter is ignored if the existing repository is reused.

Notes

If you do not specify a volume with the volume parameter, this command resets the metadata in the repository and, in effect, causes the thin volume to appear empty to the host. If you specify a volume with the volume parameter, that volume is either replaced by a newly created volume or by an existing volume if you specify one with the existingRepositoryLabel parameter. If you specify an existing volume with the existingRepositoryLabel parameter that is in a different disk pool, the thin volume will change ownership to the new disk pool.

The volume parameters of the thin volume, such as virtual capacity, quota and warning threshold, retain their previous values after you reinitialize the thin volume.

The following table lists the capacity limits for a thin volume.

Type of Capacity	Size
Minimum virtual capacity	32 MB
Maximum virtual capacity	63 TB
Minimum physical capacity	4 GB
Maximum physical capacity	64 TB

Thin volumes support all of the operations that standard volumes do with the following exceptions:

- You cannot change the segment size of a thin volume.
- You cannot enable the pre-read redundancy check for a thin volume.

- You cannot use a thin volume as the target volume in a volume copy.
- You cannot use a thin volume in a snapshot (legacy) operation.
- You cannot use a thin volume in a Synchronous Mirroring operation.

If you want to change a thin volume to a standard volume, use the volume copy operation to create a copy of the thin volume. The target of a volume copy is always a standard volume.

Minimum Firmware Level

7.83

Load Storage Array DBM Database

This command restores a Database Management (DBM) database image by retrieving the image from a file or from cache. This command restores a storage array DBM database to the exact configuration that existed when the database image was captured using the `save storageArray dbmDatabase` command. The data in a file can be just RAID configuration data or all data, including data for RAID configuration, volume groups and disk pools. The data in a cache location always includes all data.

Before using this command with the `file` option, you must first obtain a validator string (a security code) from your Technical Support Representative. To obtain a validator, use the `save storageArray dbmValidator` command to generate an XML file that contains validator information. Your Technical Support Representative uses the XML file to generate the validator string required for this command.

Syntax

```
load storageArray dbmDatabase
((file="filename" validator="validatorValue" |
sourceLocation=onboard)
[controller [(a|b)] |
contentType=(partial | all)]
```

Parameters

Parameter	Description
file	The file path and the file name of the DBM database you want to upload. Enclose the file name in double quotation marks (" "). For example: file="C:\Array Backups\DBMbackup_03302010.dbm" This command does not automatically append a file extension to the saved file. You must specify a file extension when entering the file name.

Parameter	Description
validator	The alphanumeric security code required to restore a storage array to an existing configuration. Use the <code>save storageArray dbmValidator</code> command to generate the required validation information XML file. After the validation information XML file is available, contact your Technical Support Representative to obtain the Validator.
sourceLocation	This parameter specifies the location from which to retrieve backup database information. The parameter must be included for consistency, but the only allowed value is <code>onboard</code> .
controller	This parameter specifies the controller from which data will be exclusively retrieved, if <code>sourceLocation</code> is set to <code>onboard</code> . Enclose the controller identifier in squarebrackets ([]). If the <code>controller</code> parameter is not specified, data might be retrieved from either controller. Valid controller identifiers are <code>a</code> or <code>b</code> , where <code>a</code> is the controller in slot A, and <code>b</code> is the controller in slot B.
contentType	This parameter specifies the content type of the data that will be retrieved. <ul style="list-style-type: none"> ■ If the parameter is set to <code>partial</code>, a reduced set of records for the RAID configuration data is restored. This option helps reduce the possibility of record corruption by reducing the number and type of records restored. ■ If the parameter is set to <code>all</code>, all of the data including disk pool configuration data is retrieved. <p>The default is <code>all</code>.</p>

Notes

Depending on the size of the database image, restoring the database might take up as much as 30 minutes. The host software will not show the controllers in an Optimal state until after all actions for loading the database image are completed on the controllers.

Minimum Firmware Level

7.75

7.83 adds these parameters:

- `sourceLocation`
- `controller`
- `contentType`

Locate SSD Cache

The `start locate` command identifies the Solid State Disks (SSDs) that are being used in the SSD cache by turning on the indicator lights for the SSDs. The `stop locate` command turns off the indicator lights on the SSDs.

Syntax Applicable to Starting a Locate Operation

```
start ssdCache [ssdCacheName] locate
```

Syntax Applicable to Stopping a Locate Operation

```
stop ssdCache locate
```

Parameters

Parameter	Description
ssdCache	The alphanumeric identifier (including - and _) of the SSD cache that you want to locate. Enclose the identifier in square brackets ([]). If the SSD cache name contains special characters or consists only of numbers, you also must enclose the identifier in double quotation marks (" ") inside square brackets.

Minimum Firmware Level

7.84

Recopy Volume Copy

NOTE With firmware version 7.83 the `copyType=(online | offline)` parameter is no longer used.

This command reinitiates a volume copy operation using an existing volume copy pair. This command is valid for both snapshot (legacy) volume copy pairs and snapshot image volume copy pairs.

This command works with volume copy pairs that you created with a snapshot (legacy) volume or with a snapshot image volume.

ATTENTION Starting a volume copy operation overwrites all existing data on the target volume, makes the target volume read-only to hosts, and fails all snapshot (legacy) volumes or snapshot image volumes associated with the target volume, if any exist. If you have used the target volume as a copy before, be sure you no longer need the data or have it backed up.

Syntax

```
recopy volumeCopy target [targetName]
[source [sourceName]]
[copyPriority=(highest | high | medium | low | lowest)
targetReadOnlyEnabled=(TRUE | FALSE)]
```

Parameters

Parameter	Description
target	The name of the target volume for which you want to reinitiate a volume copy operation. Enclose the target volume name in square brackets ([]). If the target volume name has special characters, you also must enclose the target volume name in double quotation marks (" ").
source	The name of the source volume for which you want to reinitiate a volume copy operation. Enclose the source volume name in square brackets ([]). If the source volume name has special characters, you also must enclose the source volume name in double quotation marks (" ").
copyPriority	The priority that the volume copy has relative to host I/O activity. Valid values are highest, high, medium, low, or lowest.
targetReadOnlyEnabled	The setting so that you can write to the target volume or only read from the target volume. To write to the target volume, set this parameter to FALSE. To prevent writing to the target volume, set this parameter to TRUE.

Notes

Copy priority defines the amount of system resources that are used to copy the data between the source volume and the target volume of a volume copy pair. If you select the highest priority level, the volume copy uses the most system resources to perform the volume copy, which decreases performance for host data transfers.

Minimum Firmware Level

6.10

7.77 adds recopying a volume copy with snapshot (legacy).

7.83 removes the `copyType=(online | offline)` parameter.

Recover RAID Volume

This command creates a RAID volume with the given properties without initializing any of the user data areas on the drives. Parameter values are derived from the Recovery Profile data file (`recoveryProfile.csv`) for the storage array. You can create the recover volume in an existing volume group or create a new volume group by using this command.

NOTE You can run this command only from a command line. You cannot run this command from the GUI script editor. You cannot use the storage management GUI to recover a volume.

Syntax

```

recover volume (drive=(trayID,drawerID,slotID) |
drives=(trayID1,drawerID1,slotID1
... trayIDn,drawerIDn,slotIDn) |
volumeGroup=volumeGroupName)
[newVolumeGroup=volumeGroupName]
userLabel=("volumeName"
volumeWWN="volumeWWN")
capacity=volumeCapacity
offset=offsetValue
raidLevel=(0 | 1 | 3 | 5 | 6)
segmentSize=segmentSizeValue
dssPreallocate=(TRUE | FALSE)
SSID=subsystemVolumeID
[owner=(a | b)
cacheReadPrefetch=(TRUE | FALSE)
dataAssurance=(none | enabled)]

```

Parameters

Parameter	Description
drive or drives	The drives that you want to assign to the volume group that will contain the volume that you want to recover. For high-capacity drive trays, specify the tray ID value, the drawer ID value, and the slot ID value for each drive that you assign to the volume. For low-capacity drive trays, specify the tray ID value and the slot ID value for each drive that you assign to the volume. Tray ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the tray ID values, the drawer ID values, and the slot ID values in square brackets ([]).
volumeGroup	The name of an existing volume group in which you want to create the volume. (To determine the names of the volume groups in your storage array, run the show storageArray profile command.)
newVolumeGroup	The name that you want to give a new volume group. Enclose the new volume group name in double quotation marks (" ").
userLabel	The name of the volume that you want to recover. Enclose the volume name in double quotation marks (" ").
volumeWWN	The world wide name of the volume that you want to recover. The name is a 16 byte identifier, for example, 60080E500017B4320000000049887D77. Enclose the identifier in double quotation marks (" ").

Parameter	Description
capacity	The size of the volume that you are adding to the storage array. Size is defined in units of bytes, KB, MB, GB, or TB.
offset	The number of blocks from the start of the volume group to the start of the referenced volume.
raidLevel	The RAID level of the volume group that contains the drives. Valid values are 0, 1, 3, 5, or 6.
segmentSize	The amount of data (in KB) that the controller writes on a single drive in a volume group before writing data on the next drive. Valid values are 8, 16, 32, 64, 128, 256, or 512.
dssPreallocate	The setting to turn on or turn off allocating volume storage capacity for future segment size changes. To turn on allocation, set this parameter to TRUE. To turn off allocation, set this parameter to FALSE.
SSID	The storage array subsystem identifier of a volume. Use the <code>show volume</code> command to determine the storage array subsystem identifier.
owner	The controller that owns the volume. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. If you do not specify an owner, the controller firmware determines the owner.
cacheReadPrefetch	The setting to turn on or turn off cache read prefetch. To turn off cache read prefetch, set this parameter to FALSE. To turn on cache read prefetch, set this parameter to TRUE.
dataAssurance	The setting to specify that a volume group, and the volumes within the volume group, has data assurance protection to make sure that the data maintains its integrity. When you use this parameter, only protected drives can be used for the volume group. These settings are valid: <ul style="list-style-type: none"> ■ none – The volume group does not have data assurance protection. ■ enabled – The volume group has data assurance protection. The volume group supports protected information and is formatted with protection information enabled.

Notes

The storage management software collects recovery profiles of the monitored storage arrays and saves the profiles on a storage management station.

The `drive` parameter supports both high-capacity drive trays and low-capacity drive trays. A high-capacity drive tray has drawers that hold the drives. The drawers slide out of the drive tray to provide access to the drives. A low-capacity drive tray does not have drawers. For a high-capacity drive tray, you must specify the identifier (ID) of the drive tray, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive tray, you need only specify the ID of the drive tray and the ID of the slot in which a drive resides. For a low-capacity drive tray, an alternative method for identifying a location for a drive is to specify the ID of the drive tray, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

If you attempt to recover a volume using the `drive` parameter or the `drives` parameter and the drives are in an unassigned state, the controller automatically creates a new volume group. Use the `newVolumeGroup` parameter to specify a name for the new volume group.

You can use any combination of alphanumeric characters, underscore (`_`), hyphen (`-`), and pound (`#`) for the names. Names can have a maximum of 30 characters.

The `owner` parameter defines which controller owns the volume. The preferred controller ownership of a volume is the controller that currently owns the volume group.

Preallocating Storage Capacity

The `dssPreallocate` parameter enables you to assign capacity in a volume for storing information that is used to rebuild a volume. When you set the `dssPreallocate` parameter to `TRUE`, the storage space allocation logic in the controller firmware preallocates the space in a volume for future segment size changes. The preallocated space is the maximum allowable segment size. The `dssPreallocate` parameter is necessary for properly recovering volume configurations that are not retrievable from the controller data base. To turn off the preallocation capability, set `dssPreallocate` to `FALSE`.

Segment Size

The size of a segment determines how many data blocks that the controller writes on a single drive in a volume before writing data on the next drive. Each data block stores 512 bytes of data. A data block is the smallest unit of storage. The size of a segment determines how many data blocks that it contains. For example, an 8-KB segment holds 16 data blocks. A 64-KB segment holds 128 data blocks.

When you enter a value for the segment size, the value is checked against the supported values that are provided by the controller at run time. If the value that you entered is not valid, the controller returns a list of valid values. Using a single drive for a single request leaves other drives available to simultaneously service other requests.

If the volume is in an environment where a single user is transferring large units of data (such as multimedia), performance is maximized when a single data transfer request is serviced with a single data stripe. (A data stripe is the segment size that is

multiplied by the number of drives in the volume group that are used for data transfers.) In this case, multiple drives are used for the same request, but each drive is accessed only once.

For optimal performance in a multiuser database or file system storage environment, set your segment size to minimize the number of drives that are required to satisfy a data transfer request.

Cache Read Prefetch

Cache read prefetch lets the controller copy additional data blocks into cache while the controller reads and copies data blocks that are requested by the host from disk into cache. This action increases the chance that a future request for data can be fulfilled from cache. Cache read prefetch is important for multimedia applications that use sequential data transfers. The configuration settings for the storage array that you use determine the number of additional data blocks that the controller reads into cache. Valid values for the `cacheReadPrefetch` parameter are `TRUE` or `FALSE`.

Minimum Firmware Level

5.43

7.10 adds RAID 6 Level capability and the `newVolumeGroup` parameter.

7.60 adds the `drawerID` user input.

7.75 adds the `dataAssurance` parameter.

Re-create External Security Key

This command regenerates a storage array security key for use with the external security key management feature.

Syntax

```
recreate storageArray securityKey
passPhrase="passPhraseString"
file="fileName"
```

Parameters

Parameter	Description
<code>passPhrase</code>	A character string that encrypts the security key so that you can store the security key in an external file.
<code>file</code>	The file path and the file name that has the security key. For example: <code>file="C:\Program Files\CLI\sup\seckey.slk"</code> IMPORTANT – The file name must have an extension of <code>.slk</code> .

Notes

Your pass phrase must meet these criteria:

- The pass phrase must be between eight and 32 characters long.
- The pass phrase must contain at least one uppercase letter.
- The pass phrase must contain at least one lowercase letter.
- The pass phrase must contain at least one number.
- The pass phrase must contain at least one non-alphanumeric character, for example, < > @ +.

NOTE If your pass phrase does not meet these criteria, you will receive an error message.

Minimum Firmware Level

7.70

Re-create Snapshot (Legacy)

This command starts a fresh copy-on-write operation by using an existing snapshot (legacy) volume. You can re-create a single snapshot (legacy) volume or re-create multiple snapshot (legacy) volumes. If you choose to re-create multiple snapshot (legacy) volumes, you can re-create from two to the maximum number of snapshot (legacy) volumes that your storage array can support.

Syntax

```
recreate snapshot (volume [volumeName] |
volumes [volumeName1 ... volumeNameN])
[userLabel="snapshotVolumeName"
warningThresholdPercent=percentValue
repositoryFullPolicy (failBaseWrites | failSnapshot)]
```

Parameters

Parameter	Description
volume or volumes	The name of the specific volume for which you want to start a fresh copy-on-write operation. You can enter more than one volume name. Enclose the volume name in square brackets ([]). If the volume name has special characters, you must also enclose the volume name in double quotation marks (" ").
userLabel	The name of the snapshot (legacy) volume. Enclose the snapshot (legacy) volume name in double quotation marks (" "). If you enter more than one snapshot (legacy) volume name, this command fails.

Parameter	Description
<code>warningThresholdPercent</code>	The percentage of repository capacity at which you receive a warning that the snapshot (legacy) repository volume is nearing full. Use integer values. For example, a value of 70 means 70 percent. The default value is 50.
<code>repositoryFullPolicy</code>	The type of processing that you want to continue if the snapshot (legacy) repository volume is full. You can choose to fail writes to the base volume (<code>failBaseWrites</code>) or fail writes to the snapshot (legacy) volume (<code>failSnapshot</code>). The default value is <code>failSnapshot</code> .

Notes

You can use any combination of alphanumeric characters, underscore (`_`), hyphen (`-`), and pound (`#`) for the names. Names can have a maximum of 30 characters.

If you do not specify a value for the `warningThresholdPercent` parameter or the `repositoryFullPolicy` parameter, the previously set value is used.

Recreating a Single Snapshot (Legacy) Volume or Multiple Snapshot (Legacy) Volumes with Optional Parameters

- If you specify one or more of the optional parameters, the re-create operation processes each snapshot (legacy) volume separately.
- If you try to use the same user label for more than one volume, the command will fail.
- If you do not set the `warningThresholdPercent` parameter or the `repositoryFullPolicy` parameter, values that you previously set are used.

Recreating Multiple Snapshot (Legacy) Volumes without Optional Parameters

- If you list multiple snapshot (legacy) volumes to be re-created but do not specify any of the optional parameters, the re-create operation processes the snapshot (legacy) volumes as a "batch" process.
- Validation checks for the necessary snapshot (legacy)-restarted preconditions are performed before restarting any snapshot (legacy). If any of the listed snapshot (legacy) volumes fail the validation, the entire command fails and the snapshot (legacy) volumes are not re-created. If the validation is successful for all of the snapshot (legacy) volumes in the list, but one or more of the snapshots (legacy) in the list fails to restart, the entire command fails and none of the snapshots (legacy) are re-created.
- During snapshot (legacy) re-creation, all affected volumes (snapshots (legacy), base, and repository) are appropriately quiesced and I/O operations are resumed to all affected volumes after all snapshots (legacy) have been successfully re-created.

**Re-create
Synchronous
Mirroring
Repository Volume**

Minimum Firmware Level

5.00

NOTE With firmware version 7.80, the `recreate storageArray mirrorRepository` command is deprecated. This command is no longer supported in either the GUI or the CLI. If you attempt to run this command, an error message will be returned indicating that this functionality is no longer supported and that no changes will be made to the specified remote mirror repositories.

This command creates a new Synchronous Mirroring repository volume (also called a mirror repository volume) by using the parameters defined for a previous mirror repository volume. The underlying requirement is that you have previously created a mirror repository volume. When you use this command, you can define the mirror repository volume in one of three ways: user-defined drives, user-defined volume group, or user-defined number of drives for the mirror repository volume. If you choose to define a number of drives, the controller firmware chooses which drives to use for the mirror repository volume.

Syntax (User-Defined Drives)

```
recreate storageArray mirrorRepository
repositoryRAIDLevel=(1 | 3 | 5 | 6)
repositoryDrives=(trayID1,slotID1 ... trayIDn,slotIDn)
[trayLossProtect=(TRUE | FALSE)
dataAssurance=(none | enabled)]
```

Syntax (User-Defined Volume Group)

```
recreate storageArray mirrorRepository
repositoryVolumeGroup=volumeGroupName
[freeCapacityArea=freeCapacityIndexNumber]
```

Syntax (User-Defined Number of Drives)

```
recreate storageArray mirrorRepository
repositoryRAIDLevel=(1 | 3 | 5 | 6)
repositoryDriveCount=numberOfDrives
[driveType=(fibre | SATA | SAS)]
[trayLossProtect=(TRUE | FALSE)
dataAssurance=(none | enabled)]
```

Parameters

Parameter	Description
repositoryRAIDLevel	The RAID level for the mirror repository volume. Valid values are 1, 3, 5, or 6.

Parameter	Description
repositoryDrives	The drives for the mirror repository volume. Specify the tray ID and slot ID for each drive that you assign to the mirror repository volume. Tray ID values are 0 to 99. Slot ID values are 1 to 32. Enclose the tray ID values and the slot ID values in parentheses.
repositoryVolumeGroup	The name of the volume group where the mirror repository volume is located.
freeCapacityArea	<p>The index number of the free space in an existing volume group that you want to use to re-create the mirror repository volume. Free capacity is defined as the free capacity between existing volumes in a volume group. For example, a volume group might have the following areas: volume 1, free capacity, volume 2, free capacity, volume 3, free capacity. To use the free capacity following volume 2, you would specify:</p> <pre>freeCapacityArea=2</pre> <p>Run the <code>show volumeGroup</code> command to determine if a free capacity area exists.</p>
repositoryDriveCount	The number of unassigned drives that you want to use for the mirror repository volume.
driveType	<p>The type of drive that you want to use for the mirror repository volume. You cannot mix drive types.</p> <p>You must use this parameter when you have more than one type of drive in your storage array.</p> <p>Valid drive types are :</p> <ul style="list-style-type: none"> ■ fibre ■ SATA ■ SAS <p>If you do not specify a drive type, the command defaults to fibre.</p>
trayLossProtect	The setting to enforce tray loss protection when you create the mirror repository volume. To enforce tray loss protection, set this parameter to TRUE. The default value is FALSE.

Parameter	Description
dataAssurance	<p>The setting to specify that a volume group, and the volumes within the volume group, has data assurance protection to make sure that the data maintains its integrity. When you use this parameter, only protected drives can be used for the volume group. These settings are valid:</p> <ul style="list-style-type: none"> ■ none – The volume group does not have data assurance protection. ■ enabled – The volume group has data assurance protection. The volume group supports protected information and is formatted with protection information enabled.

Notes

If you enter a value for the storage space of the mirror repository volume that is too small, the controller firmware returns an error message, which states the amount of space that is needed for the mirror repository volume. The command does not try to change the mirror repository volume. You can re-enter the command by using the value from the error message for the storage space value of the mirror repository volume.

When you assign the drives, if you set the `trayLossProtect` parameter to `TRUE` and have selected more than one drive from any one tray, the storage array returns an error. If you set the `trayLossProtect` parameter to `FALSE`, the storage array performs operations, but the mirror repository volume that you create might not have tray loss protection.

When the controller firmware assigns the drives, if you set the `trayLossProtect` parameter to `TRUE`, the storage array returns an error if the controller firmware cannot provide drives that result in the new mirror repository volume having tray loss protection. If you set the `trayLossProtect` parameter to `FALSE`, the storage array performs the operation even if it means that the mirror repository volume might not have tray loss protection.

Minimum Firmware Level

6.10

7.10 adds RAID Level 6 capability

7.75 adds the `dataAssurance` parameter.

Remove Drives from SSD Cache

This command decreases the capacity of the SSD cache by removing Solid State Disks (SSDs).

Syntax

```
set ssdCache [ssdCacheName]  
removeDrives=(trayID1,drawerID1,slotID1 ...  
trayIDn,drawerIDn,slotIDn)
```

Parameter

Parameter	Description
<code>ssdCache</code>	The alphanumeric identifier (including - and _) of the SSD cache from which you want to remove SSDs. Enclose the identifier in square brackets ([]). If the SSD cache name contains special characters or consists only of numbers, you also must enclose the identifier in double quotation marks (" ") inside square brackets.
<code>removeDrives</code>	The drives that you want to remove from the SSD cache. For high-capacity drive trays, specify the tray ID value, the drawer ID value, and the slot ID value for each SSD that you want to remove. For low-capacity drive trays, specify the tray ID value and the slot ID value for each SSD that you want to remove. Tray ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the tray ID values, the drawer ID values, and the slot ID values in parentheses.

Notes

You cannot remove all of the SSDs from the SSD cache using this command; at least one SSD must remain in the SSD cache. If you want to completely remove the SSD cache, use the `delete ssdCache` command instead.

Minimum Firmware Level

7.84

Remove Incomplete Asynchronous Mirrored Pair from Asynchronous Mirror Group

This command removes an orphaned mirrored pair volume on the storage array. An orphaned mirrored pair volume exists when a member volume in an asynchronous mirror group has been removed on one side of the asynchronous mirror group (either the primary side or secondary side) but not on the other side.

Orphan mirrored pair volumes are detected when inter-controller communication is restored and the two sides of the mirror configuration reconcile mirror parameters.

Use this command when the mirror relationship is successfully removed on the local or remote storage array, but cannot be removed on the corresponding storage array because of a communication problem.

Syntax

```
remove asyncMirrorGroup ["asyncMirrorGroupName"]
incompleteMirror volume="volumeName"
```

Parameters

Parameter	Description
asyncMirrorGroup	The name of the asynchronous mirror group that contains the orphaned volume that you want to remove. Enclose the asynchronous mirror group name in double quotation marks (" "). If the asynchronous mirror group name has special characters, you also must enclose the asynchronous mirror group name in double quotation marks (" ") inside square brackets.
volume	The name of the orphaned volume that you want to remove from the asynchronous mirror group. Enclose the volume name in double quotation marks (" ").

Minimum Firmware Level

7.84

Remove Member Volume from Consistency Group

This command removes a member volume from a an existing snapshot consistency group. Optionally, you can delete the repository volume members from the consistency group.

Syntax

```
set consistencyGroup ["consistencyGroupName"]
removeCGMemberVolume="memberVolumeName"
[deleteRepositoryMembers=(TRUE | FALSE)]
```

Parameters

Parameter	Description
consistencyGroupName	The name of the consistency group from which you want to remove a member. Enclose the consistency group name in double quotation marks (" ") inside square brackets ([]).
removeCGMemberVolume	The name of the member volume that you want to remove. Enclose the consistency group name in double quotation marks (" ").
deleteRepositoryMembers	Determines whether to remove all of the repository members from the consistency group.

Remove Synchronous Mirroring

Minimum Firmware Level

7.83

This command removes the mirror relationship between the primary volume and the secondary volume in a remote-mirrored pair.

NOTE In previous versions of this command the feature identifier was `remoteMirror`. This feature identifier is no longer valid and is replaced by `syncMirror`.

Syntax

```
remove syncMirror (localVolume [volumeName] |
localVolumes [volumeName1 ... volumeNameN])
```

Parameter

Parameter	Description
localVolume or localVolumes	The name of the primary volume (the volume on the local storage array) that you want to remove. You can enter more than one volume name. Enclose the volume name in square brackets ([]). If the volume name has special characters, you also must enclose the volume name in double quotation marks (" ").

Minimum Firmware Level

6.10

Remove Volume Copy

NOTE With firmware version 7.83 the `copyType=(online | offline)` parameter is no longer used.

This command removes a volume copy pair. This command is valid for both snapshot (legacy) volume copy pairs and new snapshot volume copy pairs.

Syntax

```
remove volumeCopy target [targetName]
[source [sourceName]]
```

Parameters

Parameter	Description
target	The name of the target volume that you want to remove. Enclose the target volume name in square brackets ([]). If the target volume name has special characters, you also must enclose the target volume name in double quotation marks (" ").
source	The name of the source volume that you want to remove. Enclose the source volume name in square brackets ([]). If the source volume name has special characters, you also must enclose the source volume name in double quotation marks (" ").

Minimum Firmware Level

5.40

7.77 adds creating a volume copy with snapshot (legacy).

7.83 removes the `copyType=(online | offline)` parameter.

Remove Volume from Asynchronous Mirror Group

This command removes a member volume from an existing asynchronous mirror group. Optionally, you can delete the repository volume members from the asynchronous mirror group.

This command is valid only on the local storage array that contains the asynchronous mirror group whose member volume that you want to remove.

Syntax

```
remove volume ["volumeName"]
asyncMirrorGroup="asyncMirrorGroupName"
[deleteRepositoryMembers=(TRUE | FALSE)]
```

Parameters

Parameter	Description
volume	The name of the specific volume that you want to remove from the asynchronous mirror group. Enclose the volume name in double quotation marks (" ") inside square brackets ([]).
asyncMirrorGroup	The name of the asynchronous mirror group that contains the member volume that you want to remove. Enclose the asynchronous mirror group name in double quotation marks (" ").
deleteRepositoryMembers	Determines whether to remove all of the repository members from the asynchronous mirror group.

Minimum Firmware Level

7.84

Remove Volume LUN Mapping

This command removes the logical unit number (LUN) mapping from one or more volumes.

Syntax

```
remove (allVolumes | volume [volumeName] |  
volumes [volumeName1 ... volumeNameN] | accessVolume)  
lunMapping (host="hostName" |  
hostGroup=("hostGroupName" | defaultGroup))
```

Parameters

Parameter	Description
allVolumes	This parameter removes the LUN mapping from all of the volumes.
volume or volumes	The name of the specific volume for which you are retrieving information. Enclose the volume name in square brackets ([]). If the volume name has special characters, you also must enclose the volume name in double quotation marks (" ") inside square brackets. If the volume name consists only of numbers, such as "1002," you also must enclose the volume name in double quotation marks (" ") inside square brackets. You can enter more than one volume name. Enclose all of the volume names in one set of double quotation marks (" ") inside square brackets ([]). Separate each volume name with a white space.
accessVolume	This parameter removes the access volume.
host	The name of the host to which the volume is mapped. Enclose the host name in double quotation marks (" ").
hostGroup	The name of the host group that contains the host to which the volume is mapped. Enclose the host group name in double quotation marks (" "). The defaultGroup value is the host group that contains the host to which the volume is mapped.

Notes

The access volume is the volume in a SAN environment that is used for communication between the storage management software and the storage array controller. The access volume uses a LUN address and consumes 20 MB of storage space that is not available for application data storage. An access volume is required only for in-band managed storage arrays.

ATTENTION Removing an access volume can damage your configuration – The agent uses the access volumes to communicate with a storage array. If you remove an access volume mapping for a storage array from a host that has an agent running on it, the storage management software is no longer able to manage the storage array through the agent.

You must use the `host` parameter and the `hostGroup` parameter when you specify a non-access volume or an access volume. The Script Engine ignores the `host` parameter or the `hostGroup` parameter when you use the `allVolumes` parameter or the `volumes` parameter.

Minimum Firmware Level

6.10

Rename Snapshot Volume

This command renames an existing snapshot volume.

Syntax

```
set snapVolume [ "snapVolumeName" ]  
userLabel="snapImageVolumeName "
```

Parameters

Parameter	Description
<code>snapVolume</code>	The name of the snapshot volume that you want to rename. Enclose the snapshot volume name in double quotation marks (" ") inside of square brackets ([]).
<code>userLabel</code>	A new name that you want to give to the snapshot volume. Enclose the new snapshot volume name in double quotation marks (" ").

Notes

You can use any combination of alphanumeric characters, underscore (_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

Minimum Firmware Level

7.83

Rename SSD Cache

This command changes the name of the SSD cache.

Syntax

```
set ssdCache [ old_ssdCacheName ] userLabel="new_ssdCacheName "
```

Parameter

Parameter	Description
ssdCache	The alphanumeric identifier (including - and _) of the SSD cache that you want to rename. Enclose the identifier in square brackets ([]). If the SSD cache name contains special characters or consists only of numbers, you also must enclose the identifier in double quotation marks (" ") inside square brackets.
userLabel	The new name for the SSD cache. Enclose the name in double quotation marks (" "). You can use any combination of alphanumeric characters, underscore (_), hyphen (-), and pound (#) for the identifier. Identifiers can have a maximum of 30 characters.

Minimum Firmware Level

7.84

Repair Volume Parity

This command repairs the parity errors on a volume.

Syntax

```
repair volume [volumeName] parity
parityErrorFile="filename"
[verbose=(TRUE | FALSE)]
```

Parameters

Parameter	Description
volume	The name of the specific volume for which you are retrieving information. Enclose the volume name in square brackets ([]). If the volume name has special characters, you also must enclose the volume name in double quotation marks (" ") inside square brackets. If the volume name consists only of numbers, such as "1002," you also must enclose the volume name in double quotation marks (" ") inside square brackets.
parityErrorFile	The file path and the file name that contains the parity error information that you use to repair the errors. Enclose the file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\sup\parfile.txt"
verbose	The setting to capture progress details, such as percent complete, and to show the information as the volume parity is being repaired. To capture progress details, set this parameter to TRUE. To prevent capturing progress details, set this parameter to FALSE.

Minimum Firmware Level

6.10

Replace Drive

This command redefines the composition of a volume group. You can use this command to replace a drive with either an unassigned drive or a fully integrated hot spare.

Syntax

```
replace drive([trayID,drawerID,slotID] | <"wwID">)  
replacementDrive=trayID,drawerID,slotID
```

Parameters

Parameter	Description
drive	The location of the drive that you want to reconstruct. For high-capacity drive trays, specify the tray ID value, the drawer ID value, and the slot ID value of the drive that you want to revive. For low-capacity drive trays, specify the tray ID value and the slot ID value of the drive that you want to revive. Tray ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the tray ID value, the drawer ID value, and the slot ID value in square brackets ([]).
replacementDrive	The location of the drive that you want to use for a replacement. For high-capacity drive trays, specify the tray ID value, the drawer ID value, and the slot ID value for the drive. For low-capacity drive trays, specify the tray ID value and the slot ID value for the drive. Tray ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32.

Notes

The `drive` parameter supports both high-capacity drive trays and low-capacity drive trays. A high-capacity drive tray has drawers that hold the drives. The drawers slide out of the drive tray to provide access to the drives. A low-capacity drive tray does not have drawers. For a high-capacity drive tray, you must specify the identifier (ID) of the drive tray, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive tray, you need only specify the ID of the drive tray and the ID of the slot in which a drive resides. For a low-capacity drive tray, an alternative method for identifying a location for a drive is to specify the ID of the drive tray, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

Minimum Firmware Level

7.10

7.60 adds the `drawerID` user input.

Reset Asynchronous Mirror Group Statistics

This command resets the synchronization statistics for one or more member volumes in an asynchronous mirror group to a relative 0.

Syntax

```
reset storageArray arvmStats asyncMirrorGroup
["asyncMirrorGroupName"]
volume="volumeName" sampleType=(all | mostRecent |
longestSyncTime | errors)
```

Parameters

Parameter	Description
<code>asyncMirrorGroup</code>	The name of the asynchronous mirror group for which you are resetting the synchronization statistics. Enclose the asynchronous mirror group name in square brackets ([]). If the asynchronous mirror group name has special characters, you also must enclose the asynchronous mirror group name in double quotation marks (" ") inside square brackets. If the asynchronous mirror group name consists only of numbers, such as "1002," you also must enclose the asynchronous mirror group name in double quotation marks (" ") inside square brackets.
<code>volume</code>	This parameter is optional. The name of the specific member volume in the asynchronous mirror group for which you are resetting the synchronization statistics. If no volume is specified, the statistics for every member volume in the asynchronous mirror group are reset. Enclose the volume name in in double quotation marks (" ").
<code>sampleType</code>	This parameter is optional. The default value for <code>sampleType</code> is <code>all</code> . <ul style="list-style-type: none"> ■ <code>all</code>— Data for all three sample types are reset. ■ <code>mostRecent</code> — Statistics are reset for the most recent 50 resynchronization samples. ■ <code>longestSyncTime</code> — Statistics are reset for the most recent 20 longest resynchronization samples. ■ <code>errors</code>— Statistics are reset for the most recent 20 failed resynchronization samples.

Notes

Statistics are reset for mirrored volumes in the Primary role. The statistics that are reset include the following data:

- Synchronization start time
- Synchronization type (manual or periodic)
- Synchronization duration

- Number of bytes sent
- Maximum and minimum write time (for a single write)
- Maximum and minimum synchronization data rate
- Total write time
- Repository utilization (%)
- Recovery point age

Minimum Firmware Level

7.84

Reset Controller

This command resets a controller, and it is disruptive to I/O operations.

ATTENTION When you reset a controller, the controller is removed from the data path and is not available for I/O operations until the reset operation is complete. If a host is using volumes that are owned by the controller being reset, the I/O directed to the controller is rejected. Before resetting the controller, either make sure that the volumes that are owned by the controller are not in use or make sure that a multi-path driver is installed on all of the hosts that use these volumes.

Syntax

```
reset controller [(a | b)]
```

Parameter

Parameter	Description
controller	The controller that you want to reset. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller owner, the controller firmware returns a syntax error.

Notes

The controller that receives the reset controller command resets the controller specified. For example, if the reset controller command is sent to controller A to request a reset of controller A, then controller A reboots itself by doing a soft reboot. If the reset controller command is sent to controller A to request a reset of controller B, then controller A holds controller B in reset and then releases controller B from reset, which is a hard reboot. A soft reboot in some products only resets the IOC chip. A hard reboot resets both the IOC and the expander chips in the controller.

Minimum Firmware Level

5.20

Reset iSCSI IP Address

This command resets the IP address for the remote storage array to re-establish connection with the local storage array. You can use this command to notify the local storage array that the iSCSI IP address of the remote storage array have changed and need to be updated.

When establishing an asynchronous mirroring relationship with an iSCSI connection, both storage arrays store a record of the IP address of the remote storage array of the asynchronous mirroring configuration. If the IP address of an iSCSI port changes, the remote storage array that is attempting to use that port encounters a communication error.

The storage array with the changed IP address sends a message to each remote storage array associated with the asynchronous mirror groups that are configured to mirror over an iSCSI connection. Storage arrays that receive this message automatically update their remote-target IP address.

If the storage array with the changed IP address is unable to send its inter-controller message to a remote storage array, the system sends you an alert of the connectivity issue. Use the `reset` command to re-establish connection with the local storage array.

Syntax

```
reset (remoteStorageArrayName="storageArrayName" |
remoteStorageArrayWwid=<wwID>)
iscsiIpAddress
```

Parameters

Parameter	Description
<code>remoteStorageArrayName</code>	The name for the remote storage array for which you are resetting the iSCSI IP address. Enclose the storage array name in double quotation marks (" ").
<code>remoteStorageArrayNameWwid</code>	The World Wide Identifier (WWID) of the storage array for which you are resetting the iSCSI IP address. You can use the WWID instead of the storage array name to identify the storage array. Enclose the WWID in angle brackets (<>).

Minimum Firmware Level

7.84

Reset Storage Array Battery Install Date

This command resets the age of the batteries in a storage array to zero days. You can reset the age of the batteries for an entire storage array or the age of a battery in a specific controller or in a specific battery pack.

Syntax

```
reset storageArray batteryInstallDate  
(controller=[(a | b)] | batteryPack [left | right])
```

Parameters

Parameter	Description
controller	The controller that contains the battery for which you want to reset the age. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Use the controller parameter only for controllers with batteries.
batteryPack	The battery pack contains both a left battery and a right battery. Valid identifiers are left or right, where left is the battery that supports the controller in slot A, and right is the battery that supports the controller in slot B. Use the batteryPack parameter only for controller trays with battery packs.

Notes

A controller might have a battery associated with it, so the controller is identified as either a or b. With the release of the CE7900 controller tray, battery packs inside the interconnect-battery canister are identified as either left or right. If the command statement uses the wrong parameter, an error appears.

Minimum Firmware Level

6.10

7.15 adds the ability to reset the battery installation dates on the left battery or the right battery in the CE6998-series controllers or the CE7900-series controllers.

Reset Storage Array Diagnostic Data

This command resets the NVSRAM that contains the diagnostic data for the storage array. This command does not delete the diagnostic data. This command replaces the Needs Attention status with the Diagnostic Data Available status. The old diagnostic data is written over automatically when new data is captured. The memory that contains the diagnostic data is also cleared when the controllers reboot. Before you reset the diagnostic data, use the save storageArray diagnosticData command to save the diagnostic data to a file.

ATTENTION Run this command only with the assistance of your Technical Support Representative.

Syntax

```
reset storageArray diagnosticData
```

Reset Storage Array Infiniband Statistics Baseline

Parameters

None.

Minimum Firmware Level

6.16

This command resets the Infiniband statistics baseline to 0 for the storage array.

Syntax

```
reset storageArray ibStatsBaseline
```

Parameters

None.

Notes

This command does not actually reset the raw counts maintained in the hardware and firmware. Instead, the firmware creates a snapshot of the current counter values and uses these values to report differences in the counts when the statistics are retrieved. The new baseline time is applied to both controllers so that the controller counts are synchronized with each other. If one controller resets without the other controller resetting, the counters are no longer synchronized. The client becomes aware that the controllers are not synchronized because the timestamp data reported along with the statistics is not the same for both controllers.

Minimum Firmware Level

7.10

Reset Storage Array iSCSI Baseline

This command resets the iSCSI baseline to 0 for the storage array.

Syntax

```
reset storageArray iscsiStatsBaseline
```

Parameters

None.

Notes

This command resets the baseline to 0 for both controllers in the storage array. The purpose of resetting both of the controller baselines is to help make sure that the controller counts are synchronized between the controllers. If one controller resets but the second controller does not reset, the host is informed that the controllers are out of synchronization. The host is informed by the time stamps that are reported with the statistics.

Minimum Firmware Level

7.10

Reset Storage Array RLS Baseline

This command resets the read link status (RLS) baseline for all devices by setting all of the RLS counts to 0.

Syntax

```
reset storageArray RLSBaseline
```

Parameters

None.

Minimum Firmware Level

5.00

Reset Storage Array SAS PHY Baseline

This command resets the SAS physical layer (SAS PHY) baseline for all devices, and removes the list of errors from the `.csv` file. The `.csv` file is generated when you run the `save storageArray SASPHYCounts` command.

NOTE The previous release of the `reset storageArray SASPHYBaseline` command cleared error counts for all devices except the drives. The `reset storageArray SASPHYBaseline` command now resets the SAS PHY baseline for the drives as well as the other devices. All errors are deleted from the `.csv` file.

Syntax

```
reset storageArray SASPHYBaseline
```

Parameters

None.

Minimum Firmware Level

6.10

7.83 resets the SAS PHY baseline for drives.

Reset Storage Array SOC Baseline

This command resets the baseline for all switch-on-a-chip (SOC) devices that are accessed through the controllers. This command resets the baseline by setting all of the SOC counts to 0. This command is valid only for Fibre Channel devices in an arbitrated loop topology.

Syntax

```
reset storageArray SOCBaseline
```

Parameters

None.

Minimum Firmware Level

6.16

Reset Storage Array Volume Distribution

This command reassigns (moves) all of the volumes to their preferred controller.

Syntax

```
reset storageArray volumeDistribution
```

Parameters

None.

Notes

If you use this command on a host without a multi-path driver, you must stop I/O operations to the volumes until this command has completed to prevent application errors.

Under certain host operating system environments, you might be required to reconfigure the multi-path host driver. You might also need to make operating system modifications to recognize the new I/O path to the volumes.

Minimum Firmware Level

5.20

Resume Asynchronous Mirror Group

This command resumes data transfer between all mirrored pairs in an asynchronous mirror group. Data written to the primary volumes while the asynchronous mirror group was suspended is written to the secondary volumes immediately. Periodic synchronization resumes if an automatic synchronization interval has been set.

Syntax

```
resume asyncMirrorGroup ["asyncMirrorGroupName"]
```

Parameter

Parameter	Description
asyncMirrorGroup	The name of the asynchronous mirror group that you want to resume. Enclose the asynchronous mirror group name in square brackets ([]). If the asynchronous mirror group name has special characters, you also must enclose the asynchronous mirror group name in double quotation marks (" ") inside square brackets. If the asynchronous mirror group name consists only of numbers, such as "1002," you also must enclose the asynchronous mirror group name in double quotation marks (" ") inside square brackets.

Minimum Firmware Level

7.84

Resume Consistency Group Snapshot Volume

This command restarts a copy-on-write operation for creating a consistency group snapshot volume that you stopped using the `stop cgSnapVolume` command.

Syntax

```
resume cgSnapVolume ["snapVolumeName"]
cgSnapImage="snapImageName"
```

Parameter

Parameter	Description
cgSnapVolume	The name of the consistency group snapshot volume that you want to resume. Enclose the name of the consistency group snapshot volume in double quotation marks (" ") inside square brackets ([]).
cgSnapImage	<p>The name of the snapshot image in a consistency group that you are restarting. The name of a snapshot image is comprised of two parts:</p> <ul style="list-style-type: none"> ■ The name of the consistency group ■ An identifier for the snapshot image in the consistency group. <p>The identifier for the snapshot image can be one of these:</p> <ul style="list-style-type: none"> ■ An integer value that is the sequence number of the snapshot in the consistency group. ■ NEWEST - Use this option when you want to show the latest snapshot image created in the consistency group. ■ OLDEST - Use this option when you want to show the earliest snapshot image created in the consistency group. <p>Enclose the snapshot image name in double quotation marks (" ").</p>

Notes

The name of a snapshot image has two parts separated by a colon (:):

- The identifier of the consistency group
- The identifier of the snapshot image

For example, if you want to restart a copy-on-write operation for snapshot image 12345 which is in snapshot consistency group snapgroup1 in a consistency group snapshot volume that has the name snapVol1, you would use this command:

```
resume cgSnapVolume ["snapVol1"]
cgSnapImage=["snapgroup1:12345"]
```

Minimum Firmware Level

7.83

Resume Snapshot (Legacy) Rollback

This command resumes a rollback operation that has entered a paused state. A rollback operation can enter a paused state due to processing errors, which will trigger a Needs Attention condition for the storage array.

If the rollback operation cannot be resumed, the selected snapshot (legacy) volume reverts to a paused state, and the Needs Attention condition is displayed.

Syntax

```
resume rollback volume [snapshotVolumeName]
```

Parameter

Parameter	Description
volume	The name of the specific snapshot (legacy) volume for which you want to resume a rollback operation. Enclose the snapshot (legacy) volume name in square brackets ([]). If the snapshot (legacy) volume name has special characters, you must also enclose the snapshot (legacy) volume name in double quotation marks (" ").

Minimum Firmware Level

7.80

Resume Snapshot Image Rollback

This command resumes a rollback operation that has entered a paused state. A rollback operation can enter a paused state due to processing errors, which causes a Needs Attention condition for the storage array.

If the rollback operation cannot be resumed, the selected snapshot image reverts to a paused state, and the Needs Attention condition is displayed.

NOTE You cannot use this command for snapshot images involved in online volume copy.

Syntax

```
resume snapImage [snapImageName] rollback
```

Parameter

Parameter	Description
snapImage	<p>The name of the snapshot image for which you want to restart a rollback operation. The name of a snapshot image is comprised of two parts:</p> <ul style="list-style-type: none">■ The name of the snapshot group■ An identifier for the snapshot image in the snapshot group <p>The identifier for the snapshot image can be one of these:</p> <ul style="list-style-type: none">■ An integer value that is the sequence number of the snapshot in the snapshot group.■ NEWEST - Use this option when you want to show the latest snapshot image created in the snapshot group.■ OLDEST - Use this option when you want to show the earliest snapshot image created in the snapshot group. <p>Enclose the snapshot image name in double quotation marks (" ") inside square brackets ([]).</p>

Notes

The name of a snapshot image has two parts separated by a colon (:):

- The identifier of the snapshot group
- The identifier of the snapshot image

For example, if you want to restart a rollback operation for snapshot image 12345 in a snapshot group that has the name snapgroup1, you would use this command:

```
resume snapImage ["snapgroup1:12345"] rollback;
```

Minimum Firmware Level

7.83

Resume Snapshot Volume

This command resumes a snapshot volume operation that was stopped.

Syntax

```
resume snapVolume [snapVolumeName]  
snapImage="snapCGID:imageID"
```

Parameters

Parameter	Description
snapVolume	The name of the specific snapshot volume that you want to resume. Enclose the snapshot volume name in square brackets ([]). If the snapshot volume name has special characters, you also must enclose the snapshot volume name in double quotation marks (" ").
snapImage	The name of the snapshot image for which you want to resume snapshot volume operations. The name of a snapshot image is comprised of two parts: <ul style="list-style-type: none">■ The name of the snapshot group■ An identifier for the snapshot image in the snapshot group The identifier for the snapshot image can be one of these: <ul style="list-style-type: none">■ An integer value that is the sequence number of the snapshot in the snapshot group.■ NEWEST - Use this option when you want to show the latest snapshot image created in the snapshot group.■ OLDEST - Use this option when you want to show the earliest snapshot image created in the snapshot group. Enclose the snapshot image name in double quotation marks (" ") inside square brackets ([]).

Notes

The name of a snapshot image has two parts separated by a colon (:):

- The identifier of the snapshot group
- The identifier of the snapshot image

For example, if you want to resume snapshot volume operations for snapshot image 12345 in a snapshot group that has the name snapGroup1, you would use this command:

```
resume snapVolume [snapVoll] snapImage=["snapGroup1:12345"];
```

To resume snapshot volume operations for the most recent snapshot image in a snapshot group that has the name snapGroup1, you would use this command:

```
resume snapVolume [snapVoll]  
snapImage=["snapGroup1:newest"];
```

Minimum Firmware Level

7.83

Resume SSD Cache

This command restarts the caching for all of the volumes using the SSD cache that was temporarily stopped with the `suspend ssdCache` command.

Syntax

```
resume ssdCache [ssdCacheName]
```

Parameter

Parameter	Description
ssdCache	The alphanumeric identifier (including - and _) of the SSD cache that you want to resume. Enclose the identifier in square brackets ([]). If the SSD cache name contains special characters or consists only of numbers, you also must enclose the identifier in double quotation marks (" ") inside square brackets.

Minimum Firmware Level

7.84

Resume Synchronous Mirroring

This command resumes a suspended Synchronous Mirroring operation.

NOTE In previous versions of this command the feature identifier was `remoteMirror`. This feature identifier is no longer valid and is replaced by `syncMirror`.

Syntax

```
resume syncMirror (primary [volumeName] |  
primaries ["volumeName1" ... "volumeNameN"] )  
[writeConsistency=(TRUE | FALSE)]
```

Parameters

Parameter	Description
primary or primaries	The name of the primary volume for which you want to resume operation. Enclose the primary volume name in square brackets ([]). If the primary volume name has special characters, you also must enclose the primary volume name in double quotation marks (" "). You can enter more than one volume name. Enclose all of the volume names in one set of square brackets ([]). Enclose each volume name in double quotation marks (" "). Separate each volume name with a white space.

Parameter	Description
writeConsistency	The setting to identify the volumes in this command that are in a write-consistency group or are separate. For the volumes to be in the same write-consistency group, set this parameter to TRUE. For the volumes to be separate, set this parameter to FALSE.

Notes

If you set the `writeConsistency` parameter to `TRUE`, the volumes must be in a write-consistency group (or groups). This command resumes all write-consistency groups that contain the volumes. For example, if volumes A, B, and C are in a write-consistency group and they have remote counterparts A', B', and C', the `resume syncMirror volume ["A"] writeConsistency=TRUE` command resumes A-A', B-B', and C-C'.

Minimum Firmware Level

6.10

Revive Drive

This command forces the specified drive to the Optimal state.

ATTENTION Possible loss of data access – Correct use of this command depends on the data configuration on all of the drives in the volume group. Never try to revive a drive unless you are supervised by your Technical Support Representative.

Syntax

```
revive drive [trayID,drawerID,slotID]
```

Parameter

Parameter	Description
drive	The location of the drive that you want to revive. For high-capacity drive trays, specify the tray ID value, the drawer ID value, and the slot ID value of the drive that you want to revive. For low-capacity drive trays, specify the tray ID value and the slot ID value of the drive that you want to revive. Tray ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the tray ID value, drawer ID value, and the slot ID value in square brackets ([]).

Notes

The `drive` parameter supports both high-capacity drive trays and low-capacity drive trays. A high-capacity drive tray has drawers that hold the drives. The drawers slide out of the drive tray to provide access to the drives. A low-capacity drive tray does not have drawers. For a high-capacity drive tray, you must specify the identifier (ID) of the drive tray, the ID of the drawer, and the ID of the slot in which a drive resides. For

a low-capacity drive tray, you need only specify the ID of the drive tray and the ID of the slot in which a drive resides. For a low-capacity drive tray, an alternative method for identifying a location for a drive is to specify the ID of the drive tray, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

Minimum Firmware Level

5.43

7.60 adds the `drawerID` user input.

Revive Snapshot Group

This command forces the specified snapshot group to the Optimal state. If the snapshot group is not in a Failed state, the firmware displays an error message and does not run this command.

Syntax

```
revive snapGroup ["snapGroupName"]
```

Parameter

Parameter	Description
snapGroup	The name of the snapshot group that you want to set to the Optimal state. Enclose the snapshot group name in double quotation marks (" ") inside of square brackets ([]).

Notes

You can use any combination of alphanumeric characters, underscore (_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

Minimum Firmware Level

7.83

Revive Snapshot Volume

This command forces a specified snapshot volume to the Optimal state. The snapshot volume can be one of these:

- A standalone snapshot volume
- A snapshot volume that is a member of a consistency group

If the snapshot volume is not in a Failed state, the firmware displays an error message and does not run this command.

NOTE You cannot use this command for a snapshot volume that is used in online volume copy.

Syntax

```
revive snapVolume ["snapVolumeName"]
```

Parameter

Parameter	Description
snapVolume	The name of the snapshot volume that you want to set to the Optimal state. Enclose the snapshot group name in double quotation marks (" ") inside of square brackets ([]).

Notes

You can use any combination of alphanumeric characters, underscore (_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

Minimum Firmware Level

7.83

Revive Volume Group

This command forces the specified volume group and its associated failed drives to the Optimal state.

ATTENTION Possible loss of data access – Correct use of this command depends on the data configuration on all of the drives in the volume group. Never try to revive a drive unless you are supervised by your Technical Support Representative.

Syntax

```
revive volumeGroup [volumeGroupName]
```

Parameter

Parameter	Description
volumeGroup	The alphanumeric identifier (including - and _) of the volume group to be set to the Optimal state. Enclose the volume group identifier in square brackets ([]).

Minimum Firmware Level

6.10

Save Asynchronous Mirror Group Statistics

This command saves to a file the synchronization statistics for one or more member volumes in an asynchronous mirror group. The statistics collected are available only for asynchronous mirror group member volumes in a primary role on the local storage array.

A set of data is collected during the synchronization process that can be used to evaluate how well the mirror configuration is functioning. The data is collected as a set of *samples*. A sample is created at the beginning of a synchronization process and updated regularly while the synchronization process proceeds.

A sample collects data until the synchronization process completes or until a disruption to the synchronization process occurs such as a volume ownership transfer or a read-write error. When a synchronization process disruption is resolved (for example, the volume is transferred to the alternate controller), a new sample is created and updated as the synchronization process continues.

Syntax

```
save storageArray arvmStats asyncMirrorGroup
["asyncMirrorGroupName"]
arvmStats file="filename"
[volume="volumeName" |
sampleType=(all | mostRecent | longestSyncTime | errors) |
recordLimit=(1-90)]
```

Parameters

Parameter	Description
asyncMirrorGroup	The name of the asynchronous mirror group for which you are retrieving synchronization statistics. Enclose the asynchronous mirror group name in square brackets ([]). If the asynchronous mirror group name has special characters, you also must enclose the asynchronous mirror group name in double quotation marks (" ") inside square brackets. If the asynchronous mirror group name consists only of numbers, such as "1002," you also must enclose the asynchronous mirror group name in double quotation marks (" ") inside square brackets.
file	The file path and the file name to which you want to save the synchronization statistics. Enclose the file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\logs\iscsistat.csv" This command does not automatically append a file extension to the saved file. You can use any file name, but you must use the .csv extension.

Parameter	Description
volume	<p>This parameter is optional. The name of the specific member volume in the asynchronous mirror group for which you are retrieving synchronization statistics. If no volume is specified, the statistics for every member volume in the asynchronous mirror group are saved.</p> <p>If statistics are collected for more than one volume in an asynchronous mirror group, all the data will be written to the same file.</p> <p>Enclose the volume name in double quotation marks (" ").</p>
sampleType	<p>This parameter is optional. The default value for <code>sampleType</code> is <code>all</code>.</p> <ul style="list-style-type: none"> ■ <code>all</code>— Data for all three sample types are collected and written to the same file. ■ <code>mostRecent</code> — Statistics are recorded for the most recent 50 resynchronization samples. ■ <code>longestSyncTime</code> — Statistics are collected for the most recent 20 longest resynchronization samples. ■ <code>errors</code>— Statistics are recorded for the most recent 20 failed resynchronization samples. These samples include a failure code.
recordLimit	<p>This parameter is optional. The default value for record limit is no limit. The <code>recordLimit</code> must be greater than 0 and less than or equal to 90.</p>

Notes

Statistics are captured for mirrored volumes in the Primary role. The statistics collected include the following data:

- Synchronization start time
- Synchronization type (manual or periodic)
- Synchronization duration
- Number of bytes sent
- Maximum write time (for a single write)
- Minimum write time (for a single write)
- Minimum synchronization data rate
- Maximum synchronization data rate
- Total write time
- Repository utilization (%)
- Recovery point age

During initial synchronization, statistics samples are captured approximately every 15 minutes.

The synchronization statistics are included in the support bundle.

Minimum Firmware Level

7.84

Save Controller NVSRAM

This command saves a copy of the controller NVSRAM values to a file. This command saves all of the regions.

Syntax

```
save controller [(a | b)] NVSRAM file="filename"
```

Parameters

Parameter	Description
controller	The controller with the NVSRAM values that you want to save. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]).
file	The file path and the file name to which you want to save the NVSRAM values. Enclose the NVSRAM file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\logs\nvsramb.txt" The default name of the file that contains the NVSRAM values is nvsram-data.txt. This command does not automatically append a file extension to the saved file. You must specify a file extension when entering the file name.

Minimum Firmware Level

6.10

Save Drive Channel Fault Isolation Diagnostic Status

This command saves the drive channel fault isolation diagnostic data that is returned from the `start driveChannel faultDiagnostics` command. You can save the diagnostic data to a file as standard text or as XML.

See "[Start Drive Channel Fault Isolation Diagnostics](#)" for more information.

Syntax

```
save driveChannel faultDiagnostics file="filename"
```

Parameter

Parameter	Description
file	<p>The file path and the file name to which you want to save the results of the fault isolation diagnostics test on the drive channel. Enclose the file name in double quotation marks (" "). For example:</p> <pre>file="C:\Program Files\CLI\sup\fltdiag.txt"</pre> <p>This command does not automatically append a file extension to the saved file. You must specify a file extension when entering the file name.</p>

Notes

A file extension is not automatically appended to the saved file. You must specify the applicable format file extension for the file. If you specify a file extension of .txt, the output will be in a text file format. If you specify a file extension of .xml, the output will be in an XML file format.

Minimum Firmware Level

7.15 introduces this new capability for the CE7900 controller tray.

Save Drive Log

This command saves the log sense data to a file. Log sense data is maintained by the storage array for each drive.

Syntax

```
save allDrives logFile="filename"
```

Parameter

Parameter	Description
logFile	<p>The file path and the file name to which you want to save the log sense data. Enclose the file name in double quotation marks (" "). For example:</p> <pre>file="C:\Program Files\CLI\logs\lgsendat.txt"</pre> <p>This command does not automatically append a file extension to the saved file. You must specify a file extension when entering the file name.</p>

Minimum Firmware Level

6.10

Save Storage Array Configuration

This command creates a script file that you can use to re-create the current storage array volume configuration.

Syntax

```
save storageArray configuration file="filename"
[(allconfig | globalSettings=(TRUE | FALSE)
volumeConfigAndSettings=(TRUE | FALSE)
hostTopology=(TRUE | FALSE)
lunMappings=(TRUE | FALSE))]
```

Parameters

Parameter	Description
file	The file path and the file name to which you want to save the configuration settings. Enclose the file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\logs\sacnf.cfg" The default name of the file that contains the configuration settings is storage-array-configuration.cfg. This command does not automatically append a file extension to the saved file. You must specify a file extension when entering the file name.
allConfig	The setting to save all of the configuration values to the file. (If you choose this parameter, all of the configuration parameters are set to TRUE.)
globalSettings	The setting to save the global settings to the file. To save the global settings, set this parameter to TRUE. To prevent saving the global settings, set this parameter to FALSE. The default value is TRUE.
volumeConfigAndSettings	The setting to save the volume configuration settings and all of the global settings to the file. To save the volume configuration settings and global settings, set this parameter to TRUE. To prevent saving the volume configuration settings and global settings, set this parameter to FALSE. The default value is TRUE.
hostTopology	The setting to save the host topology to the file. To save the host topology, set this parameter to TRUE. To prevent saving the host topology, set this parameter to FALSE. The default value is FALSE.

Parameter	Description
lunMappings	The setting to save the LUN mapping to the file. To save the LUN mapping, set this parameter to TRUE. To prevent saving the LUN mapping, set this parameter to FALSE. The default value is FALSE.

Notes

When you use this command, you can specify any combination of the parameters for the global setting, the volume configuration setting, the host topology, or the LUN mapping. If you want to enter all settings, use the `allConfig` parameter. The parameters are all optional.

Minimum Firmware Level

6.10

Save Storage Array Core Dump

This command saves a core dump to a specified file on a host.

Syntax

```
save storageArray coreDump file="filename"
```

Parameter

Parameter	Description
file	The file path and the file name to which you save the core dump. Enclose the file name in double quotation marks (" "). For example: file="C:\Array Backups\DBMbackup_03302010.core"

Notes

A core dump file might be one to two gigabytes in size, and take several minutes to save. In addition to saving the core dump itself, this command generates an XML descriptor file based on the core dump metadata. This descriptor file is saved in ZIP file format to the same path as the core dump. The following example show the XML format for the descriptor file.

```
<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
- <DPLcoreDumpInfo>
<dplcoreDumpTag>someTag#</dplcoreDumpTag>
<captureTime>12/22/10 3:58:53 PM IST</captureTime>
- <StorageArrayData>
<ArraySAID>600a0b80006e006a000000004c65efc1</ArraySAID>
<ManagementApiVersion>devmgr.v1083api01.Manager</ManagementA
piVersion>
<fwVersion>07.83.01.00</fwVersion>
<platformType>7091</platformType>
```

```

</StorageArrayData>
<fullcoreDumpCtrl>controllerserialNumber1</fullcoreDumpCtrl>
<fullcoreDumpSize>fullCaptureSize</fullcoreDumpSize>
<altcoreDumpCtrl>controllerserialNumber2</altcoreDumpCtrl>
<altcoreDumpSize>altCaptureSize</altcoreDumpSize>
<triggerReason>Exception</triggerReason>-<DPLcoreDumpDetail>
<dplcoreDumpReason>DPLcoreDumpReason</dplcoreDumpReason>
- <fwLocation >
<filename>filename</filename>
<lineNumber>line#</lineNumber>
</fwLocation >
<panicReason>panicString</panicReason>
</DPLcoreDumpDetail>
</DPLcoreDumpInfo>

```

When you retrieve a core dump from the controller cache to a host, a flag is set on the controller to indicate that the core dump does not need to be retrieved. This setting persists for 48 hours. If a new core dump occurs during that period, the new core dump is saved to the controller cache and overwrites any previous core dump data in cache.

Minimum Firmware Level

7.83

Save Storage Array DBM Database

Use this command to back up RAID configuration data or all data to a file on the host. You can specify multiple data locations and controller.

Syntax

```

save storageArray dbmDatabase
file="fileName"
[sourceLocation=(disk | onboard) |
controller [(a|b)] |
contentType=(partial | all)]

```

Parameters

Parameter	Description
file	<p>The file path and the file name to which you want to save the database. Enclose the file name in double quotation marks (" "). For example:</p> <pre>file="C:\Program Files\CLI\logs\dbmdata.txt"</pre> <p>This parameter must appear last, after any of the optional parameters.</p> <p>This command does not automatically append a file extension to the saved file. You must specify a file extension when entering the file name.</p>

Parameter	Description
sourceLocation	<p>This parameter specifies the location from which to obtain backup database information.</p> <ul style="list-style-type: none"> ■ disk indicates that data comes directly from the database on the drive ■ onboard indicates that data comes from the RPA memory location <p>The default location is onboard.</p>
controller	<p>This parameter specifies the controller from which data will be exclusively retrieved, if sourceLocation is set to onboard. If the controller parameter is not specified, data might be retrieved from either controller.</p> <p>Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B.</p>
contentType	<p>This parameter specifies the content type of the data that will be retrieved.</p> <ul style="list-style-type: none"> ■ If the parameter is set to partial, a reduced set of records for the RAID configuration data is restored. This option helps reduce the possibility of record corruption by reducing the number and type of records restored. ■ If the parameter is set to all, all of the data including disk pool configuration data is retrieved. <p>The default is all.</p>

Notes

The data that you save to the host using this command can, if needed, be restored to the controller. A validator, however, is required to restore data from the file on the host.

Minimum Firmware Level

7.75

7.83 adds these parameters:

- sourceLocation
- controller
- contentType

Save Storage Array DBM Validator Information File

This command saves the database management (DBM) validation information for a storage array in an XML file. A Technical Support Representative can use this XML file to generate a validator string (a security code). The validator string must be included in the `load storageArray dbmDatabase` command when restoring a storage array back to a pre-existing configuration.

Syntax

```
save storageArray dbmValidatorInfo file="filename"  
dbmDatabase="filename"
```

Parameters

Parameter	Description
file	The file path and the file name of the DBM validator required for Technical Support. Enclose the file name in double quotation marks (" "). For example: file="C:\Array Backups \DBMvalidator.xml" This command does not automatically append a file extension to the saved file. You must specify the .xml file extension when you enter the file name.
dbmDatabase	The file path and the file name of the DBM database from which you want to generate the XML information file. Enclose the file name in double quotation marks (" "). For example: dbmDatabase="C:\Array Backups \DBMbackup_03302010.dbm" This command does not automatically append a file extension to the saved file. You must specify the .dbm file extension when you enter the file name.

Notes

If you are generating the XML information file to obtain a validator string, you must run this command while you are connected to the controller where you intend to restore the database. The following example shows the format of the XML file:

```
<?xml version="1.0" encoding="utf-8"?>  
<DbmImageMetadata>  
<Controllers>  
<A>1IT0nnnnnnnnnABCD</A>  
<B>1T0nnnnnnnnnABCD</B>  
</Controllers>  
<Header>  
<ImageVersion>1</ImageVersion>  
<TimeStamp>1269388283</TimeStamp>  
</Header>  
<Trailer>  
<CRC>nnnnnnnnnn</CRC><  
/Trailer>  
</DbmImageMetadata>
```

Minimum Firmware Level

7.75

Save Storage Array Diagnostic Data

This command saves the storage array diagnostic data from either the controllers or the environmental services modules (ESMs) to a file. You can review the file contents at a later time. You can also send the file to your Technical Support Representative for further review.

After you have saved the diagnostic data, you can reset the NVSRAM registers that contain the diagnostic data so that the old data can be overwritten. Use the `reset storageArray diagnosticData` command to reset the diagnostic data registers.

ATTENTION Run this command only with the assistance of your Technical Support Representative.

Syntax

```
save storageArray diagnosticData [(controller | tray)]  
file="filename"
```

Parameters

Parameter	Description
diagnosticData	This parameter allows you to download the diagnostic data from either the controllers or the ESMs.
file	The file path and the file name to which you want to save the storage array diagnostic data. Enclose the file name in double quotation marks (" "). For example: <code>file="C:\Program Files\CLI\logs\sadiag.zip"</code> This command automatically saves the data to a compressed file; however, this command does not automatically append a file extension to the saved file. You must specify the <code>.zip</code> extension when entering the file name.

Notes

In versions of this command before 10.77, the user option was `esm` instead of `tray`. Starting in 10.77 `tray` replaces `esm`. The use of `esm` is still supported, but for best compatibility with future releases, replace `esm` with `tray`.

Minimum Firmware Level

6.16

7.77 `tray` replaces `esm`.

Save Storage Array Events

This command saves events from the Major Event Log to a file. You can save these events:

- **Critical events** – An error occurred on the storage array that needs to be addressed immediately. Loss of data access might occur if you do not immediately correct the error.
- **Warning events** – An error occurred on the storage array that results in degraded performance or reduced ability to recover from another error. Access to data has not been lost, but you must correct the error to prevent possible loss of data access if another error would occur.
- **Informational events** – An event occurred on the storage array that does not impact normal operations. The event is reporting a change in configuration or other information that might be useful in evaluating how well the storage array is performing.
- **Debug events** – An event occurred on the storage array that provides information that you can use to help determine the steps or states that led to an error. You can send a file with this information to your Technical Support Representative to help determine the cause of an error.

NOTE Some storage arrays might not be able to support all four types of events.

Syntax

```
save storageArray (allEvents | criticalEvents |  
warningEvents | infoEvents | debugEvents)  
file="filename"  
[count=numberOfEvents  
forceSave=(TRUE | FALSE)]
```

Parameters

Parameter	Description
allEvents	The parameter to save all of the events to a file.
criticalEvents	The parameter to save only the critical events to a file.
warningEvents	The parameter to save only the warning events to a file.
infoEvents	The parameter to save only the informational events to a file.
debugEvents	The parameter to save only the debug events to a file.

Parameter	Description
file	The file path and the file name to which you want to save the events. Enclose the file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\logs\events.txt" The default name of the file that contains the contents of the Major Event Log is major-event-log.txt. This command does not automatically append a file extension to the saved file. You must specify a file extension when entering the file name.
count	The number of events or critical events that you want to save to a file. If you do not enter a value for the count, all events or all critical events are saved to the file. If you enter a value for the count, only that number of events or critical events (starting with the last event entered) are saved to the file. Use integer values.
forceSave	The parameter to force saving the critical events to a file. To force saving the events, set this parameter to TRUE. The default value is FALSE.

Notes

You have the option to save all events (`allEvents`) or only the critical events (`criticalEvents`).

Minimum Firmware Level

6.10

7.77 add these parameters:

- `warningEvents`
- `infoEvents`
- `debugEvents`
- `forceSave`

Save Storage Array Firmware Inventory

This command saves a report to a file of all of the firmware currently running on the storage array. The report lists the firmware for these components:

- Controllers
- Drives
- Drawers (if applicable)
- Environmental services modules (ESMs)

You can use the information to help identify out-of-date firmware or firmware that does not match the other firmware in your storage array. You can also send the report to your Technical Support Representative for further review.

Syntax

```
save storageArray firmwareInventory file="filename"
```

Parameter

Parameter	Description
file	<p>The file path and the file name to which you want to save the firmware inventory. Enclose the file name in double quotation marks (" "). For example:</p> <pre>file="C:\Program Files\CLI\logs\fwinvent.txt"</pre> <p>The default name of the file that contains the firmware inventory is <code>firmware-inventory.txt</code>. This command does not automatically append a file extension to the saved file. You must specify a file extension when entering the file name.</p>

Minimum Firmware Level

7.70

Save Storage Array InfiniBand Statistics

This command saves the InfiniBand performance statistics of the storage array to a file.

Syntax

```
save storageArray ibStats [raw | baseline]
file="filename"
```

Parameters

Parameter	Description
raw	The statistics that are collected are all statistics from the controller start-of-day. Enclose the parameter in square brackets ([]).
baseline	The statistics that are collected are all statistics from the time the controllers were reset to zero using the <code>reset storageArray ibStatsBaseline</code> command. Enclose the parameter in square brackets ([]).
file	<p>The file path and the file name to which you want to save the performance statistics. Enclose the file name in double quotation marks (" "). For example:</p> <pre>file="C:\Program Files\CLI\sup\ibstat.txt"</pre> <p>The default name of the file that contains the InfiniBand performance statistics is <code>ib-statistics.csv</code>. This command does not automatically append a file extension to the saved file. You must specify a file extension when entering the file name.</p>

Notes

If you have not reset the InfiniBand baseline statistics since the controller start-of-day, the time at the start-of-day is the default baseline time.

Minimum Firmware Level

7.32

Save Storage Array iSCSI Statistics

This command saves the iSCSI performance of the storage array to a file.

Syntax

```
save storageArray iscsiStatistics [raw | baseline]  
file="filename"
```

Parameters

Parameter	Description
raw	The statistics collected are all statistics from the controller start-of-day. Enclose the parameter in square brackets ([]).
baseline	The statistics that are collected are all statistics from the time the controllers were reset to zero using the <code>reset storageArray ibStatsBaseline</code> command. Enclose the parameter in square brackets ([]).
file	The file path and the file name to which you want to save the performance statistics. Enclose the file name in double quotation marks (" "). For example: <code>file="C:\Program Files\CLI\logs\iscsistat.csv"</code> The default name of the file that contains the iSCSI performance statistics is <code>iscsi-statistics.csv</code> . You can use any file name but you must use the <code>.csv</code> extension.

Notes

If you have not reset the iSCSI baseline statistics since the controller start-of-day, the time at the start-of-day is the default baseline time.

Minimum Firmware Level

7.10

Save Storage Array Performance Statistics

This command saves the performance statistics to a file. Before you use this command, run the `set session performanceMonitorInterval` command and the `set session performanceMonitorIterations` command to specify how often statistics are collected.

Syntax

```
save storageArray performanceStats file="filename"
```

Parameter

Parameter	Description
file	The file path and the file name to which you want to save the performance statistics. Enclose the file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\logs\sastat.csv" The default name of the file that contains the performance statistics is performanceStatistics.csv. You can use any file name, but you must use the .csv extension.

Minimum Firmware Level

6.10

Save Storage Array RLS Counts

This command saves the read link status (RLS) counters to a file.

Syntax

```
save storageArray RLSCounts file="filename"
```

Parameter

Parameter	Description
file	The file path and the file name to which you want to save the RLS counters. Enclose the file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\logs\rlscnt.csv" The default name of the file that contains the RLS counts is readLinkStatus.csv. You can use any file name, but you must use the .csv extension.

Notes

To more effectively save RLS counters to a file, perform these steps:

1. Run the `reset storageArray RLSBaseline` command to set all of the RLS counters to 0.
2. Run the storage array for a predetermined amount of time (for instance, two hours).
3. Run the `save storageArray RLSCounts file="filename"` command.

Minimum Firmware Level

6.10

Save Storage Array SAS PHY Counts

This command saves the SAS physical layer (SAS PHY) counters to a file. To reset the SAS PHY counters, run the `reset storageArray SASPHYBaseline` command.

Syntax

```
save storageArray SASPHYCounts file="filename"
```

Parameter

Parameter	Description
file	The file path and the file name to which you want to save the SAS PHY counters. Enclose the file path and the file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\logs\sasphy.csv" The default name of the file that contains the SAS PHY error statistics is <code>sas-phy-error-log.csv</code> . You can use any file name but you must use the <code>.csv</code> extension.

Minimum Firmware Level

6.10

Save Storage Array SOC Counts

This command saves the SOC error statistics to a file. This command is valid only for Fibre Channel devices in an arbitrated loop topology.

Syntax

```
save storageArray SOCCounts file="filename"
```

Parameter

Parameter	Description
file	The file path and the file name to which you want to save the SOC error statistics. Enclose the file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\logs\socstat.csv" The default name of the file that contains the SOC error statistics is <code>soc-statistics.csv</code> . You can use any file name but you must use the <code>.csv</code> extension.

Notes

To more effectively save SOC error statistics to a file, perform these steps:

1. Run the `reset storageArray SOCBaseline` command to set all of the SOC counters to 0.
2. Run the storage array for a predetermined amount of time (for example, two hours).
3. Run the `save storageArray SOCCounts file="filename"` command.

Minimum Firmware Level

6.16

Save Storage Array State Capture

This command saves the state capture of a storage array to a file.

Syntax

```
save storageArray stateCapture file="filename"
```

Parameter

Parameter	Description
file	The file path and the file name to which you want to save the state capture. Enclose the file name in double quotation marks (" "). For example: <code>file="C:\Program Files\CLI\logs\state.zip"</code> The default name of the file that contains the state capture is <code>state-capture-data.txt</code> . This command does not automatically append a file extension to the saved file. You must specify a file extension when entering the file name.

Minimum Firmware Level

6.10

Save Storage Array Support Data

This command saves the support-related information of the storage array to a file. Support-related information includes these items:

- The storage array profile
- The Major Event Log information
- The read link status (RLS) data
- The NVSRAM data
- Current problems and associated recovery information
- The performance statistics for the entire storage array
- The persistent registration information and the persistent reservation information
- Detailed information about the current status of the storage array

- The diagnostic data for the drive
- A recovery profile for the storage array
- The unreadable sectors that are detected on the storage array
- The state capture data
- An inventory of the versions of the firmware running on the controllers, the drives, the drawers, and the environmental services modules (ESMs)

Syntax

```
save storageArray supportData file="filename"
```

Parameters

Parameter	Description
file	<p>The file path and the file name to which you want to save the support-related data for the storage array. Enclose the file path and the file name in double quotation marks (" "). For example:</p> <pre>file="C:\Program Files\CLI\logs\supdat.7z"</pre> <p>This command automatically saves the data to a compressed file and adds the file extension .7z to the saved file.</p> <p>If you are running firmware version 7.86 the file extension must be .7z.</p> <p>If you are running a firmware version earlier than 7.86 the file extension must be .zip.</p>

Notes

Starting with firmware level 7.86 the file name extension must be .7z on systems running Windows. If you are running a firmware version earlier than 7.86 on a Windows system, the files extension must be .zip. On systems that are not running Windows you can use any file extension that works for your system.

Minimum Firmware Level

6.10

7.86 Requires that the file name uses a .7z extension.

Save Tray Log

This command saves the log sense data to a file. Log sense data is maintained by the environmental cards for each tray. Not all of the environmental cards contain log sense data.

Syntax

```
save allTrays logFile="filename"
```

Parameter

Parameter	Description
logFile	The file path and the file name to which you want to save the log sense data. Enclose the file name in double quotation marks ("). For example: file="C:\Program Files\CLI\logs\traylogdat.txt" This command does not automatically append a file extension to the saved file. You must specify a file extension when entering the file name.

Minimum Firmware Level

6.50

Schedule Automatic Support Bundle Collection Configuration

NOTE This command is an SMcli command, not a script command. You must run this command from a command line. You cannot run this command from the script editor in the storage management software

This command creates or removes a schedule for saving a support bundle on one or more storage arrays.

Syntax

```
SMcli -supportBundle schedule (enable|disable)
(all|storageArrayName)
[data=pathName |
startTime=HH:MM |
startDate=MM:DD:YYYY] |
endDate=MM:DD:YYYY] |
(daysOfWeek=(Sunday Monday Tuesday Wednesday Thursday Friday
Saturday) |
dayOfWeek=(Sunday|Monday|Tuesday|Wednesday|Thursday|Friday|S
aturday) |
months=(January February March April May June July August
September October November December)
onDays=(1-31) |
weekNumber=(First|Second|Third|Fourth|Last))
```

Parameters

Parameter	Description
enable	Enable automatic collection of support bundles when a critical MEL event is detected. This parameter resumes collecting support bundle data if it has been suspended. You must use the parameter whenever you create a new schedule or modify an existing schedule.

Parameter	Description
disable	<p>Disables the automatic collection of support bundles and deletes any previously defined schedules immediately.</p> <hr/> <p>NOTE Disabling a schedule also deletes the schedule.</p> <hr/>
all	Use this parameter if you want to set a collection schedule for all of the storage arrays detected from the host.
<i>storageArrayName</i>	The name of a specific storage array for which you want to set a scheduled.
data	<p>The file path and the file name to which you want to save the support bundle data. For example: file="C:\Program Files\CLI\sup\data.txt"</p> <hr/> <p>NOTE You can use any file extension.</p> <hr/> <p>Enclose the file path and name in double quotation marks (" ").</p>
startTime	<p>The time of a day that you want the collection of a support bundle to start. The format for entering the time is HH:MM, where HH is the hour and MM is the minute past the hour. Use a 24-hour clock.</p> <p>The default time is 00:00, midnight. If you do not enter a time to start and you have enabled support bundle collection, collection of the data automatically occurs at midnight.</p>
startDate	<p>A specific date on which you want to start collecting support bundle data. The format for entering the date is MM:DD:YY.</p> <p>The default date is the current date.</p>
endDate	<p>A specific date on which you want to stop collecting support bundle data. The format for entering the date is MM:DD:YY.</p> <p>The default is none.</p>
daysOfWeek	<p>Specific days of the week on which you want to collect support bundle data. Use this parameter when you want to collect support bundle data on several days of week. For example: daysOfWeek=Monday Wednesday Friday</p> <p>Enclose the days of the week in parentheses and separate each day with a white space.</p>
dayOfWeek	<p>A specific day of the week on which you want to collect support bundle data. Use this parameter when you want to collect support bundle data on only one day of the week. For example: dayOfWeek=Wednesday</p>

Parameter	Description
months	<p>Specific months of the year on which you want to collect support bundle data. You can use this parameter when you want to collect support bundle data on several days of week or for only a single month. An example or setting a schedule for several month is this: months=(January April July October)</p> <p>Enclose the months in parentheses and separate each month with a white space.</p> <p>An example or setting a schedule for a single month is this: months=June</p> <p>You do not need to enclose the month in parentheses.</p>
onDays	<p>A day in the month on which you want to collect support bundle data. For example: onDays=15</p>
weekNumber	<p>A week in the month during which you want to collect support bundle data. For example: weekNumber=first</p>

Notes

When you use the `all` parameter to set a common schedule for all of the storage arrays, the schedules for individual storage arrays are deleted. If a schedule is set for all storage arrays, newly discovered storage arrays will follow the same schedule. If a schedule is set for all storage arrays and then a schedule is set for a single storage array, newly discovered storage arrays will not have a schedule defined.

Minimum Firmware Level

7.83

Set Alert Severities

NOTE This command is an SMcli command, not a script command. You must run this command from a command line. You cannot run this command from the script editor in the storage management software

This command enables you to set the level of severities that causes an alert to be sent to the Windows event log. The alert severities apply to all of the storage arrays in the entire storage system.

Syntax

```
SMcli -alertSeverities (severity |
[severity1, ... severityN])
```

Parameter

Parameter	Description
severity	<p>Use this parameter to set the severity for returning alerts. The alert severity values that you can set are these:</p> <ul style="list-style-type: none">■ <code>critical</code> – Alerts will be sent■ <code>warning</code> – Alerts will be sent■ <code>informational</code> – Alerts will not be sent■ <code>debug</code> – Alerts will be sent <hr/> <p>NOTE The debug value is for Technical Support only. Do not attempt to use this value.</p> <hr/> <p>You can set one or more alert severities values. If you set more than one alert severities value, enclose all of the values in square brackets ([]) and separate the values by a comma.</p>

Minimum Firmware Level

7.83

Set Asynchronous Mirror Group

Use this command to change the synchronization settings and warning thresholds that the controller owner of the primary side of the asynchronous mirror group uses when it performs an initial synchronization or resynchronization. Changing the synchronization settings affects the synchronization operations of all mirrored pairs within the asynchronous mirror group.

Syntax

```
set asyncMirrorGroup ["asyncMirrorGroupName"]
[syncInterval=integer (minutes | hours | days)
warningSyncThreshold=integer (minutes | hours | days)
warningRecoveryThreshold=integer (minutes | hours | days)
warningThresholdPercent=percentValue
userLabel="New_asyncMirrorGroupName"
autoResync=(TRUE | FALSE)
volume="repos_xxxx" increaseRepositoryCapacity
(repositoryVolume=("repos_xxxx") |
repositoryVolumes=("repos_xxxx" "repos_xxxx"))
role=(primary | secondary)
(force=TRUE | FALSE | noSync=TRUE | FALSE)
```

Parameters

Parameter	Description
asyncMirrorGroup	The name of the asynchronous mirror group that you want to modify. Enclose the asynchronous mirror group name in square brackets ([]). If the asynchronous mirror group name has special characters, you also must enclose the name in double quotation marks (" ") inside square brackets. If the asynchronous mirror group name consists only of numbers, such as "1002," you also must enclose the name in double quotation marks (" ") inside square brackets.
syncInterval	Specify the length of time between automatically sending updates of modified data from the local storage array to the remote storage array. You can specify the length of time in minutes, hours, or days.
warningSyncThreshold	Specify the length of time to wait until a warning is triggered for cases in which the synchronization of all of the volumes within the asynchronous mirror group takes longer than the defined time. You can specify the length of time in minutes, hours, or days.
warningRecoveryThreshold	Specify the length of time to wait until a warning is triggered when the automatic data update for the point-in-time image on the remote storage array is older than the defined time. Define the threshold from the end of the previous update. You can specify the length of time in minutes, hours, or days. NOTE You must set the Recovery Point Threshold to be twice as large as the synchronization interval threshold.
warningThresholdPercent	Specify the percent value that determines when a warning is triggered when the capacity of a mirror repository volume reaches the defined percentage. Define the threshold by percentage (%) of the capacity remaining.
userLabel	The new name that you want to give to the asynchronous mirror group. Use this parameter when you want to rename the asynchronous mirror group. Enclose the new asynchronous mirror group name in double quotation marks (" ").

Parameter	Description
autoResync	<p>The settings for automatic resynchronization between the primary volumes and the secondary volumes of an asynchronous mirrored pair within an asynchronous mirror group. This parameter has these values:</p> <ul style="list-style-type: none"> ■ <code>enabled</code>– Automatic resynchronization is turned on. You do not need to do anything further to resynchronize the primary volume and the secondary volume. ■ <code>disabled</code>– Automatic resynchronization is turned off. To resynchronize the primary volume and the secondary volume, you must run the <code>resume asyncMirrorGroup</code> command.
volume	<p>The name of an asynchronous mirror group repository volume for which you want to increase the capacity.</p> <p>The name of a repository volume is comprised of two parts:</p> <ul style="list-style-type: none"> ■ The term <i>repos</i> ■ A four digit numerical identifier that the storage management software assigns to the repository volume name <p>Enclose the name of the repository volume in double quotation marks (" ").</p>

Parameter	Description
<p>repositoryVolume or repositoryVolumes</p>	<p>The name of an unused repository volume that you want to use to increase the capacity of another repository volume. An available standard volume is added to the repository volume to increase the capacity of the repository volume.</p> <p>The name of the unused repository volume is comprised of two parts:</p> <ul style="list-style-type: none"> ■ The term <i>repos</i> ■ A four digit numerical identifier that the storage management software assigns to the repository volume name <p>Enclose the name of the existing repository volume in double quotation marks (" ").</p> <p>You can enter more than one repository volume name. Enclose each repository volume name in double quotation marks (" ") inside parentheses (()). Separate each repository volume name with a white space.</p>
<p>role</p>	<p>Use this parameter to promote the asynchronous mirror group to a primary role or demote the asynchronous mirror group to a secondary role. To define the asynchronous mirror group as the primary role, set this parameter to <code>primary</code>. To define the asynchronous mirror group as the secondary role, set this parameter to <code>secondary</code>.</p>
<p>force</p>	<p>The role reversal is forced if the communications link between the storage arrays is down and promotion or demotion on the local side results in a dual-primary condition or a dual-secondary condition. To force a role reversal, set this parameter to <code>TRUE</code>. The default value is <code>FALSE</code>.</p>
<p>nosync</p>	<p>This parameter defines whether to perform an initial synchronization between the primary storage array and the secondary storage array before a role reversal operation is performed. To perform an initial synchronization, set this parameter to <code>TRUE</code>. The default value is <code>FALSE</code>.</p>

Notes

You can use any combination of alphanumeric characters, underscore (`_`), hyphen (`-`), and pound (`#`) for the names. Names can have a maximum of 30 characters.

When you use this command, you can specify one or more of the parameters. You do not, however, need to use all of the parameters.

An asynchronous mirror group repository volume is an expandable volume that is structured as a concatenated collection of up to 16 standard volume entities. Initially, an expandable repository volume has only a single volume. The capacity of the expandable repository volume is exactly that of the single volume. You can increase the capacity of an expandable repository volume by attaching additional unused repository volumes to it. The composite expandable repository volume capacity then becomes the sum of the capacities of all of the concatenated standard volumes.

To increase the capacity of a repository volume use the `set asyncMirrorGroup` command as shown in this example:

```
set asyncMirrorGroup ["amg_001"] volume="repos_006"
increaseRepositoryCapacity
repositoryVolumes=("repos_0021" "repos_0021");
```

An asynchronous mirror group repository volume must satisfy a minimum capacity requirement that is the sum of the following:

- 32 MB to support fixed overhead for the asynchronous mirror group and for copy-on-write processing.
- Capacity for rollback processing, which is 1/5000th of the capacity of the base volume.

The minimum capacity is enforced by the controller firmware and the storage management software.

Minimum Firmware Level

7.84

Set Consistency Group Attributes

This command defines the properties for a snapshot consistency group.

Syntax

```
set consistencyGroup ["consistencyGroupName"]
[userLabel="consistencyGroupName" |
repositoryFullPolicy=(failBaseWrites | purgeSnapImages)|
repositoryFullLimit=percentValue|
autoDeleteLimit=numberOfSnapImages|
rollbackPriority=(lowest | low | medium | high | highest)]
```

Parameters

Parameter	Description
<code>consistencyGroupName</code>	The name of the consistency group for which you are setting properties. Enclose the consistency group name in double quotation marks (" ") inside square brackets ([]).
<code>userLabel</code>	The new name that you want to give to the snapshot consistency group. Enclose the new consistency group name in double quotation marks (" ").
<code>repositoryFullPolicy</code>	How you want snapshot processing to continue if the snapshot repository volumes are full. You can choose to fail writes to the base volume (<code>failBaseWrites</code>) or delete (purge) the snapshot images (<code>purgeSnapImages</code>). The default action is <code>purgeSnapImages</code> .
<code>repositoryFullLimit</code>	The percentage of repository capacity at which you receive a warning that the snapshot repository volume is nearing full. Use integer values. For example, a value of 70 means 70 percent. The default value is 75.
<code>autoDeleteLimit</code>	Each consistency group can be configured to perform automatic deletion of its snapshot images to keep the total number of snapshot images in the snapshot group at or below a designated level. When this option is enabled, then any time a new snapshot image is created in the snapshot group, the system automatically deletes the oldest snapshot image in the group to comply with the limit value. This action frees repository capacity so it can be used to satisfy ongoing copy-on-write requirements for the remaining snapshot images.
<code>rollbackPriority</code>	The priority for rollback operations for a consistency group while the storage array is operational. Valid values are <code>highest</code> , <code>high</code> , <code>medium</code> , <code>low</code> , or <code>lowest</code> .

Notes

You can use any combination of alphanumeric characters, underscore (`_`), hyphen (`-`), and pound (`#`) for the names. Names can have a maximum of 30 characters.

When you use this command, you can specify one or more of the parameters. You do not, however, need to use all of the parameters.

The rollback priority defines the amount of system resources that should be allocated to the rollback operation at the expense of system performance. A value of `high` indicates that the rollback operation is prioritized over all other host I/O. A value of `low` indicates that the rollback operation should be performed with minimal impact to host I/O.

Auto Delete

You can configure each snapshot group to perform automatic deletion of its snapshot images to keep the total number of snapshot images in the snapshot group at or below a maximum number of images. When the number of snapshot images in the snapshot group is at the maximum limit, the `autoDeleteLimit` parameter automatically deletes snapshot images whenever a new snapshot image is created in the snapshot group. The `autoDeleteLimit` parameter deletes the oldest snapshot images in the snapshot group until the maximum number of images defined with the parameter is met. Deleting snapshot images in this way frees repository capacity so it can be used to satisfy ongoing copy-on-write requirements for the remaining snapshot images.

Minimum Firmware Level

7.83

Set Consistency Group Snapshot Volume

This command creates a unique name for a snapshot volume of a consistency group.

Syntax

```
set cgSnapVolume [cgSnapVolumeName]  
userLabel="cgSnapVolumeName"
```

Parameter

Parameter	Description
<code>cgSnapVolume</code>	The name of the consistency group volume that you want to rename. Enclose the name of the consistency group snapshot volume in square brackets ([]).
<code>userLabel</code>	The new name that you want to give to the snapshot volume in the consistency group. Enclose the new snapshot volume name in double quotation marks (" ").

Notes

You can use any combination of alphanumeric characters, underscore (`_`), hyphen (`-`), and pound (`#`) for the names. Names can have a maximum of 30 characters.

Minimum Firmware Level

7.83

Set Controller

This command defines the attributes for the controllers.

Syntax

```
set controller [(a | b)]
availability=(online | offline | serviceMode) |
ethernetPort [(1 | 2)] ethernetPortOptions |
globalNVSramByte [nvsramOffset]=(nvsramByteSetting |
nvsramBitSetting) |
hostNVSramByte [hostType, nvsramOffset]=(nvsramByteSetting |
nvsramBitSetting) |
IPv4GatewayIP=ipAddress |
IPv6RouterAddress=ipv6Address |
iscsiHostPort [(1 | 2 | 3 | 4)] iscsiHostPortOptions |
rloginEnabled=(TRUE | FALSE) |
serviceAllowedIndicator=(on | off)
```

Parameters

Parameter	Description
controller	The controller for which you want to define properties. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the identifier for the controller in square brackets ([]). If you do not specify a controller, the firmware for the controller returns a syntax error.
availability	The mode for the controller, which you can set to online, offline, or serviceMode (service).
ethernetPort	The attributes (options) for the management Ethernet ports. The entries to support this parameter are listed in the Syntax Element Statement Data table that follows. Many settings are possible, including setting the IP address, the gateway address, and the subnet mask address.
globalNVSramByte	A portion of the controller NVSRAM. Specify the region to be modified using the starting byte offset within the region and the byte value or bit value of the new data to be stored into the NVSRAM.
hostNVSramByte	The NVSRAM for the host-specific region. The setting specifies the host index for the specific host, the starting offset within the region, the number of bytes, and the the byte value or bit value of the new data to be stored into the NVSRAM.
IPv4GatewayIP	The IP address of the node that provides the interface to the network. The address format for the IPv4 gateway is (0–255).(0–255).(0–255).(0–255)

Parameter	Description
IPv6RouterAddress	The IP address of IPv6 router that connects two or more logical subnets. The address format for the IPv6 router is (0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF).
iscsiHostPort	The values that support this parameter are listed in the Syntax Element Statement Data table that follows. Many settings are possible, including setting the IP address, the gateway address, the subnet mask address, the IPv4 priority, and the IPv6 priority.
rloginEnabled	The setting for whether the remote login feature is turned on or turned off. To turn on the remote login feature, set this parameter to TRUE. To turn off the remote login feature, set this parameter to FALSE.
serviceAllowedIndicator	The setting for whether the Service Action Allowed indicator light is turned on or turned off. To turn on the Service Action Allowed indicator light, set this parameter to on. To turn off the Service Action Allowed indicator light, set this parameter to off.

Syntax Element Statement Data

Options for the ethernetPort Parameter
enableIPv4=(TRUE FALSE)
enableIPv6=(TRUE FALSE)
IPv6LocalAddress=(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF)
IPv6RoutableAddress=(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF)
IPv4Address=(0-255).(0-255).(0-255).(0-255)
IPv4ConfigurationMethod=[(static dhcp)]
IPv4SubnetMask=(0-255).(0-255).(0-255).(0-255)
duplexMode=(TRUE FALSE)
portSpeed=[(autoNegotiate 10 100 1000)]

Options for the `iscsiHostPort` Parameter

```
IPv4Address=(0-255).(0-255).(0-255).(0-255) |
IPv6LocalAddress=(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):
(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF) |
IPv6RoutableAddress=(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):
(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF) |
IPv6RouterAddress=(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):
(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF) |
enableIPv4=(TRUE | FALSE) |
enableIPv6=(TRUE | FALSE) |
enableIPv4Vlan=(TRUE | FALSE) |
enableIPv6Vlan=(TRUE | FALSE) |
enableIPv4Priority=(TRUE | FALSE) |
enableIPv6Priority=(TRUE | FALSE) |
IPv4ConfigurationMethod=(static | dhcp) |
IPv6ConfigurationMethod=(static | auto) |
IPv4GatewayIP=(TRUE | FALSE) |
IPv6HopLimit=[0-255] |
IPv6NdDetectDuplicateAddress=[0-256] |
IPv6NdReachableTime=[0-65535] |
IPv6NdRetransmitTime=[0-65535] |
IPv6NdTimeOut=[0-65535] |
IPv4Priority=[0-7] |
IPv6Priority=[0-7] |
IPv4SubnetMask=(0-255).(0-255).(0-255).(0-255) |
IPv4VlanId=[1-4094] |
IPv6VlanId=[1-4094] |
maxFramePayload=[frameSize] |
tcpListeningPort=[3260, 49152-65536] |
portSpeed=[(autoNegotiate | 1 | 10)]
```

Notes

NOTE Before firmware version 7.75, the `set controller` command supported an `NVSRAMByte` parameter. The `NVSRAMByte` parameter is deprecated and must be replaced with either the `hostNVSRAMByte` parameter or the `globalNVSRAMByte` parameter.

When you use this command, you can specify one or more of the parameters. You do not need to use all of the parameters.

Setting the `availability` parameter to `serviceMode` causes the alternate controller to take ownership of all of the volumes. The specified controller no longer has any volumes and refuses to take ownership of any more volumes. Service mode is persistent across reset cycles and power cycles until the `availability` parameter is set to `online`.

Use the `show controller NVSRAM` command to show the NVSRAM information. Before making any changes to the NVSRAM, contact your Technical Support Representative to learn what regions of the NVSRAM you can modify.

When the `duplexMode` option is set to `TRUE`, the selected Ethernet port is set to full duplex. The default value is half duplex (the `duplexMode` parameter is set to `FALSE`).

To make sure that the IPv4 settings or the IPv6 settings are applied, you must set these `iscsiHostPort` options:

- `enableIPV4=TRUE`
- `enableIPV6=TRUE`

The IPv6 address space is 128 bits. It is represented by eight 16-bit hexadecimal blocks separated by colons.

The `maxFramePayload` option is shared between IPv4 and IPv6. The payload portion of a standard Ethernet frame is set to 1500, and a jumbo Ethernet frame is set to 9000. When using jumbo frames, all of the devices that are in the network path should be capable of handling the larger frame size.

The `portSpeed` option is expressed as megabits per second (Mb/s).

Values for the `portSpeed` option of the `iscsiHostPort` parameter are in megabits per second (Mb/s).

The following values are the default values for the `iscsiHostOptions`:

- The `IPv6HopLimit` option is 64.
- The `IPv6NdReachableTime` option is 30000 milliseconds.
- The `IPv6NdRetransmitTime` option is 1000 milliseconds.
- The `IPv6NdTimeOut` option is 30000 milliseconds.
- The `tcpListeningPort` option is 3260.

Minimum Firmware Level

7.15 removed the `bootp` parameter, and added the new Ethernet port options and the new iSCSI host port options.

7.50 moved the `IPV4Gateway` parameter and the `IPV6RouterAddress` parameter from the iSCSI host port options to the command.

7.60 adds the `portSpeed` option of the `iscsiHostPort` parameter.

7.75 deprecates the `NVSRAMByte` parameter.

Set Controller Service Action Allowed Indicator

This command turns on or turns off the Service Action Allowed indicator light on a controller in a controller tray or a controller-drive tray. If the storage array does not support the Service Action Allowed indicator light feature, this command returns an error. If the storage array supports the command but is unable to turn on or turn off the indicator light, this command returns an error. (To turn on or turn off the Service Action Allowed indicator light on the power-fan canister or the interconnect-battery canister, use the `set tray serviceAllowedIndicator` command.)

Syntax

```
set controller=[(a | b)]
serviceAllowedIndicator=(on | off)
```

Parameters

Parameter	Description
<code>controller</code>	The controller that has the Service Action Allowed indicator light that you want to turn on or turn off. Valid controller identifiers are <code>a</code> or <code>b</code> , where <code>a</code> is the controller in slot A, and <code>b</code> is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the controller firmware returns a syntax error.
<code>serviceAllowedIndicator</code>	The setting to turn on or turn off the Service Action Allowed indicator light. To turn on the Service Action Allowed indicator light, set this parameter to <code>on</code> . To turn off the Service Action Allowed indicator light, set this parameter to <code>off</code> .

Notes

This command was originally defined for use with the CE6998 controller tray. This command is not supported by controller trays that were shipped before the introduction of the CE6998 controller tray. The 3992 and 3994 controllers also support this command.

Minimum Firmware Level

6.14

Set Disk Pool

This command sets the attributes associated with a disk pool based on the specified parameters.

Syntax

```
set diskPool (diskPool=[diskPoolName] |
diskPools=["diskPoolName1" ... "diskPoolNameN"] |
allDiskPools)
[reservedDriveCount=reservedDriveCountValue |
warningThreshold=(warningThresholdValue | default) |
criticalThreshold=(criticalThresholdValue | default) |
criticalPriority=(highest|high|medium|low|lowest) |
degradedPriority=(highest|high|medium|low|lowest) |
backgroundPriority=(highest|high|medium|low|lowest) |
userLabel=diskPoolName]
```

Parameters

Parameter	Description
<code>diskPool</code> or <code>diskPools</code>	The name of the disk pool for which you are setting attributes. Enclose the disk pool name in square brackets ([]). You can enter more than one disk pool name. Enclose all of the disk pool names in one set of square brackets ([]). Enclose each disk pool name in double quotation marks (" "). Separate each disk pool name with a white space.
<code>allDiskPools</code>	This parameter sets attributes for all of the disk pools in the storage array.
<code>reservedDriveCount</code>	This parameter sets the number of drives to be reserved in the disk pool for drive failures or drive removals. Setting this value to 0 results in a possible loss of data if a drive failure occurs in a disk pool.

Parameter	Description
warningThreshold	<p>The percentage of disk pool capacity at which you receive a warning alert that the disk pool is nearing full. Use integer values. For example, a value of 70 means 70 percent. For best operation, the value for this parameter must be less than the value for the <code>criticalThreshold</code> parameter.</p> <p>Valid values are from 0 to 100.</p> <p>The default value is 50.</p> <p>Setting this parameter to 0 disables warning alerts.</p> <p>If you set this to <code>default</code>, the warning alert threshold value is determined by the controller firmware.</p>
criticalThreshold	<p>The percentage of disk pool capacity at which you receive a critical alert that the disk pool is nearing full. Use integer values. For example, a value of 70 means 70 percent. For best operation, the value for this parameter must be greater than the value for the <code>warningThreshold</code> parameter.</p> <p>Valid values are from 0 to 100.</p> <p>The default value is 85 percent.</p> <p>Setting this parameter to 0 disables both warning alerts and critical alerts.</p> <p>If you set this to <code>default</code>, the critical alert threshold value is determined by the controller firmware.</p>
criticalPriority	<p>The priority for reconstruction operations for critical events on the disk pool. For example, disk pool reconstruction after at least two drive failures.</p> <p>Valid values are <code>highest</code>, <code>high</code>, <code>medium</code>, <code>low</code>, and <code>lowest</code>. The default value is <code>highest</code>.</p>
degradedPriority	<p>The priority for reconstruction operations for degraded events on the disk pool. For example, disk pool reconstruction after at one drive failure.</p> <p>Valid values are <code>highest</code>, <code>high</code>, <code>medium</code>, <code>low</code>, and <code>lowest</code>. The default value is <code>high</code>.</p>
backgroundPriority	<p>The priority for background operations on the disk pool.</p> <p>Valid values are <code>highest</code>, <code>high</code>, <code>medium</code>, <code>low</code>, and <code>lowest</code>. The default value is <code>low</code>.</p>
userLabel	<p>The new alphanumeric identifier (including - and _) that you want to give the disk pool. Enclose the disk pool identifier in double quotation marks (" ").</p>

Notes

Each disk pool name must be unique. You can use any combination of alphanumeric characters, underscore (`_`), hyphen (`-`), and pound (`#`) for the user label. User labels can have a maximum of 30 characters.

You can specify an arbitrary set of disk pools. If you select multiple disk pools, setting a value for the `userLabel` causes an error.

If you do not specify a value for an optional parameter, a default value is assigned.

Disk Pool Alert Thresholds

Each disk pool has two progressively severe levels of alerts to inform users when the storage capacity of the disk pool is approaching full. The threshold for an alert is a percent of the used capacity to the total usable capacity in the disk pool. The alerts are:

- Warning – This is the first level of alert that the used capacity in a disk pool is approaching full. When the threshold for the warning alert is reached a Needs Attention condition is generated and an event is posted to the storage management software. The warning threshold is superseded by the critical threshold. The default warning threshold is 50 percent.
- Critical – This is the most severe level of alert that the used capacity in a disk pool is approaching full. When the threshold for the critical alert is reached a Needs Attention condition is generated and an event is posted to the storage management software. The warning threshold is superseded by the critical threshold. The default threshold for the critical alert is 85 percent.

To be effective, the value for a warning alert must always be less than the value for a critical alert. If the value for the warning alert is the same as the value for a critical alert, only the critical alert is sent.

Disk Pool Background Operations

Disk pools support these background operations:

- Reconstruction
- Instant Availability Format (IAF)
- Format
- Dynamic Capacity Expansion (DCE)
- Dynamic Volume Expansion (DVE) (For disk pools, DVE is actually not a background operation, but DVE is supported as a synchronous operation.)

Disk pools do not queue background commands. You can start several background commands sequentially, but starting more than one background operation at a time delays the completion of commands that you started previously. The relative priority levels for the supported background operations are:

1. Reconstruction
2. Format

3. IAF
4. DCE

Minimum Firmware Level

7.83

**Set Disk Pool
(Modify Disk Pool)**

This command adds capacity to a disk pool or changes the controller ownership for the entire disk pool. These two operations are mutually exclusive.

Syntax

```
set diskPool [diskPoolName]
((addDrives=[trayID1,drawerID1,slotID1 ...
trayIDn,drawerIDn,slotIDn] |
addCapacity=(diskPoolCapacity)) | owner=(a | b))
```

Parameters

Parameter	Description
diskPool	The name of the disk pool that you want to modify. Enclose the disk pool name in square brackets ([]).
addDrives	The drives that you want to add to the disk pool. For high-capacity drive trays, specify the tray ID value, the drawer ID value, and the slot ID value of the drive that you want to add. For low-capacity drive trays, specify the tray ID value and the slot ID value of the drive that you want to add. Tray ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the tray ID value, the drawer ID value, and the slot ID value in square brackets ([]).
addCapacity	The amount of additional storage capacity that you want to add to the disk pool. This parameter automatically selects the drives to meet the capacity that you want to add. The capacity is defined in units of bytes, KB, MB, GB, or TB.
owner	The controller that owns the disk pool. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. If you do not specify an owner, the controller firmware determines the owner.

Notes

Volumes already in the disk pool remain on line and available for I/O operations while you add new drives. The disk pool must be in the Complete state before you add capacity. If the disk pool is not in the Complete state, run the `set diskPool complete` command before you add new drives.

To add capacity, specify individual drives with the `addDrives` parameter, or an amount of drive capacity with the `addCapacity` parameter. If you use `addDrives`, the host must validate the drive set before allowing the operation to run. If you use the `addCapacity` parameter, the capacity you specify is taken as the minimum capacity to be added. The candidate drives with the best match for quality of service and a capacity greater than or equal to what you specified are used. If no candidate is available with a minimum match, or the drive list specified is not available or attribute mismatches are detected, the operation fails.

You also can use this command to change ownership of a disk pool from one controller in the storage array to the other. Using this command to change ownership is mutually exclusive with using the command to add drives or to add capacity.

Minimum Firmware Level

7.83

Set Drawer Service Action Allowed Indicator

This command turns on or turns off the Service Action Allowed indicator light on a drawer that holds drives. Drawers are used in high-capacity drive trays. The drawers slide out of the drive tray to provide access to the drives. Use this command only for drive trays that use drawers. If the storage array does not support the Service Action Allowed indicator light feature, this command returns an error. If the storage array supports the command but is unable to turn on or turn off the indicator light, this command returns an error.

Syntax

```
set tray [trayID] drawer [drawerID]
serviceAllowedIndicator=(on | off | forceOnWarning)
```

Parameters

Parameter	Description
tray	The tray where the drawer resides. Tray ID values are 0 to 99. Enclose the tray ID value in square brackets ([]). If you do not enter a tray ID value, the tray ID of the controller tray is the default value.
drawer	The location of the drawer for which you want to turn on or turn off the Service Action Allowed Indicator light. Drawer ID values are 1 to 5. Enclose the drawer ID value in square brackets ([]).

Parameter	Description
serviceAllowedIndicator	<p>The setting to turn on or turn off the Service Action Allowed indicator light. To turn on the Service Action Allowed indicator light, set this parameter to <code>on</code>. To turn off the Service Action Allowed indicator light, set this parameter to <code>off</code>.</p> <p>For information about using <code>forceOnWarning</code>, see the Notes.</p>

Notes

Before you can enter this command, the drive tray must meet these conditions:

- The drive tray cannot be over temperature.
- The fans must have a status of Optimal.
- All drive tray components must be in place.
- The volumes in the drive drawer cannot be in a Degraded state. If you remove drives from the drive drawer and a volume is already in a Degraded state, the volume can fail.

ATTENTION Do not issue this command if you cannot meet any of these conditions.

All volumes with drives in the affected drive drawer are checked to make sure that the volumes have drawer loss protection before the command is sent. If the volumes have drawer loss protection, the Set Service Action Allowed command proceeds without stopping I/O activity to the volume.

If any volumes in the affected drawer do not have drawer loss protection, you must stop I/O activity to those volumes. A warning appears, which indicates that this command should not be completed.

If you are preparing a component for removal and want to override the warning that the volumes do not have drawer loss protection, enter this parameter:

```
serviceAllowedIndicator=forceOnWarning
```

`forceOnWarning` sends the request to prepare to remove a component to the controller firmware, and forces the `set drawer serviceAllowedIndicator` command to proceed.

To turn on or turn off the Service Action Allowed indicator light for the entire high-capacity drive tray, use the `set tray serviceAllowedIndicator` command.

Minimum Firmware Level

7.60

Set Drive Channel Status

This command defines how the drive channel performs.

Syntax

```
set driveChannel [(1 | 2 | 3 | 4 | 5 | 6 | 7 | 8)]
status=(optimal | degraded)
```

Parameters

Parameter	Description
driveChannel	The identifier number of the drive channel for which you want to set the status. Valid drive channel values are 1, 2, 3, 4, 5, 6, 7, or 8. Enclose the drive channel number in square brackets ([]).
status	The condition of the drive channel. You can set the drive channel status to <i>optimal</i> or <i>degraded</i> .

Notes

Use the *optimal* option to move a degraded drive channel back to the *Optimal* state. Use the *degraded* option when the drive channel is experiencing problems, and the storage array requires additional time for data transfers.

Minimum Firmware Level

6.10

7.15 adds the update to the drive channel identifier.

Set Drive Hot Spare

This command assigns or unassigns one or more drives as a hot spare.

Syntax

```
set (drive [trayID,drawerID,slotID] |
drives [trayID1,drawerID1,slotID1 ...
trayIDn,drawerIDn,slotIDn])
hotSpare=(TRUE | FALSE)
```

Parameters

Parameter	Description
drive or drives	The location of the drive that you want to use for a hot spare. For high-capacity drive trays, specify the tray ID value, the drawer ID value, and the slot ID value for the drive. For low-capacity drive trays, specify the tray ID value and the slot ID value for the drive. Tray ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the tray ID value, the drawer ID value, and the slot ID value in square brackets ([]).

Parameter	Description
hotSpare	The setting to assign the drive as the hot spare. To assign the drive as the hot spare, set this parameter to <code>TRUE</code> . To remove a hot spare assignment from a drive, set this parameter to <code>FALSE</code> .

Notes

The `drive` parameter supports both high-capacity drive trays and low-capacity drive trays. A high-capacity drive tray has drawers that hold the drives. The drawers slide out of the drive tray to provide access to the drives. A low-capacity drive tray does not have drawers. For a high-capacity drive tray, you must specify the identifier (ID) of the drive tray, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive tray, you need only specify the ID of the drive tray and the ID of the slot in which a drive resides. For a low-capacity drive tray, an alternative method for identifying a location for a drive is to specify the ID of the drive tray, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

Minimum Firmware Level

6.10

7.60 adds the `drawerID` user input.

Set Drive Service Action Allowed Indicator

This command turns on or turns off the Service Action Allowed indicator light on a drive in drive trays that support the Service Action Allowed indicator light feature. If the storage array does not support the Service Action Allowed indicator light feature, this command returns an error. If the storage array supports the command but is unable to turn on or turn off the indicator light, this command returns an error.

Syntax

```
set (drive [trayID,drawerID,slotID] |
drives [trayID1,drawerID1,slotID1 ...
trayIDn,drawerIDn,slotIDn])
serviceAllowedIndicator=(on | off)
```

Parameters

Parameter	Description
<code>drive</code> or <code>drives</code>	The location of the drive that you want to turn on or turn off the service action allowed indicator. For high-capacity drive trays, specify the tray ID value, the drawer ID value, and the slot ID value for the drive. For low-capacity drive trays, specify the tray ID value and the slot ID value for the drive. Tray ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the tray ID value, the drawer ID value, and the slot ID value in square brackets ([]).

Parameter	Description
serviceAllowedIndicator	The setting to turn on or turn off the Service Action Allowed indicator light. To turn on the Service Action Allowed indicator light, set this parameter to <code>on</code> . To turn off the Service Action Allowed indicator light, set this parameter to <code>off</code> .

Notes

The `drive` parameter supports both high-capacity drive trays and low-capacity drive trays. A high-capacity drive tray has drawers that hold the drives. The drawers slide out of the drive tray to provide access to the drives. A low-capacity drive tray does not have drawers. For a high-capacity drive tray, you must specify the identifier (ID) of the drive tray, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive tray, you need only specify the ID of the drive tray and the ID of the slot in which a drive resides. For a low-capacity drive tray, an alternative method for identifying a location for a drive is to specify the ID of the drive tray, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

Minimum Firmware Level

6.16

7.60 adds the `drawerID` user input.

Set Drive State

This command sets a drive to the Failed state. (To return a drive to the Optimal state, use the `revive drive` command.)

Syntax

```
set drive [trayID,drawerID,slotID]
operationalState=failed
```

Parameter

Parameter	Description
<code>drive</code>	The location of the drive that you want to set to the Failed state. For high-capacity drive trays, specify the tray ID value, the drawer ID value, and the slot ID value for the drive. For low-capacity drive trays, specify the tray ID value and the slot ID value for the drive. Tray ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the tray ID value, the drawer ID value, and the slot ID value in square brackets ([]).

Notes

The `drive` parameter supports both high-capacity drive trays and low-capacity drive trays. A high-capacity drive tray has drawers that hold the drives. The drawers slide out of the drive tray to provide access to the drives. A low-capacity drive tray does not

have drawers. For a high-capacity drive tray, you must specify the identifier (ID) of the drive tray, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive tray, you need only specify the ID of the drive tray and the ID of the slot in which a drive resides. For a low-capacity drive tray, an alternative method for identifying a location for a drive is to specify the ID of the drive tray, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

Minimum Firmware Level

5.20

7.60 adds the `drawerID` user input.

Set Foreign Drive to Native

A drive is considered to be native when it is a part of a volume group in a storage array. A drive is considered to be foreign when it does not belong to a volume group in a storage array or when it fails to be imported with the drives of a volume group that are transferred to a new storage array. The latter failure creates an incomplete volume group on the new storage array.

Run this command to add the missing (foreign) drives back into their original volume group and to make them part of the volume group in the new storage array.

Use this operation for emergency recovery only: when one or more drives need to be changed from a foreign drive status and returned to a native status within their original volume group.

ATTENTION Possible data corruption or data loss – Using this command for reasons other than what is stated previously might result in data loss without notification.

Syntax

```
set (drive [trayID,drawerID,slotID] |
drives [trayID1,drawerID1,slotID1
... trayIDn,drawerIDn,slotIDn] |
allDrives) nativeState
```

Parameters

Parameter	Description
<code>drive</code> or <code>drives</code>	The location of the foreign drive that you want to add to the volume group in a storage array. For high-capacity drive trays, specify the tray ID value, the drawer ID value, and the slot ID value for the drive. For low-capacity drive trays, specify the tray ID value and the slot ID value for the drive. Tray ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the tray ID value, the drawer ID value, and the slot ID value in square brackets ([]).
<code>allDrives</code>	The setting to select all of the drives.

Notes

The `drive` parameter supports both high-capacity drive trays and low-capacity drive trays. A high-capacity drive tray has drawers that hold the drives. The drawers slide out of the drive tray to provide access to the drives. A low-capacity drive tray does not have drawers. For a high-capacity drive tray, you must specify the identifier (ID) of the drive tray, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive tray, you need only specify the ID of the drive tray and the ID of the slot in which a drive resides. For a low-capacity drive tray, an alternative method for identifying a location for a drive is to specify the ID of the drive tray, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

Minimum Firmware Level

7.10

7.60 adds the `drawerID` user input.

Set Host

This command assigns a host to a host group or moves a host to a different host group. You can also create a new host group and assign the host to the new host group with this command. The actions performed by this command depend on whether the host has individual mappings or does not have individual mappings.

Syntax

```
set host [hostName]
hostGroup=("hostGroupName" | none | defaultGroup)
userLabel="newHostName"
hostType=(hostTypeIndexLabel | hostTypeIndexNumber)
```

Parameters

Parameter	Description
<code>host</code>	The name of the host that you want to assign to a host group. Enclose the host name in square brackets ([]). If the host name has special characters, you also must enclose the host name in double quotation marks (" ").
<code>hostGroup</code>	The name of the host group to which you want to assign the host. (The following table defines how the command runs if the host does or does not have individual mappings.) Enclose the host group name in double quotation marks (" "). The <code>defaultGroup</code> option is the host group that contains the host to which the volume is mapped.
<code>userLabel</code>	The new host name. Enclose the host name in double quotation marks (" ").
<code>hostType</code>	The index label or number of the host type for the host port. Use the <code>show storageArray hostTypeTable</code> command to generate a list of available host type identifiers. If the host type has special characters, enclose the host type in double quotation marks (" ").

Host Group Parameter	Host Has Individual Mappings	Host Does Not Have Individual Mappings
<i>hostGroupName</i>	The host is removed from the present host group and is placed under the new host group defined by <i>hostGroupName</i> .	The host is removed from the present host group and is placed under the new host group defined by <i>hostGroupName</i> .
none	The host is removed from the host group as an independent partition and is placed under the root node.	The host is removed from the present host group and is placed under the default group.
defaultGroup	The command fails.	The host is removed from the present host group and is placed under the default group.

Notes

When you use this command, you can specify one or more of the optional parameters.

For the names, you can use any combination of alphanumeric characters, hyphens, and underscores. Names can have a maximum of 30 characters.

Minimum Firmware Level

6.10

Set Host Channel

This command defines the loop ID for the host channel.

Syntax

```
set hostChannel [hostChannelNumber]
preferredID=portID
```

Parameters

Parameter	Description
hostChannel	The identifier number of the host channel for which you want to set the loop ID. Enclose the host channel identifier number in square brackets ([]). Use a host channel value that is appropriate for your particular controller model. A controller tray might support one host channel or as many as eight host channels. Valid host channel values are a1, a2, a3, a4, a5, a6, a7, a8, b1, b2, b3, b4, b5, b6, b7, or b8.
preferredID	The port identifier for the specified host channel. Port ID values are 0 to 127.

Minimum Firmware Level

6.10

6.14 adds an update to the host channel identifier.

7.15 adds an update to the host channel identifier.

Set Host Group

This command renames a host group.

Syntax

```
set hostGroup [hostGroupName]  
userLabel="newHostGroupName"
```

Parameters

Parameter	Description
hostGroup	The name of the host group that you want to rename. Enclose the host group name in square brackets ([]). If the host group name has special characters, you also must enclose the host group name in double quotation marks (" ").
userLabel	The new name for the host group. Enclose the new host group name in double quotation marks (" ").

Notes

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

Minimum Firmware Level

6.10

Set Host Port

This command changes the host type for a host port. You can also change a host port label with this command.

Syntax

```
set hostPort [portLabel] host="hostName"  
userLabel="newPortLabel"
```

Parameters

Parameter	Description
hostPort	The name of the host port for which you want to change the host type, or for which you want to create a new name. Enclose the host port name in square brackets ([]). If the host port label has special characters, enclose the host port label in double quotation marks (" ").

Parameter	Description
host	The name of the host to which the host port is connected. Enclose the host name in double quotation marks (" ").
userLabel	The new name that you want to give to the host port. Enclose the new name of the host port in double quotation marks (" ").

Notes

When you use this command, you can specify one or more of the optional parameters.

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

Minimum Firmware Level

6.10

Set iSCSI Initiator

This command sets the attributes for an iSCSI initiator.

Syntax

```
set iscsiInitiator (["iscsiID"] |
userLabel="newName" |
host="newHostName" |
chapSecret="newSecurityKey")
```

Parameters

Parameter	Description
iscsiInitiator	The name of the iSCSI initiator for which you want to set attributes. Enclose the iSCSI initiator name in double quotation marks (" ") and square brackets ([]).
userLabel	The new name that you want to use for the iSCSI initiator. Enclose the new name in double quotation marks (" ").
host	The name of the new host to which the host port is connected. Enclose the host name in double quotation marks (" ").
chapSecret	The security key that you want to use to authenticate a peer connection. Enclose the security key in double quotation marks (" ").

Notes

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

Challenge Handshake Authentication Protocol (CHAP) is a protocol that authenticates the peer of a connection. CHAP is based upon the peers sharing a *secret*. A secret is a security key that is similar to a password.

Use the `chapSecret` parameter to set up the security keys for initiators that require a mutual authentication. The CHAP secret must be between 12 characters and 57 characters. This table lists the valid characters.

Space	!	"	#	\$	%	&	'	()	*	+
,	-	.	/	0	1	2	3	4	5	6	7
8	9	:	;	<	=	>	?	@	A	B	C
D	E	F	G	H	I	J	K	L	M	N	O
P	Q	R	S	T	U	V	W	X	Y	Z	[
\]	^	_	'	a	b	c	d	e	f	g
h	i	j	k	l	m	n	o	p	q	r	s
t	u	v	w	x	y	z	{		}	~	

Minimum Firmware Level

7.10

Set iSCSI Target Properties

This command defines properties for an iSCSI target.

Syntax

```
set iscsiTarget ["userLabel"]
authenticationMethod=(none | chap) |
chapSecret=securityKey |
targetAlias="userLabel"
```

Parameters

Parameter	Description
<code>iscsiTarget</code>	The iSCSI target for which you want to set properties. Enclose the <code>userLabel</code> in double quotation marks (" "). You must also enclose the <code>userLabel</code> in either square brackets ([]) or angle brackets (< >).
<code>authenticationMethod</code>	The means of authenticating your iSCSI session.
<code>chapSecret</code>	The security key that you want to use to authenticate a peer connection.
<code>targetAlias</code>	The new name that you want to use for the target. Enclose the name in double quotation marks (" ").

Notes

Challenge Handshake Authentication Protocol (CHAP) is a protocol that authenticates the peer of a connection. CHAP is based upon the peers sharing a *secret*. A secret is a security key that is similar to a password.

Use the `chapSecret` parameter to set up the security keys for initiators that require a mutual authentication. The CHAP secret must be between 12 characters and 57 characters. This table lists the valid characters.

Space	!	"	#	\$	%	&	'	()	*	+
,	-	.	/	0	1	2	3	4	5	6	7
8	9	:	;	<	=	>	?	@	A	B	C
D	E	F	G	H	I	J	K	L	M	N	O
P	Q	R	S	T	U	V	W	X	Y	Z	[
\]	^	_	'	a	b	c	d	e	f	g
h	i	j	k	l	m	n	o	p	q	r	s
t	u	v	w	x	y	z	{		}	~	

Minimum Firmware Level

7.10

Set Read-Only Snapshot Volume to Read/Write Volume

This command changes a snapshot volume that is a read-only volume to a snapshot volume that is read/write volume. You also can use this command to identify a new repository volume for the read/write volume, or to set a full limit warning level for the repository volume.

Syntax

```
set snapVolume ["snapImageVolumeName"] convertToReadWrite  
[(repositoryVolume="repos_xxxx" |  
repositoryVolume=(volumeGroupName [capacity=capacityValue])  
repositoryVolume=(diskPoolName [capacity=capacityValue]))  
repositoryFullLimit=percentValue]
```

Parameters

Parameter	Description
<code>snapVolume</code>	The name of the snapshot volume that you want to change from read-only to read/write. Enclose the snapshot volume identifier in double quotation marks (" ") inside of square brackets ([]).

Parameter	Description
repositoryVolume	<p>The name of the repository volume that you want to use for the read/write volume.</p> <p>You have two options for defining the name of a repository volume:</p> <ul style="list-style-type: none"> ■ Use an existing repository volume: name ■ Create a new repository volume when you run this command <p>The name of an existing repository volume is comprised of two parts:</p> <ul style="list-style-type: none"> ■ The term <i>repos</i> ■ A four digit numerical identifier that the storage management software assigns to the repository volume name <p>Enclose the name of the existing repository volume in double quotation marks (" ").</p> <p>If you want to create a new repository volume when you run this command you must enter the name of either a volume group or a disk pool in which you want the repository volume. Optionally, you can also define the capacity of the repository volume. If you want to define the capacity you can use these values:</p> <ul style="list-style-type: none"> ■ An integer value that represents a percentage of the base volume capacity ■ A decimal fraction value that represents a percentage of the base volume capacity ■ A specific size for the repository volume. Size is defined in units of bytes, KB, MB, GB, or TB. <p>If you do not use the capacity option, the storage management software sets the capacity to 20 percent of the base volume capacity.</p> <p>When you run this command the storage management software creates the repository volume for the snapshot volume.</p>
repositoryFullLimit	<p>The percentage of repository volume capacity at which you receive a warning that the snapshot image repository volume is nearing full. Use integer values. For example, a value of 70 means 70 percent. The default value is 75.</p>

Notes

You can use any combination of alphanumeric characters, underscore (_), hyphen(-), and pound (#) for the user label. User labels can have a maximum of 30 characters.

The repository volume name is automatically created by the storage management software and the firmware when you create a new snapshot group. You cannot rename the repository volume because renaming the repository volume will break the linkage with the snapshot images.

A snapshot group repository volume is an expandable volume that is structured as a concatenated collection of up to 16 standard volume entities. Initially, an expandable repository volume has only a single element. The capacity of the expandable repository volume is exactly that of the single element. You can increase the capacity of an expandable repository volume by attaching additional standard volumes to it. The composite expandable repository volume capacity then becomes the sum of the capacities of all of the concatenated standard volumes.

A snapshot group repository volume must satisfy a minimum capacity requirement that is the sum of the following:

- 32 MB to support fixed overhead for the snapshot group and for copy-on-write processing.
- Capacity for rollback processing, which is 1/5000th of the capacity of the base volume.

The minimum capacity is enforcement by the controller firmware and the storage management software.

Minimum Firmware Level

7.83

Set Session

This command defines how you want the current script engine session to run. For the purpose of this command a session is the duration for the running of the commands. This command does not permanently set the parameters for the storage array.

Syntax

```
set session errorAction=(stop | continue)
password="storageArrayPassword"
userRole=(admin | monitor)
performanceMonitorInterval=intervalValue
performanceMonitorIterations=iterationValue
```

Parameters

Parameter	Description
<code>errorAction</code>	How the session responds if an error is encountered during processing. You can choose to stop the session if an error is encountered, or you can continue the session after encountering an error. The default value is <code>stop</code> . (This parameter defines the action for execution errors, not syntax errors. Some error conditions might override the <code>continue</code> value.)
<code>password</code>	The password for the storage array. Enclose the password in double quotation marks (" ").
<code>userRole</code>	<p>Defines the user role for the password. The user role can be either:</p> <ul style="list-style-type: none"> ■ <code>admin</code> – Enables you to view and modify storage array configurations ■ <code>monitor</code> – Enables you to view storage array configurations and monitor storage array health conditions, but not modify the configuration <p>The <code>userRole</code> parameter is optional. If you do not use the <code>userRole</code> parameter, the storage array uses any password that you define as the administrator password.</p>
<code>performanceMonitorInterval</code>	The frequency of gathering performance data. Enter an integer value for the polling interval, in seconds, for which you want to capture data. The range of values is 3 to 3600 seconds. The default value is 5 seconds.
<code>performanceMonitorIterations</code>	The number of samples to capture. Enter an integer value. The range of values for samples captured is 1 to 3600. The default value is 5.

Notes

When you use this command, you can specify one or more of the optional parameters.

Passwords are stored on each storage array in a management domain. If a password was not previously set, you do not need a password. The password can be any combination of alphanumeric characters with a maximum of 30 characters. (You can define a storage array password by using the `set storageArray` command.)

The polling interval and the number of iterations that you specify remain in effect until you end the session. After you end the session, the polling interval and the number of iterations return to the default values.

Minimum Firmware Level

5.20

Set Snapshot (Legacy) Volume

This command defines the properties for a snapshot (legacy) volume and lets you rename a snapshot (legacy) volume.

Syntax

```
set (volume [volumeName] | volumes ["volumeName1" ...
"volumeNameN" ])
userLabel="snapshotVolumeName"
warningThresholdPercent=percentValue
repositoryFullPolicy=(failBaseWrites | failSnapshot)
enableSchedule=(TRUE | FALSE)
schedule (immediate | snapshotSchedule)
rollbackPriority=(0 | 1 | 2 | 3 | 4)
```

Parameters

Parameter	Description
volume or volumes	The name of the snapshot (legacy) volume for which you want to define properties. Enclose the snapshot (legacy) volume name in double quotation marks (" ") inside of square brackets ([]). You can enter more than one snapshot (legacy) volume name. Enclose all of the snapshot (legacy) volume names in one set of square brackets ([]). Enclose each volume name in double quotation marks (" "). Separate each volume name with a white space.
userLabel	A new name that you want to give to a snapshot (legacy) volume. Enclose the new snapshot (legacy) volume name in double quotation marks (" ").
warningThresholdPercent	The percentage of repository capacity at which you receive a warning that the snapshot (legacy) repository is nearing full. Use integer values. For example, a value of 70 means 70 percent. The default value is 50.

Parameter	Description
repositoryFullPolicy	How you want snapshot (legacy) processing to continue if the snapshot (legacy) repository volume is full. You can choose to fail writes to the base volume (<code>failBaseWrites</code>) or fail writes to the snapshot (legacy) volume (<code>failSnapshot</code>). The default value is <code>failSnapshot</code> .
enableSchedule	Use this parameter to turn on or to turn off the ability to schedule a snapshot (legacy) operation. To turn on snapshot (legacy) scheduling, set this parameter to <code>TRUE</code> . To turn off snapshot (legacy) scheduling, set this parameter to <code>FALSE</code> .
schedule	<p>Use this parameter to schedule a snapshot operation.</p> <p>You can use one of these options for setting a schedule for a snapshot operation:</p> <ul style="list-style-type: none"> ■ <code>immediate</code> ■ <code>startDate</code> ■ <code>scheduleDay</code> ■ <code>startTime</code> ■ <code>scheduleInterval</code> ■ <code>endDate</code> ■ <code>timesPerDay</code> ■ <code>timeZone</code> ■ <code>scheduleDate</code> ■ <code>month</code> <p>See the "Notes" section for information explaining how to use these options.</p>
rollbackPriority	Use this parameter to determine whether system resources should be allocated to the rollback operation at the expense of system performance. A value of 0 indicates that the rollback operation is prioritized over all other host I/O. A value of 4 indicates that the rollback operation should be performed with minimal impact to host I/O.

Notes

When you use this command, you can specify one or more of the optional parameters.

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

You can set the `warningThresholdPercent` parameter and the `repositoryFullPolicy` parameter for both the snapshot (legacy) repository volume or the snapshot (legacy) volume.

Scheduling Snapshots (Legacy)

The `enableSchedule` parameter and the `schedule` parameter provide a way for you to schedule automatic snapshots (legacy). Using these parameters, you can schedule snapshots (legacy) daily, weekly, or monthly (by day or by date). The `enableSchedule` parameter turns on or turns off the ability to schedule snapshots (legacy). When you enable scheduling, you use the `schedule` parameter to define when you want the snapshots (legacy) to occur.

This list explains how to use the options for the `schedule` parameter:

- `immediate` – As soon as you enter the command, a snapshot (legacy) volume is created, and a copy-on-write operation begins.
- `startDate` – A specific date on which you want to create a snapshot (legacy) volume and perform a copy-on-write operation. The format for entering the date is `MM:DD:YY`. If you do not provide a start date, the current date is used. An example of this option is `startDate=06:27:11`.
- `scheduleDay` – A day of the week on which you want to create a snapshot (legacy) volume and perform a copy-on-write operation. You can enter these values: `monday`, `tuesday`, `wednesday`, `thursday`, `friday`, `saturday`, `sunday`, and `all`. An example of this option is `scheduleDay=wednesday`.
- `startTime` – The time of a day that you want to create a snapshot (legacy) volume and start performing a copy-on-write operation. The format for entering the time is `HH:MM`, where `HH` is the hour and `MM` is the minute past the hour. Use a 24-hour clock. For example, 2:00 in the afternoon is 14:00. An example of this option is `startTime=14:27`.
- `scheduleInterval` – An amount of time, in minutes, that you want to have as a minimum between copy-on-write operation. You can create a schedule in which you have overlapping copy-on-write operations because of the duration a copy operation. You can make sure that you have time between copy-on-write operations by using this option. The maximum value for the `scheduleInterval` option is 1440 minutes. An example of this option is `scheduleInterval=180`.
- `endDate` – A specific date on which you want to stop creating a snapshot (legacy) volume and end the copy-on-write operation. The format for entering the date is `MM:DD:YY`. An example of this option is `endDate=11:26:11`.

- `timesPerDay` – The number of times that you want the schedule to run in a day. An example of this option is `timesPerDay=4`.
- `timeZone` – Use this parameter to define the time zone in which the storage array is operating. You can define the time zone in one of two ways:
 - `GMT±HH:MM` – The time zone offset from GMT. Enter the offset in hours and minutes. For example `GMT-06:00` is the central time zone in the United States.
 - `Text string` – Standard time zone text strings. For example: `"USA/Chicago"` or `"Germany/Berlin"`. Time zone text strings are case sensitive. If you enter an incorrect text string, GMT time is used. Enclose the text string in double quotation marks.
- `scheduleDate` – A day of the month on which you want to create a snapshot volume and perform a copy-on-write operation. The values for the days are numerical and in the range of 1–31. Enclose the value for the day in double quotation marks inside parentheses. If you want to create a snapshot volume on a specific date of a specific month use the `scheduleDate` option with the `month` option. If you have set up a weekly schedule, you cannot use the `scheduleDate` option. An example of the `scheduleDate` option is `scheduleDate=("15")`.
- `month` – A specific month on which you want to create a snapshot volume and perform a copy-on-write operation. The values for the months are: `jan`, `feb`, `mar`, `apr`, `may`, `jun`, `jul`, `aug`, `sep`, `oct`, `nov`, and `dec`. Enclose the value in parentheses. You can enter more than one month by enclosing the months in a single set of parentheses and separating each month with a white space. If you want to create a snapshot volume on a specific date of a specific month use the `scheduleDate` option with the `month` option. These are examples of the `month` option:
 - `month=(mar)`
 - `month=(mar apr may)`

The code string for defining a schedule is similar to these examples:

```
enableSchedule=true schedule startTime=14:27
enableSchedule=true schedule scheduleInterval=180
enableSchedule=true schedule timeZone=GMT-06:00
enableSchedule=true schedule timeZone="USA/Chicago"
enableSchedule=true schedule month=(mar) scheduleDate=( "15" )
```

If you also use the `scheduleInterval` option, the firmware chooses between the `timesPerDay` option and the `scheduleInterval` option by selecting the lowest value of the two options. The firmware calculates an integer value for the `scheduleInterval` option by dividing 1440 by the `scheduleInterval`

option value that you set. For example, $1440/180 = 8$. The firmware then compares the `timesPerDay` integer value with the calculated `scheduleInterval` integer value and uses the smaller value.

To remove a schedule, use the `delete volume` command with the `schedule` parameter. The `delete volume` command with the `schedule` parameter deletes only the schedule, not the snapshot (legacy) volume.

Minimum Firmware Level

6.10

7.77 adds scheduling.

7.80 adds the `rollback` parameter.

7.83 removes the `noEndDate` option. This option is not supported.

7.86 adds the `scheduleDate` option and the `month` option.

Set Snapshot Group Attributes

This command defines the properties for a snapshot group.

Syntax

```
set snapGroup ["snapGroupName"]
[userLabel="snapGroupName" |
repositoryFullPolicy=(failBaseWrites | purgeSnapImages) |
repositoryFullLimit=percentValue |
autoDeleteLimit=numberOfSnapImages |
rollbackPriority=(lowest | low | medium | high | highest)]
```

Parameters

Parameter	Description
<code>snapGroup</code>	The name of the snapshot group for which you are setting properties. Enclose the snapshot group name in double quotation marks (" ") inside of square brackets ([]).
<code>userLabel</code>	The new name that you want to give to the snapshot group. Use this parameter when you want to rename the snapshot group. Enclose the new snapshot group name in double quotation marks (" ").
<code>repositoryFullPolicy</code>	How you want snapshot image processing to continue if the repository volumes for the snapshot image group are full. You can choose to fail writes to the base volume (<code>failBaseWrites</code>) or delete (<code>purge</code>) the snapshot images (<code>purgeSnapImages</code>). The default action is <code>purgeSnapImages</code> .

Parameter	Description
repositoryFullLimit	The percentage of repository volume capacity at which you receive a warning that the snapshot image repository volume is nearing full. Use integer values. For example, a value of 70 means 70 percent. The default value is 75.
autoDeleteLimit	The maximum number of snapshot images that you want to automatically delete if you have selected to purge the snapshot images for a repository full policy. Use integer values. The default value is 32.
rollbackPriority	Use this parameter to determine whether system resources should be allocated to the rollback operation at the expense of system performance. Valid values are <code>highest</code> , <code>high</code> , <code>medium</code> , <code>low</code> , or <code>lowest</code> . A value of <code>high</code> indicates that the rollback operation is prioritized over all other host I/O. A value of <code>lowest</code> indicates that the rollback operation should be performed with minimal impact to host I/O.

Notes

You can use any combination of alphanumeric characters, underscore (`_`), hyphen (`-`), and pound (`#`) for the names. Names can have a maximum of 30 characters.

When you use this command, you can specify one or more of the parameters. You do not, however, need to use all of the parameters.

Minimum Firmware Level

7.83

Set Snapshot Group Media Scan

This command runs a media scan on a snapshot group.

Syntax

```
set snapGroup ["snapGroupName"]
mediaScanEnabled=(TRUE | FALSE)
redundancyCheckEnabled=(TRUE | FALSE)
```

Parameters

Parameter	Description
snapGroup	The name of the snapshot group on which you want to run a media scan. Enclose the snapshot group name in double quotation marks (" ") inside of square brackets ([]).

Parameter	Description
mediaScanEnabled	The setting to turn on or turn off media scan for the volume. To turn on media scan, set this parameter to TRUE. To turn off media scan, set this parameter to FALSE. (If media scan is disabled at the storage array level, this parameter has no effect.)
redundancyCheckEnabled	The setting to turn on or turn off redundancy checking during a media scan. To turn on redundancy checking, set this parameter to TRUE. To turn off redundancy checking, set this parameter to FALSE.

Notes

You can use any combination of alphanumeric characters, underscore (_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

Minimum Firmware Level

7.83

Set Snapshot Group Repository Volume Capacity

This command increases or decreases the capacity of a snapshot group repository volume.

Syntax for Increasing Capacity

```
set snapGroup ["snapGroupName"] increaseRepositoryCapacity
(repositoryVolumes="repos_xxxx" |
repositoryVolumes=((volumeGroupName
[capacity=capacityValue])) |
repositoryVolumes=((diskPoolName
[capacity=capacityValue])))
```

Syntax for Decreasing Capacity

```
set snapGroup ["snapGroupName"]
decreaseRepositoryCapacity
count=numberOfVolumes
```

Parameters

Parameter	Description
snapGroup	The name of the snapshot group on which you want to run a media scan. Enclose the snapshot group name in double quotation marks (" ") inside of square brackets ([]).

Parameter	Description
repositoryVolume	<p>The name of the repository volume for which you want to increase capacity. An available standard volume is added to the repository volume to increase the capacity of the repository volume.</p> <p>You have two options for defining the name of a repository volume:</p> <ul style="list-style-type: none"> ■ Use an existing repository volume: name ■ Create a new repository volume when you run this command <p>The name of an existing repository volume is comprised of two parts:</p> <ul style="list-style-type: none"> ■ The term <i>repos</i> ■ A four digit numerical identifier that the storage management software assigns to the repository volume name <p>Enclose the name of the existing repository volume in double quotation marks (" ").</p> <p>If you want to create a new repository volume when you run this command you must enter the name of either a volume group or a disk pool in which you want the repository volume. Optionally, you can also define the capacity of the repository volume. If you want to define the capacity you can use these values:</p> <ul style="list-style-type: none"> ■ An integer value that represents a percentage of the base volume capacity ■ A decimal fraction value that represents a percentage of the base volume capacity ■ A specific size for the repository volume. Size is defined in units of bytes, KB, MB, GB, or TB. <p>If you do not use the capacity option, the storage management software sets the capacity to 20 percent of the base volume capacity.</p> <p>When you run this command the storage management software creates the repository volume for the snapshot volume.</p>
count	<p>The number of repository volumes that you want to remove from the snapshot group. Use integer values.</p>

Notes

You can use any combination of alphanumeric characters, underscore (_), hyphen(-), and pound (#) for the user label. User labels can have a maximum of 30 characters.

The repository volume name is automatically created by the storage management software and the firmware when you create a new snapshot group. You cannot rename the repository volume because renaming the repository volume breaks the linkage with the snapshot images.

A snapshot group repository volume is an expandable volume that is structured as a concatenated collection of up to 16 standard volume entities. Initially, an expandable repository volume has only a single element. The capacity of the expandable repository volume is exactly that of the single element. You can increase the capacity of an expandable repository volume by attaching additional standard volumes to it. The composite expandable repository volume capacity then becomes the sum of the capacities of all of the concatenated standard volumes.

A snapshot group repository volume must satisfy a minimum capacity requirement that is the sum of the following:

- 32 MB to support fixed overhead for the snapshot group and for copy-on-write processing.
- Capacity for rollback processing, which is 1/5000th of the capacity of the base volume.

The minimum capacity is enforcement by the controller firmware and the storage management software.

Minimum Firmware Level

7.83

Set Snapshot Group Schedule

This command defines the the schedule for taking snapshot images for a snapshot group.

Syntax

```
set snapGroup ["snapGroupName"]
enableSchedule=(TRUE | FALSE)
schedule (immediate | snapshotSchedule)
```

Parameters

Parameter	Description
snapGroup	The name of the snapshot group for which you are setting properties. Enclose the snapshot group name in double quotation marks (" ") inside of square brackets ([]).

Parameter	Description
enableSchedule	Use this parameter to turn on or to turn off the ability to schedule a snapshot operation. To turn on snapshot scheduling, set this parameter to TRUE. To turn off snapshot scheduling, set this parameter to FALSE.
schedule	<p>Use this parameter to schedule a snapshot operation.</p> <p>You can use one of these options for setting a schedule for a snapshot operation:</p> <ul style="list-style-type: none"> ■ immediate ■ startDate ■ scheduleDay ■ startTime ■ scheduleInterval ■ endDate ■ timesPerDay ■ timeZone ■ scheduleDate ■ month <p>See the "Notes" section for information explaining how to use these options.</p>

Notes

You can use any combination of alphanumeric characters, underscore (`_`), hyphen (`-`), and pound (`#`) for the names. Names can have a maximum of 30 characters.

Scheduling Snapshots

The `enableSchedule` parameter and the `schedule` parameter provide a way for you to schedule creating snapshot images for a snapshot group. Using these parameters, you can schedule snapshots daily, weekly, or monthly (by day or by date). The `enableSchedule` parameter turns on or turns off the ability to schedule snapshots. When you enable scheduling, you use the `schedule` parameter to define when you want the snapshots to occur.

This list explains how to use the options for the `schedule` parameter:

- `immediate` – As soon as you enter the command, a snapshot image is created, and a copy-on-write operation begins.
- `startDate` – A specific date on which you want to create a snapshot image and perform a copy-on-write operation. The format for entering the date is `MM:DD:YY`. If you do not provide a start date, the current date is used. An example of this option is `startDate=06:27:11`.
- `scheduleDay` – A day of the week on which you want to create a snapshot image and perform a copy-on-write operation. You can enter these values: `monday`, `tuesday`, `wednesday`, `thursday`, `friday`, `saturday`, `sunday`, and `all`. An example of this option is `scheduleDay=wednesday`.
- `startTime` – The time of a day that you want to create a snapshot image and start performing a copy-on-write operation. The format for entering the time is `HH:MM`, where `HH` is the hour and `MM` is the minute past the hour. Use a 24-hour clock. For example, 2:00 in the afternoon is 14:00. An example of this option is `startTime=14:27`.
- `scheduleInterval` – An amount of time, in minutes, that you want to have as a minimum between copy-on-write operations. You can possibly create a schedule in which you have overlapping copy-on-write operations because of the duration of a copy operation. You can make sure that you have time between copy-on-write operations by using this option. The maximum value for the `scheduleInterval` option is 1440 minutes. An example of this option is `scheduleInterval=180`.
- `endDate` – A specific date on which you want to stop creating a snapshot image and end the copy-on-write operation. The format for entering the date is `MM:DD:YY`. An example of this option is `endDate=11:26:11`.
- `noEndDate` – Use this option if you do not want your scheduled copy-on-write operation to end. If you later decide to end the copy-on-write operations you must re-enter the `set snapGroup` command and specify an end date.
- `timesPerDay` – The number of times that you want the schedule to run in a day. An example of this option is `timesPerDay=4`.
- `timeZone` – Use this parameter to define the time zone in which the storage array is operating. You can define the time zone in one of two ways:
 - `GMT±HH:MM` – The time zone offset from GMT. Enter the offset in hours and minutes. For example `GMT-06:00` is the central time zone in the United States.
 - `Text string` – Standard time zone text strings. For example: `"USA/Chicago"` or `"Germany/Berlin"`. Time zone text strings are case sensitive. If you enter an incorrect text string, GMT time is used. Enclose the text string in double quotation marks.

- `scheduleDate` – A day of the month on which you want to create a snapshot volume and perform a copy-on-write operation. The values for the days are numerical and in the range of 1–31. Enclose the value for the day in double quotation marks inside parentheses. If you want to create a snapshot volume on a specific date of a specific month use the `scheduleDate` option with the `month` option. If you have set up a weekly schedule, you cannot use the `scheduleDate` option. An example of the `scheduleDate` option is `scheduleDate=("15")`.
- `month` – A specific month on which you want to create a snapshot volume and perform a copy-on-write operation. The values for the months are: `jan`, `feb`, `mar`, `apr`, `may`, `jun`, `jul`, `aug`, `sep`, `oct`, `nov`, and `dec`. Enclose the value in parentheses. You can enter more than one month by enclosing the months in a single set of parentheses and separating each month with a white space. If you want to create a snapshot volume on a specific date of a specific month use the `scheduleDate` option with the `month` option. These are examples of the `month` option:
 - `month=(mar)`
 - `month=(mar apr may)`

The code string for defining a schedule is similar to these examples:

```
enableSchedule=true schedule startTime=14:27
enableSchedule=true schedule scheduleInterval=180
enableSchedule=true schedule timeZone=GMT-06:00
enableSchedule=true schedule timeZone="USA/Chicago"
enableSchedule=true schedule month=(mar) scheduleDate=( "15" )
```

If you also use the `scheduleInterval` option, the firmware chooses between the `timesPerDay` option and the `scheduleInterval` option by selecting the lowest value of the two options. The firmware calculates an integer value for the `scheduleInterval` option by dividing 1440 by the `scheduleInterval` option value that you set. For example, $1440/180 = 8$. The firmware then compares the `timesPerDay` integer value with the calculated `scheduleInterval` integer value and uses the smaller value.

To remove a schedule, use the `delete volume` command with the `schedule` parameter. The `delete volume` command with the `schedule` parameter deletes only the schedule, not the snapshot volume.

Minimum Firmware Level

7.83

7.86 adds the `scheduleDate` option and the `month` option.

Set Snapshot Volume Media Scan

This command runs a media scan on the drives used for a snapshot volume. Optionally, you also can perform a redundancy check on the data.

Syntax

```
set snapVolume ["snapVolumeName"]
mediaScanEnabled=(TRUE | FALSE)
[redundancyCheckEnabled=(TRUE | FALSE)]
```

Parameters

Parameter	Description
snapVolume	The name of the snapshot volume for which you are setting properties. Enclose the snapshot volume identifier in double quotation marks (" ") inside of square brackets ([]).
mediaScanEnabled	The setting to turn on or turn off media scan for the snapshot volume. To turn on media scan, set this parameter to TRUE. To turn off media scan, set this parameter to FALSE. (If media scan is disabled at the storage array level, this parameter has no effect.)
redundancyCheckEnabled	The setting to turn on or turn off redundancy checking during a media scan. To turn on redundancy checking, set this parameter to TRUE. To turn off redundancy checking, set this parameter to FALSE.

Notes

You can use any combination of alphanumeric characters, underscore (_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

Minimum Firmware Level

7.83

Set Snapshot Volume Repository Volume Capacity

This command increases or decreases the capacity of a snapshot volume repository volume.

Syntax for Increasing Capacity

```
set snapVolume["snapVolumeName"] increaseRepositoryCapacity
repositoryVolumes=("repos_xxxx" |
repositoryVolumes=(volumeGroupName
[capacity=capacityValue]) repositoryVolumes=(diskPoolName
[capacity=capacityValue]))
```

Syntax for Decreasing Capacity

```
set snapVolume ["snapVolumeName"] decreaseRepositoryCapacity
count=numberOfVolumes
```

Parameters

Parameter	Description
snapVolume	The name of the snapshot volume for which you are setting properties. Enclose the snapshot volume identifier in double quotation marks (" ") inside of square brackets ([]).
repositoryVolume	<p>The name of the repository volume for which you want to increase capacity. An available standard volume is added to the repository volume to increase the capacity of the repository volume.</p> <p>You have two options for defining the name of a repository volume:</p> <ul style="list-style-type: none"> ■ Use an existing repository volume: name ■ Create a new repository volume when you run this command <p>The name of an existing repository volume is comprised of two parts:</p> <ul style="list-style-type: none"> ■ The term <i>repos</i> ■ A four digit numerical identifier that you assign to the repository volume name <p>Enclose the name of the existing repository volume in double quotation marks (" ").</p> <p>If you want to create a new repository volume when you run this command you must enter the name of either a volume group or a disk pool in which you want the repository volume. Optionally, you can also define the capacity of the repository volume. If you want to define the capacity you can use these values:</p> <ul style="list-style-type: none"> ■ An integer value that represents a percentage of the base volume capacity ■ A decimal fraction value that represents a percentage of the base volume capacity ■ A specific size for the repository volume. Size is defined in units of bytes, KB, MB, GB, or TB. <p>If you do not use the capacity option, the storage management software sets the capacity to 20 percent of the base volume capacity.</p> <p>Enclose the name of the new repository volume in parenthesis ().</p>
count	The number of volumes that you want to remove . Use integer values.

Notes

You can use any combination of alphanumeric characters, underscore (`_`), hyphen (`-`), and pound (`#`) for the names. Names can have a maximum of 30 characters.

A snapshot repository volume is an expandable volume that is structured as a concatenated collection of up to 16 standard volume entities. Initially, an expandable repository volume has only a single element. The capacity of the expandable repository volume is exactly that of the single element. You can increase the capacity of an expandable repository volume by attaching additional standard volumes to it. The composite expandable repository volume capacity then becomes the sum of the capacities of all of the concatenated standard volumes.

A snapshot group repository volume must satisfy a minimum capacity requirement that is the sum of the following:

- 32 MB to support fixed overhead for the snapshot group and for copy-on-write processing.
- Capacity for rollback processing, which is 1/5000th of the capacity of the base volume.

The minimum capacity is enforcement by the controller firmware and the storage management software.

Minimum Firmware Level

7.83

Set Storage Array

This command defines the properties of the storage array.

Syntax

```
set storageArray (alarm=(enable | disable | mute) |
autoSupportConfig (enable | disable) |
cacheBlockSize=cacheBlockSizeValue |
cacheFlushStart=cacheFlushStartSize |
cacheFlushStop=cacheFlushStopSize |
defaultHostType=("hostTypeName" | hostTypeIdentifier)
failoverAlertDelay=delayValue |
mediaScanRate=(disabled | 1-30) |
password="password" |
userRole=(admin | monitor)
userLabel="storageArrayName"
isnsRegistration=(TRUE | FALSE))
```

Parameters

Parameter	Description
alarm	<p>The setting for the audible alarm. This parameter has these values:</p> <ul style="list-style-type: none"> ■ enable – The audible alarm is turned on and sounds if a fault occurs. ■ disable – The audible alarm is turned off and does not sound if a fault occurs. ■ mute – The audible alarm is turned off if it is sounding. <p>If another fault occurs after you set the audible alarm to <code>mute</code>, the audible alarm sounds again.</p>
cacheBlockSize	The cache block size that is used by the controller for managing the cache. Valid values are 4 (4 KB), 8 (8 KB), 16 (16 KB), or 32 (32 KB).
cacheFlushStart	The percentage of unwritten data in the cache that causes a cache flush. Use integer values from 0 to 100 to define the percentage. The default value is 80.
cacheFlushStop	The percentage of unwritten data in the cache that stops a cache flush in progress. Use integer values from 0 to 100 to define the percentage. This value must be less than the value of the <code>cacheFlushStart</code> parameter.
defaultHostType	The default host type of any unconfigured host port to which the controllers are connected. To generate a list of valid host types for the storage array, run the <code>show storageArray hostTypeTable</code> command. Host types are identified by a name or a numerical index. Enclose the host type name in double quotation marks (" "). Do not enclose the host type numerical identifier in double quotation marks.
failoverAlertDelay	The failover alert delay time in minutes. The valid values for the delay time are 0 to 60 minutes. The default value is 5.
mediaScanRate	The number of days over which the media scan runs. Valid values are <code>disabled</code> , which turns off the media scan, or 1 day to 30 days, where 1 day is the fastest scan rate, and 30 days is the slowest scan rate. A value other than <code>disabled</code> or 1 to 30 does not allow the media scan to function.
password	The password for the storage array. Enclose the password in double quotation marks (" ").

Parameter	Description
<code>userRole</code>	<p>Defines the user role for the password. The user role can be either:</p> <ul style="list-style-type: none"> ■ <code>admin</code> – Enables you to view and modify storage array configurations ■ <code>monitor</code> – Enables you to view storage array configurations and monitor storage array health conditions, but not modify the configuration <p>The <code>userRole</code> parameter is optional. If you do not use the <code>userRole</code> parameter, the storage array uses any password that you define as the administrator password.</p>
<code>userLabel</code>	The name for the storage array. Enclose the storage array name in double quotation marks (" ").
<code>isnsRegistration</code>	The means of listing the iSCSI target on the iSNS server. Set the parameter to <code>TRUE</code> to list it.

Notes

When you use this command, you can specify one or more of the optional parameters.

Auto Support Data

NOTE You cannot use this parameter with the ASUP commands.

When enabled, the `set storageArray autoSupportConfig` command causes all configuration and state information for the storage array to be returned each time a critical Major Event Log (MEL) event is detected. The configuration and state information is returned in the form of an object graph. The object graph contains all relevant logical and physical objects and their associated state information for the storage array.

The `set storageArray autoSupportConfig` command collects configuration and state information in this way:

- Automatic collection of the configuration and state information occurs every 72 hours. The configuration and state information is saved to the storage array zip archive file. The archive file has a time stamp that is used to manage the archive files.
- Two storage array zip archive files are maintained for each storage array. The zip archive files are kept on a drive. After the 72-hour time period is exceeded, the oldest archive file is always overwritten during the new cycle.
- After you enable automatic collection of the configuration and state information using this command, an initial collection of information starts. Collecting information after the you issue the command makes sure that one archive file is available and starts the time stamp cycle.

You can run the `set storageArray autoSupportConfig` command on more than one storage array.

Cache Block Size

When you define cache block sizes, use the 4-KB cache block size for storage arrays that require I/O streams that are typically small and random. Use the 8-KB cache block size when the majority of your I/O streams are larger than 4 KB but smaller than 8 KB. Use the 16-KB cache block size or the 32-KB cache block size for storage arrays that require large data transfer, sequential, or high-bandwidth applications.

The `cacheBlockSize` parameter defines the supported cache block size for all of the volumes in the storage array. Not all controller types support all cache block sizes. For redundant configurations, this parameter includes all of the volumes that are owned by both controllers within the storage array.

Cache Flush Start and Cache Flush Stop

When you define values to start a cache flush, a value that is too low increases the chance that data needed for a host read is not in the cache. A low value also increases the number of drive writes that are necessary to maintain the cache level, which increases system overhead and decreases performance.

When setting storage array cache settings, the value of the `cacheFlushStart` parameter must always be greater than or equal to the value of the `cacheFlushStop` parameter. For example, if the value of the `cacheFlushStart` parameter is set to 80, you may set the value of the `cacheFlushStop` parameter within the range of 0 to 80.

When you define values to stop a cache flush, the lower the value, the higher the chance that the data for a host read requires a drive read rather than reading from the cache.

Default Host Type

When you define host types, if Storage Partitioning is enabled, the default host type affects only those volumes that are mapped in the default group. If Storage Partitioning is not enabled, all of the hosts that are attached to the storage array must run the same operating system and be compatible with the default host type.

Media Scan Rate

Media scan runs on all of the volumes in the storage array that have Optimal status, do not have modification operations in progress, and have the `mediaScanRate` parameter enabled. Use the `set volume` command to enable or disable the `mediaScanRate` parameter.

Password

Passwords are stored on each storage array. For best protection, the password must meet these criteria:

- The password must be between eight and 30 characters long.
- The password must contain at least one uppercase letter.
- The password must contain at least one lowercase letter.
- The password must contain at least one number.
- The password must contain at least one non-alphanumeric character, for example, <> @ +.

NOTE If you are using full disk encryption drives in your storage array, you must use these criteria for your storage array password.

NOTE You must set a password for your storagearray before you can create a security key for encrypted full disk encryption drives.

Minimum Firmware Level

5.00 adds the `defaultHostType` parameter.

5.40 adds the `failoverAlertDelay` parameter.

6.14 adds the `alarm` parameter.

7.15 adds more cache block sizes.

Set Storage Array AutoSupport Bundle Disable

This command turns off the AutoSupport (ASUP) bundle collection and transmission for the storage array. You can run this version of the command from the script editor or in a script file.

Syntax

```
set storageArray autoSupport disable
```

Parameters

None.

Minimum Firmware Level

7.86

Set Storage Array AutoSupport Bundle Enable

This command turns on the AutoSupport (ASUP) bundle collection and transmission for the storage array.

Syntax

```
set storageArray autoSupport enable
```

Parameters

None.

Minimum Firmware Level

7.86

Set Storage Array ICMP Response

This command returns the default values for negotiable settings for sessions and connections, which represent the starting point for the storage array for negotiations.

Syntax

```
set storageArray icmpPingResponse=(TRUE | FALSE)
```

Parameter

Parameter	Description
icmpPingResponse	This parameter turns on or turns off Echo Request messages. Set the parameter to TRUE to turn on Echo Request messages. Set the parameter to FALSE to turn off Echo Request messages.

Notes

The Internet Control Message Protocol (ICMP) is used by operating systems in a network to send error messages, test packets, and informational messages related to the IP, such as a requested service is not available or that a host or router could not be reached. The ICMP response command sends ICMP Echo Request messages and receives ICMP Echo Response messages to determine if a host is reachable and the time it takes for packets to get to and from that host.

Minimum Firmware Level

7.10

Set Storage Array iSNS Server IPv4 Address

This command sets the configuration method and address for an IPv4 Internet Storage Name Service (iSNS).

Syntax

```
set storageArray isnsIPv4ConfigurationMethod=[static | dhcp]
isnsIPv4Address=ipAddress
```

Parameters

Parameters	Description
<code>isnsIPv4ConfigurationMethod</code>	The method that you want to use to define the iSNS server configuration. You can enter the IP address for the IPv4 iSNS servers by selecting <code>static</code> . For IPv4, you can choose to have a Dynamic Host Configuration Protocol (DHCP) server select the iSNS server IP address by entering <code>dhcp</code> . To enable DHCP, you must set the <code>isnsIPv4Address</code> parameter to <code>0.0.0.0</code> .
<code>isnsIPv4Address</code>	The IP address that you want to use for the iSNS server. Use this parameter with the <code>static</code> value for IPv4 configurations. If you choose to have a DHCP server set the IP address for an IPv4 Internet iSNS server, you must set the <code>isnsIPv4Address</code> parameter to <code>0.0.0.0</code> .

Notes

The iSNS protocol facilitates the automated discovery, management, and configuration of iSCSI devices and Fibre Channel devices on a TCP/IP network. iSNS provides intelligent storage discovery and management services comparable to those found in Fibre Channel networks, which allow a commodity IP network to function in a similar capacity as a storage area network. iSNS also facilitates a seamless integration of IP networks and Fibre Channel networks, due to its ability to emulate Fibre Channel fabric services and manage both iSCSI devices and Fibre Channel devices.

The DHCP server passes configuration parameters, such as network addresses, to IP nodes. DHCP enables a client to acquire all of the IP configuration parameters that it needs to operate. DHCP lets you automatically allocate reusable network addresses.

Minimum Firmware Level

7.10

Set Storage Array iSNS Server IPv6 Address

This command sets the IPv6 address for the iSNS server.

Syntax

```
set storageArray isnsIPv6Address=ipAddress
```

Parameter

Parameter	Description
<code>isnsIPv6Address</code>	The IPv6 address that you want to use for the iSNS server.

Notes

The iSNS protocol facilitates the automated discovery, management, and configuration of iSCSI devices and Fibre Channel devices on a TCP/IP network. iSNS provides intelligent storage discovery and management services comparable to those found in Fibre Channel networks, which permits a commodity IP network to function in a similar capacity as a storage area network. iSNS also facilitates a seamless integration of IP networks and Fibre Channel networks, due to its ability to emulate Fibre Channel fabric services, and manage both iSCSI devices and Fibre Channel devices. iSNS provides value in any storage network that has iSCSI devices, Fibre Channel devices, or any combination.

Minimum Firmware Level

7.10

Set Storage Array iSNS Server Listening Port

This command sets the iSNS server listening port.

Syntax

```
set storageArray isnsListeningPort=listeningPortIPAddress
```

Parameter

Parameter	Description
isnsListeningPort	The IP address that you want to use for the iSNS server listening port. The range of values for the listening port is 49152 to 65535. The default value is 3205.

Notes

A listening port resides on the database server and is responsible for these activities:

- Listening (monitoring) for incoming client connection requests
- Managing the traffic to the server

When a client requests a network session with a server, a listener receives the actual request. If the client information matches the listener information, then the listener grants a connection to the database server.

Minimum Firmware Level

7.10

Set Storage Array iSNS Server Refresh

This command refreshes the network address information for the iSNS server. This command is valid for only IPv4.

Syntax

```
set storageArray isnsServerRefresh
```

Parameters

None.

Notes

If the DHCP server is not operating at full capability, or if the DHCP server is unresponsive, the refresh operation can take between two and three minutes to complete.

The `set storageArray isnsServerRefresh` command returns an error if you did not set the configuration method to DHCP. To set the configuration method to DHCP, use the `set storageArray isnsIPV4ConfigurationMethod` command.

Minimum Firmware Level

7.10

Set Storage Array Learn Cycle

This command sets the learn cycle for the battery backup unit. The learn cycle enables the storage management software to predict the remaining battery life. Learn cycles run at set intervals and store the results for software analysis.

Syntax

```
set storageArray learnCycleDate  
(daysToNextLearnCycle=numberOfDays |  
day=dayOfTheWeek) time=HH:MM
```

Parameters

Parameter	Description
<code>daysToNextLearnCycle</code>	Valid values are 0 through 7, where 0 is immediately and 7 is in seven days. The <code>daysToNextLearnCycle</code> parameter takes place up to seven days after the next scheduled learn cycle.
<code>day</code>	Valid values for the <code>day</code> parameter include the days of the week (Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, and Saturday). Setting the <code>day</code> causes the next learn cycle to be scheduled on the specified day, after the currently scheduled learn cycle.
<code>time</code>	The time in 24-hour format; for example 8:00 a.m. is entered as 08:00. Nine o'clock p.m. is entered as 21:00, and 9:30 p.m. is entered as 21:30.

Notes

You can set the learn cycle to occur only once during a seven-day period.

The `time` parameter selects a specific time that you want to run the learn cycle. If a value is not entered, the command uses a default value of `00:00` (midnight).

If the day and time specified are in the past, the next learn cycle takes place on the next possible day specified.

Minimum Firmware Level

7.15

Set Storage Array Redundancy Mode

This command sets the redundancy mode of the storage array to either simplex or duplex.

Syntax

```
set storageArray redundancyMode=(simplex | duplex)
```

Parameter

Parameter	Description
<code>redundancyMode</code>	Use <code>simplex</code> mode when you have a single controller. Use <code>duplex</code> mode when you have two controllers.

Minimum Firmware Level

6.10

Set Storage Array Security Key

Use this command to set the security key that is used throughout the storage array to implement the Drive Security premium feature. When any security-capable drive in the storage array is assigned to a secured volume group or disk pool, that drive will be security-enabled using the security key. Before you can set the security key, you must use the `create storageArray securityKey` command to create the security key.

Syntax

```
set storageArray securityKey
```

Parameters

None.

Notes

Security-capable drives have hardware to accelerate cryptographic processing and each has a unique drive key. A security-capable drive behaves like any other drive until it is added to a secured volume group, at which time the security-capable drive becomes security-enabled.

Whenever a security-enabled drive is powered on, it requires the correct security key from the controller before it can read or write data. So, a security-enabled drive uses two keys: the drive key that encrypts and decrypts the data and the security key that authorizes the encryption and decryption processes. The `set storageArray securityKey` command commits the security key to all of the controllers and security-enabled drives in the storage array. The full disk encryption feature ensures that if a security-enabled drive is physically removed from a storage array, its data cannot be read by any other device unless the security key is known.

Minimum Firmware Level

7.50

Set Storage Array Time

This command sets the clocks on both controllers in a storage array by synchronizing the controller clocks with the clock of the host from which you run this command.

Syntax

```
set storageArray time
```

Parameters

None.

Minimum Firmware Level

6.10

Set Storage Array Tray Positions

This command defines the position of the trays in a storage array. You must include all of the trays in the storage array when you enter this command.

Syntax

```
set storageArray trayPositions=(controller |
trayID ... trayIDn)
```

Parameter

Parameter	Description
trayPositions	A list of all of the tray IDs. The sequence of the tray IDs in the list defines the positions for the controller tray and the drive trays in a storage array. Valid values are 0 to 99. Enter the tray ID values separated with a space. Enclose the list of tray ID values in parentheses. For storage arrays where the controller tray has a predefined identifier that is not in the range of valid tray position values, use the <code>controller</code> value.

Notes

This command defines the position of a tray in a storage array by the position of the tray ID in the `trayPositions` list. For example, if you have a controller tray with an ID set to 84 and drive trays with IDs set to 1, 12, and 50, the `trayPositions` sequence (84 1 12 50) places the controller tray in the first position, drive tray 1 in the second position, drive tray 12 in the third position, and drive tray 50 in the fourth position. The `trayPositions` sequence (1 84 50 12) places the controller tray in the second position, drive tray 1 in the first position, drive tray 50 in the third position, and drive tray 12 in the fourth position.

NOTE You must include all of the trays in the storage array in the list defined by the `trayPositions` parameter. If the number of trays in the list does not match the total number of trays in the storage array, an error message appears.

Minimum Firmware Level

6.10

For 6.14 and 6.16, `controller` is not a valid value.

Set Storage Array Unnamed Discovery Session

This command enables the storage array to participate in unnamed discovery sessions.

Syntax

```
set storageArray unnamedDiscoverySession=(TRUE | FALSE)
```

Parameter

Parameter	Description
<code>unnamedDiscoverySession</code>	This parameter turns on or turns off unnamed discovery sessions. Set the parameter to <code>TRUE</code> to turn on unnamed discovery sessions. Set the parameter to <code>FALSE</code> to turn off unnamed discovery sessions.

Notes

Discovery is the process where initiators determine the targets that are available. Discovery occurs at power-on/initialization and also if the bus topology changes, for example, if an extra device is added.

An unnamed discovery session is a discovery session that is established without specifying a target ID in the login request. For unnamed discovery sessions, neither the target ID nor the target portal group ID are available to the targets.

Minimum Firmware Level

7.10

Set Synchronous Mirroring

This command defines the properties for a remote-mirrored pair.

NOTE In previous versions of this command the feature identifier was `remoteMirror`. This feature identifier is no longer valid and is replaced by `syncMirror`.

Syntax

```
set syncMirror (localVolume [volumeName] |
localVolumes ["volumeName1" ... "volumeNameN"])
role=(primary | secondary)
[force=(TRUE | FALSE)]
syncPriority=(highest | high | medium | low | lowest)
autoResync=(enabled | disabled)
writeOrder=(preserved | notPreserved)
writeMode=(synchronous | asynchronous)
```

Parameters

Parameter	Description
localVolume or localVolumes	The name of the primary volume for which you want to define properties. Enclose the primary volume name in square brackets ([]). If the primary volume name has special characters, you also must enclose the primary volume name in double quotation marks (" "). You can enter more than one primary volume name. Enclose all of the primary volume names in one set of square brackets ([]). Enclose each primary volume name in double quotation marks (" "). Separate each primary volume name with a white space.
role	The setting for the volume to act as the primary volume or the secondary volume. To define the volume as the primary volume, set this parameter to <code>primary</code> . To define the volume as the secondary volume, set this parameter to <code>secondary</code> . This parameter applies only when the volume is part of a mirror relationship.
force	The role reversal is forced if the communications link between the storage arrays is down and promotion or demotion on the local side results in a dual-primary condition or a dual-secondary condition. To force a role reversal, set this parameter to <code>TRUE</code> . The default value is <code>FALSE</code> .
syncPriority	The priority that full synchronization has relative to host I/O activity. Valid values are <code>highest</code> , <code>high</code> , <code>medium</code> , <code>low</code> , or <code>lowest</code> .

Parameter	Description
<code>autoResync</code>	The settings for automatic resynchronization between the primary volumes and the secondary volumes of a remote-mirrored pair. This parameter has these values: <ul style="list-style-type: none"> ■ <code>enabled</code> – Automatic resynchronization is turned on. You do not need to do anything further to resynchronize the primary volume and the secondary volume. ■ <code>disabled</code> – Automatic resynchronization is turned off. To resynchronize the primary volumes and the secondary volume, you must run the <code>resume syncMirror</code> command.
<code>writeOrder</code>	This parameter defines write order for data transmission between the primary volume and the secondary volume. Valid values are <code>preserved</code> or <code>notPreserved</code> .
<code>writeMode</code>	This parameter defines how the primary volume writes to the secondary volume. Valid values are <code>synchronous</code> or <code>asynchronous</code> .

Notes

When you use this command, you can specify one or more of the optional parameters.

Synchronization priority defines the amount of system resources that are used to synchronize the data between the primary volumes and the secondary volumes of a mirror relationship. If you select the highest priority level, the data synchronization uses the most system resources to perform the full synchronization, which decreases the performance for host data transfers.

The `writeOrder` parameter applies only to asynchronous mirrors and makes them become part of a consistency group. Setting the `writeOrder` parameter to `preserved` causes the remote-mirrored pair to transmit data from the primary volume to the secondary volume in the same order as the host writes to the primary volume. In the event of a transmission link failure, the data is buffered until a full synchronization can occur. This action can require additional system overhead to maintain the buffered data, which slows operations. Setting the `writeOrder` parameter to `notPreserved` frees the system from having to maintain data in a buffer, but it requires forcing a full synchronization to make sure that the secondary volume has the same data as the primary volume.

Minimum Firmware Level

6.10

Set Thin Volume Attributes

This command defines the properties for a thin volume. You can use the parameters to define properties for one or more thin volumes.

Syntax

```
set (volume ["volumeName"] |
volumes ["volumeName1" ... "volumeNameN"] | volume <wwID>)
[newCapacity=capacityValue |
repositoryMaxCapacity=capacityValue |
repositoryExpansionPolicy=(automatic|manual)|
warningThresholdPercent=warningThresholdPercentValue |
addRepositoryCapacity=capacity-spec]
```

Parameters

Parameter	Description
volume or volumes	The name of the specific volume for which you want to define properties. Enclose the volume name in square brackets ([]). If the volume name has special characters, you also must enclose the volume name in double quotation marks (" ") inside square brackets. If the volume name consists only of numbers, such as "1002," you also must enclose the volume name in double quotation marks (" ") inside square brackets. You can enter more than one volume name. Enclose all of the volume names in one set of double quotation marks (" ") inside square brackets ([]). Separate each volume name with a white space.
volume	The World Wide Identifier (WWID) of the volume for which you are setting properties. You can use the WWID instead of the volume name to identify the volume. Enclose the WWID in angle brackets (<>).
newCapacity	This parameter increases the virtual capacity of the thin volume. The virtual capacity is the value that the volume will report to a host that is mapped to the volume. Values smaller or equal to the existing capacity will cause an error. Size is defined in units of bytes, KB, MB, GB, or TB. The minimum virtual capacity is 32 MB. The maximum virtual capacity is 63 TB.

Parameter	Description
repositoryMaxCapacity	This parameter sets the maximum capacity of the repository volume. The value must not be smaller than the physical capacity of the repository volume. If the new value results in a reduction in capacity to a level below the warning threshold, the command will produce an error.
repositoryExpansionPolicy	This parameter sets the expansion policy to <code>automatic</code> or <code>manual</code> . When you change the policy from <code>automatic</code> to <code>manual</code> , the maximum capacity value (quota) changes to the physical capacity of the repository volume.
warningThresholdPercent	The percentage of thin volume capacity at which you receive a warning alert that the thin volume is nearing full. Use integer values. For example, a value of 70 means 70 percent. Valid values are from 1 to 100. Setting this parameter to 100 disables warning alerts.
addRepositoryCapacity	This parameter allocates capacity from the free extent of the disk pool. If insufficient space is available the command fails.

Notes

When you use this command, you can specify one or more of the optional parameters. The following table lists the capacity limits for a thin volume.

Type of Capacity	Size
Minimum virtual capacity	32 MB
Maximum virtual capacity	63 TB
Minimum physical capacity	4 GB
Maximum physical capacity	64 TB

Thin volumes support all of the operations that standard volumes do with the following exceptions:

- You cannot change the segment size of a thin volume.
- You cannot enable the pre-read redundancy check for a thin volume.
- You cannot use a thin volume as the target volume in a volume copy.
- You cannot use a thin volume in a snapshot (legacy) operation.
- You cannot use a thin volume in a Synchronous Mirroring operation.

If you want to change a thin volume to a standard volume, use the volume copy operation to create a copy of the thin volume. The target of a volume copy is always a standard volume.

Minimum Firmware Level

7.83

Set Tray Alarm

This command turns on, turns off, or mutes the audible alarm for a specific tray or all of the trays in a storage array.

Syntax

```
set (allTrays | tray [trayID]
    alarm=(enable | disable | mute))
```

Parameters

Parameter	Description
allTrays	The setting to select all of the trays in a storage array that have audible alarms that you want to turn on, turn off, or mute.
tray	The specific tray that has the audible alarm that you want to turn on, turn off, or mute. Tray ID values are 0 to 99. Enclose the tray ID value in square brackets ([]).
alarm	The setting for the audible alarm. This alarm has these values: <ul style="list-style-type: none"> ■ enable – The audible alarm is turned on and sounds if a fault occurs. ■ disable – The audible alarm is turned off and does not sound if a fault occurs. ■ mute – The audible alarm is turned off if it is sounding. (If another fault occurs after you set the audible alarm to <code>mute</code> , the audible alarm sounds again.)

Minimum Firmware Level

6.16

Set Tray Identification

This command sets the tray ID of a controller tray, a controller-drive tray, or a drive tray in a storage array. This command is valid only for controller trays, controller-drive trays, or drive trays that have tray IDs that you can set through the controller firmware. You cannot use this command for controller trays, controller-drive trays, or drive trays that have a tray ID that you set with a switch.

Syntax

```
set tray ["serialNumber"] id=trayID
```

Parameters

Parameter	Description
tray	The serial number of the controller tray, controller-drive tray, or the drive tray for which you are setting the tray ID. Serial numbers can be any combination of alphanumeric characters and any length. Enclose the serial number in double quotation marks (" ").
id	The value for the controller tray tray ID, controller-drive tray tray ID, or the drive tray tray ID. Tray ID values are 0 through 99. You do not need to enclose the tray ID value in parentheses.

Notes

This command originally supported the CE6998 controller tray. The CE6998-series controller trays can connect to a variety of drive trays, including those whose tray IDs are set by switches. When connecting a CE6998-series controller tray to drive trays whose tray IDs are set by switches, valid values for tray IDs for the controller tray are 80 through 99. This range avoids conflicts with tray IDs that are used for attached drive trays.

Minimum Firmware Level

6.14 adds support for the CE6998 controller tray.

6.16 adds support for controller trays, controller-drive trays, and drive trays that have tray IDs set through the controller firmware.

Set Tray Service Action Allowed Indicator

This command turns on or turns off the Service Action Allowed indicator light on a power-fan canister, an interconnect-battery canister, or an environmental services module (ESM) canister. If the storage array does not support the Service Action Allowed indicator light feature, this command returns an error. If the storage array supports the command but is unable to turn on or turn off the indicator light, this command returns an error.

To turn on or turn off the Service Action Allowed indicator light on the controller canister, use the `set controller serviceAllowedIndicator` command.

Syntax

```
set tray [trayID]
(powerFan [(left | right | top | bottom)] |
interconnect |
esm [(left | right | top | bottom)] |
battery [(left | right)] |
serviceAllowedIndicator=(on | off))
```

Parameters

Parameter	Description
tray	The tray where the power-fan canister, the interconnect canister, the ESM canister, or the battery canister resides. Tray ID values are 0 to 99. Enclose the tray ID value in square brackets ([]). If you do not enter a tray ID value, the tray ID of the controller tray is the default value.
powerFan	The Service Action Allowed indicator light on the power-fan canister that you want to turn on or turn off. Valid power-fan canister identifiers are <code>left</code> , <code>right</code> , <code>top</code> , or <code>bottom</code> . Enclose the power-fan canister identifier in square brackets ([]).
interconnect	The Service Action Allowed indicator light for the interconnect-battery canister.
esm	The Service Action Allowed indicator light for an ESM canister. Valid ESM canister identifiers are <code>left</code> , <code>right</code> , <code>top</code> , or <code>bottom</code> .
battery	The Service Action Allowed indicator light for a battery. Valid battery identifiers are <code>left</code> or <code>right</code> .
serviceAllowedIndicator	The setting to turn on or turn off the Service Action Allowed indicator light. To turn on the Service Action Allowed indicator light, set this parameter to <code>on</code> . To turn off the Service Action Allowed indicator light, set this parameter to <code>off</code> .

Example

This command turns on the Service Action Allowed indicator light for the left ESM in tray 5 with the IP address of 155.155.155.155.

```
SMcli 123.145.167.214 123.145.167.215 -c "set tray [5]
ESM [left] serviceAllowedIndicator=on;"
```

Notes

This command was originally defined for use with the CE6998 controller tray. This command is not supported by controller trays that were shipped before the introduction of the CE6998 controller tray.

Minimum Firmware Level

6.14 adds these parameters:

- `powerFan`
- `interconnect`

Set Volume Attributes for a Volume in a Volume Group

6.16 adds these parameters:

- tray
- esm

7.60 adds the identifiers `top` and `bottom`.

This command defines the properties for a volume in a volume group. You can use most parameters to define properties for one or more volumes. You also can use some parameters to define properties for only one volume. The syntax definitions are separated to show which parameters apply to several volumes and which apply to only one volume.

NOTE In configurations where volume groups consist of more than 32 volumes, the operation can result in host I/O errors or internal controller reboots due to the expiration of the timeout period before the operation completes. If you experience host I/O errors or internal controller reboots, quiesce the host I/O and try the operation again.

Syntax Applicable to One or More Volumes

```
set (allVolumes | volume ["volumeName"] |
volumes ["volumeName1" ... "volumeNameN"] | volume <wwID>)
cacheFlushModifier=cacheFlushModifierValue
cacheWithoutBatteryEnabled=(TRUE | FALSE)
mediaScanEnabled=(TRUE | FALSE)
mirrorCacheEnabled=(TRUE | FALSE)
modificationPriority=(highest | high | medium | low |
lowest)
owner=(a | b)
preReadRedundancyCheck=(TRUE | FALSE)
readCacheEnabled=(TRUE | FALSE)
writeCacheEnabled=(TRUE | FALSE)
cacheReadPrefetch=(TRUE | FALSE)
dataAssuranceDisabled=(TRUE | FALSE)
```

Syntax Applicable to Only One Volume

```
set (volume ["volumeName"] | volume <wwID>)
addCapacity=volumeCapacity
[addDrives=(trayID1,drawerID1,slotID1 ...
trayIDn,drawerIDn,slotIDn)]
redundancyCheckEnabled=(TRUE | FALSE)
segmentSize=segmentSizeValue
userLabel=volumeName
preReadRedundancyCheck=(TRUE | FALSE)
```

Parameters

Parameter	Description
<code>allVolumes</code>	The properties for all volumes in the storage array.
<code>volume</code> or <code>volumes</code>	<p>The name of the specific volume for which you are retrieving information. Enclose the volume name in square brackets ([]). If the volume name has special characters, you also must enclose the volume name in double quotation marks (" ") inside square brackets. If the volume name consists only of numbers, such as "1002," you also must enclose the volume name in double quotation marks (" ") inside square brackets.</p> <p>You can enter more than one volume name. Enclose all of the volume names in one set of double quotation marks (" ") inside square brackets ([]). Separate each volume name with a white space.</p>
<code>volume</code>	The World Wide Identifier (WWID) of the volume for which you are setting properties. You can use the WWID instead of the volume name to identify the volume. Enclose the WWID in angle brackets (<>).
<code>cacheFlushModifier</code>	The maximum amount of time that data for the volume stays in cache before the data is flushed to physical storage. Valid values are listed in the Notes section.
<code>cacheWithoutBatteryEnabled</code>	The setting to turn on or turn off caching without batteries. To turn on caching without batteries, set this parameter to <code>TRUE</code> . To turn off caching without batteries, set this parameter to <code>FALSE</code> .
<code>mediaScanEnabled</code>	The setting to turn on or turn off media scan for the volume. To turn on media scan, set this parameter to <code>TRUE</code> . To turn off media scan, set this parameter to <code>FALSE</code> . (If media scan is disabled at the storage array level, this parameter has no effect.)
<code>mirrorCacheEnabled</code>	The setting to turn on or turn off the mirror cache. To turn on the mirror cache, set this parameter to <code>TRUE</code> . To turn off the mirror cache, set this parameter to <code>FALSE</code> .
<code>modificationPriority</code>	The priority for volume modifications while the storage array is operational. Valid values are <code>highest</code> , <code>high</code> , <code>medium</code> , <code>low</code> , or <code>lowest</code> .

Parameter	Description
owner	The controller that owns the volume. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Use this parameter only if you want to change the volume owner.
preReadRedundancyCheck	<p>The setting to turn on or turn off pre-read redundancy checking. Turning on pre-read redundancy checking verifies the consistency of RAID redundancy data for the stripes containing the read data. Pre-read redundancy checking is performed on read operations only. To turn on pre-read redundancy checking, set this parameter to TRUE. To turn off pre-read redundancy checking, set this parameter to FALSE.</p> <hr/> <p>NOTE Do not use this parameter on non-redundant volumes, such as RAID 0 volumes.</p> <hr/>
readCacheEnabled	The setting to turn on or turn off the read cache. To turn on the read cache, set this parameter to TRUE. To turn off the read cache, set this parameter to FALSE.
writeCacheEnabled	The setting to turn on or turn off the write cache. To turn on the write cache, set this parameter to TRUE. To turn off the write cache, set this parameter to FALSE.
cacheReadPrefetch	The setting to turn on or turn off cache read prefetch. To turn off cache read prefetch, set this parameter to FALSE. To turn on cache read prefetch, set this parameter to TRUE.

Parameter	Description
dataAssuranceDisabled	<p>The setting to turn off data assurance for a specific volume.</p> <p>For this parameter to have meaning, your volume must be capable of data assurance. This parameter changes a volume from one that supports data assurance to a volume that cannot support data assurance.</p> <p>To remove data assurance from a volume that supports data assurance, set this parameter to <code>TRUE</code>.</p> <hr/> <p>NOTE If you remove data assurance from a volume, you cannot reset data assurance for that volume.</p> <hr/> <p>To reset data assurance for the data on a volume, from which you removed data assurance, perform these steps:</p> <ol style="list-style-type: none"> 1. Remove the data from the volume. 2. Delete the volume. 3. Recreate a new volume with the properties of the deleted volume. 4. Set data assurance for the new volume. 5. Move the data to the new volume.
addCapacity	<p>The setting to increase the storage size (capacity) of the volume for which you are defining properties. Size is defined in units of <code>bytes</code>, <code>KB</code>, <code>MB</code>, <code>GB</code>, or <code>TB</code>. The default value is <code>bytes</code>.</p>
addDrives	<p>The setting to add new drives to the volume. For high-capacity drive trays, specify the tray ID value, the drawer ID value, and the slot ID value for the drive. For low-capacity drive trays, specify the tray ID value and the slot ID value for the drive. Tray ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the tray ID value, drawer ID value, and the slot ID value in parentheses. Use this parameter with the <code>addCapacity</code> parameter if you need to specify additional drives to accommodate the new size.</p>

Parameter	Description
<code>redundancyCheckEnabled</code>	The setting to turn on or turn off redundancy checking during a media scan. To turn on redundancy checking, set this parameter to <code>TRUE</code> . To turn off redundancy checking, set this parameter to <code>FALSE</code> .
<code>segmentSize</code>	The amount of data (in KB) that the controller writes on a single drive in a volume before writing data on the next drive. Valid values are 8, 16, 32, 64, 128, 256, or 512.
<code>userLabel</code>	The new name that you want to give an existing volume. Enclose the new volume name in double quotation marks (" ").
<code>preReadRedundancyCheck</code>	The setting to check the consistency of RAID redundancy data on the stripes during read operations. Do not use this operation for non-redundant volumes, for example RAID Level 0. To check redundancy consistency, set this parameter to <code>TRUE</code> . For no stripe checking, set this parameter to <code>FALSE</code> .

Notes

Host I/O errors might result in volume groups with more than 32 volumes. This operation might also result in internal controller reboots due to the expiration of the timeout period before the operation completes. If you experience this issue, quiesce host I/O, and try the operation again.

When you use this command, you can specify one or more of the optional parameters.

You can apply these parameters to only one volume at a time:

- `addCapacity`
- `segmentSize`
- `userLabel`
- `logicalUnitNumber`

Add Capacity, Add Drives, and Segment Size

Setting the `addCapacity` parameter, the `addDrives` parameter, or the `segmentSize` parameter starts a long-running operation that you cannot stop. These long-running operations are performed in the background and do not prevent you from running other commands. To show the progress of long-running operations, use the `show volume actionProgress` command.

The `addDrives` parameter supports both high-capacity drive trays and low-capacity drive trays. A high-capacity drive tray has drawers that hold the drives. The drawers slide out of the drive tray to provide access to the drives. A low-capacity drive tray does not have drawers. For a high-capacity drive tray, you must specify the identifier (ID) of the drive tray, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive tray, you need only specify the ID of the drive tray and the ID of the slot in which a drive resides. For a low-capacity drive tray, an alternative method for identifying a location for a drive is to specify the ID of the drive tray, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

Cache Flush Modifier

Valid values for the cache flush modifier are listed in this table.

Value	Description
<code>Immediate</code>	Data is flushed as soon as it is placed into the cache.
250	Data is flushed after 250 ms.
500	Data is flushed after 500 ms.
750	Data is flushed after 750 ms.
1	Data is flushed after 1 s.
1500	Data is flushed after 1500 ms.
2	Data is flushed after 2 s.
5	Data is flushed after 5 s.
10	Data is flushed after 10 s.
20	Data is flushed after 20 s.
60	Data is flushed after 60 s (1 min.).
120	Data is flushed after 120 s (2 min.).
300	Data is flushed after 300 s (5 min.).
1200	Data is flushed after 1200 s (20 min.).
3600	Data is flushed after 3600 s (1 hr).
<code>Infinite</code>	Data in cache is not subject to any age or time constraints. The data is flushed based on other criteria that are managed by the controller.

Cache Without Battery Enabled

Write caching without batteries enables write caching to continue if the controller batteries are completely discharged, not fully charged, or not present. If you set this parameter to `TRUE` without an uninterruptible power supply (UPS) or other backup power source, you can lose data if the power to the storage array fails. This parameter has no effect if write caching is disabled.

Modification Priority

Modification priority defines the amount of system resources that are used when modifying volume properties. If you select the highest priority level, the volume modification uses the most system resources, which decreases the performance for host data transfers.

Cache Read Prefetch

The `cacheReadPrefetch` parameter enables the controller to copy additional data blocks into cache while the controller reads and copies data blocks that are requested by the host from the drive into cache. This action increases the chance that a future request for data can be fulfilled from cache. Cache read prefetch is important for multimedia applications that use sequential data transfers. The configuration settings for the storage array that you use determine the number of additional data blocks that the controller reads into cache. Valid values for the `cacheReadPrefetch` parameter are `TRUE` or `FALSE`.

Segment Size

The size of a segment determines how many data blocks that the controller writes on a single drive in a volume before writing data on the next drive. Each data block stores 512 bytes of data. A data block is the smallest unit of storage. The size of a segment determines how many data blocks that it contains. For example, an 8-KB segment holds 16 data blocks. A 64-KB segment holds 128 data blocks.

When you enter a value for the segment size, the value is checked against the supported values that are provided by the controller at run time. If the value that you entered is not valid, the controller returns a list of valid values. Using a single drive for a single request leaves other drives available to simultaneously service other requests.

If the volume is in an environment where a single user is transferring large units of data (such as multimedia), performance is maximized when a single data transfer request is serviced with a single data stripe. (A data stripe is the segment size that is multiplied by the number of drives in the volume group that are used for data transfers.) In this case, multiple drives are used for the same request, but each drive is accessed only once.

For optimal performance in a multiuser database or file system storage environment, set your segment size to minimize the number of drives that are required to satisfy a data transfer request.

Minimum Firmware Level

5.00 adds the `addCapacity` parameter.

7.10 adds the `preReadRedundancyCheck` parameter.

7.60 adds the `drawerID` user input.

7.75 adds the `dataAssuranceDisabled` parameter.

Set Volume Attributes for a Disk Pool

This command defines the properties for a volume in a disk pool.

NOTE In configurations where disk pools consist of more than 32 volumes, the operation can result in host I/O errors or internal controller reboots due to the expiration of the timeout period before the operation completes. If you experience host I/O errors or internal controller reboots, bring the host to a quiescent state and try the operation again.

Syntax

```
set (allVolumes | volume ["volumeName"] |
volumes ["volumeName1" ... "volumeNameN"] | volume <wwID>)
[(addCapacity = capacityValue[KB|MB|GB|TB|Bytes] |
[addDrives = (drive-spec-list)]) |
cacheFlushModifier=cacheFlushModifierValue |
cacheReadPrefetch = (TRUE | FALSE) |
cacheWithoutBatteryEnabled=(TRUE | FALSE) |
mediaScanEnabled=(TRUE | FALSE) |
mirrorCacheEnabled=(TRUE | FALSE) |
owner=(a | b) |
preReadRedundancyCheck = (TRUE | FALSE) |
readCacheEnabled=(TRUE | FALSE) |
redundancyCheckEnabled = (TRUE | FALSE) |
segmentSize = segmentSizeValue
userLabel = userLabelValue
writeCacheEnabled=(TRUE | FALSE) |
dataAssuranceDisabled=(TRUE | FALSE)]
```

Parameters

Parameter	Description
<code>allVolumes</code>	The properties for all volumes in the storage array.

Parameter	Description
<p>volume or volumes</p>	<p>The name of the specific volume for which you want to define properties. Enclose the volume name in square brackets ([]). If the volume name has special characters, you also must enclose the volume name in double quotation marks (" ") inside square brackets. If the volume name consists only of numbers, such as "1002," you also must enclose the volume name in double quotation marks (" ") inside square brackets.</p> <p>You can enter more than one volume name. Enclose all of the volume names in one set of double quotation marks (" ") inside square brackets ([]). Separate each volume name with a white space.</p>
<p>volume</p>	<p>The World Wide Identifier (WWID) of the volume for which you are setting properties. You can use the WWID instead of the volume name to identify the volume. Enclose the WWID in angle brackets (<>).</p>
<p>addCapacity</p>	<p>The setting to increase the storage size (capacity) of the volume for which you are defining properties. Size is defined in units of bytes, KB, MB, GB, or TB. The default value is bytes.</p> <p>This parameter is not valid for thin volumes.</p>
<p>addDrives</p>	<p>The setting to add new drives to the volume. For high-capacity drive trays, specify the tray ID value, the drawer ID value, and the slot ID value for the drive. For low-capacity drive trays, specify the tray ID value and the slot ID value for the drive. Tray ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the tray ID value, drawer ID value, and the slot ID value in parentheses. Use this parameter with the addCapacity parameter if you need to specify additional drives to accommodate the new size.</p> <p>This parameter is not valid for disk pool volumes or for thin volumes.</p>
<p>cacheFlushModifier</p>	<p>The maximum amount of time that data for the volume stays in cache before the data is flushed to physical storage. Valid values are listed in the Notes section.</p>
<p>cacheReadPrefetch</p>	<p>The setting to turn on or turn off cache read prefetch. To turn off cache read prefetch, set this parameter to FALSE. To turn on cache read prefetch, set this parameter to TRUE.</p>

Parameter	Description
cacheWithoutBatteryEnabled	The setting to turn on or turn off caching without batteries. To turn on caching without batteries, set this parameter to TRUE. To turn off caching without batteries, set this parameter to FALSE.
mediaScanEnabled	The setting to turn on or turn off media scan for the volume. To turn on media scan, set this parameter to TRUE. To turn off media scan, set this parameter to FALSE. (If media scan is disabled at the storage array level, this parameter has no effect.)
mirrorCacheEnabled	The setting to turn on or turn off the mirror cache. To turn on the mirror cache, set this parameter to TRUE. To turn off the mirror cache, set this parameter to FALSE.
owner	The controller that owns the volume. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Use this parameter only if you want to change the volume owner.
preReadRedundancyCheck	The setting to turn on or turn off pre-read redundancy checking. Turning on pre-read redundancy checking verifies the consistency of RAID redundancy data for the stripes containing the read data. Pre-read redundancy checking is performed on read operations only. To turn on pre-read redundancy checking, set this parameter to TRUE. To turn off pre-read redundancy checking, set this parameter to FALSE. NOTE Do not use this parameter on non-redundant volumes, such as RAID 0 volumes.
readCacheEnabled	The setting to turn on or turn off the read cache. To turn on the read cache, set this parameter to TRUE. To turn off the read cache, set this parameter to FALSE.
redundancyCheckEnabled	The setting to turn on or turn off redundancy checking during a media scan. To turn on redundancy checking, set this parameter to TRUE. To turn off redundancy checking, set this parameter to FALSE.
segmentSize	The amount of data (in KB) that the controller writes on a single drive in a volume before writing data on the next drive. Valid values are 8, 16, 32, 64, 128, 256, or 512
userLabel	The new name that you want to give an existing volume. Enclose the new volume name in double quotation marks (" ").

Parameter	Description
<code>writeCacheEnabled</code>	The setting to turn on write cache capability.
<code>ssdCacheEnabled</code>	The setting to turn on flash cache capability.

Notes

When you use this command, you can specify one or more of the optional parameters.

You can apply these parameters to only one volume at a time:

- `addCapacity`
- `segmentSize`
- `userLabel`

Add Capacity, Add Drives, and Segment Size

Setting the `addCapacity` parameter, the `addDrives` parameter, or the `segmentSize` parameter starts a long-running operation that you cannot stop. These long-running operations are performed in the background and do not prevent you from running other commands. To show the progress of long-running operations, use the `show volume actionProgress` command.

The `addDrives` parameter supports both high-capacity drive trays and low-capacity drive trays. A high-capacity drive tray has drawers that hold the drives. The drawers slide out of the drive tray to provide access to the drives. A low-capacity drive tray does not have drawers. For a high-capacity drive tray, you must specify the identifier (ID) of the drive tray, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive tray, you need only specify the ID of the drive tray and the ID of the slot in which a drive resides. For a low-capacity drive tray, an alternative method for identifying a location for a drive is to specify the ID of the drive tray, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

Cache Flush Modifier

Valid values for the cache flush modifier are listed in this table.

Value	Description
<code>Immediate</code>	Data is flushed as soon as it is placed into the cache.
<code>250</code>	Data is flushed after 250 ms.
<code>500</code>	Data is flushed after 500 ms.
<code>750</code>	Data is flushed after 750 ms.
<code>1</code>	Data is flushed after 1 s.
<code>1500</code>	Data is flushed after 1500 ms.
<code>2</code>	Data is flushed after 2 s.

Value	Description
5	Data is flushed after 5 s.
10	Data is flushed after 10 s.
20	Data is flushed after 20 s.
60	Data is flushed after 60 s (1 min.).
120	Data is flushed after 120 s (2 min.).
300	Data is flushed after 300 s (5 min.).
1200	Data is flushed after 1200 s (20 min.).
3600	Data is flushed after 3600 s (1 hr.).
Infinite	Data in cache is not subject to any age or time constraints. The data is flushed based on other criteria that are managed by the controller.

Cache Without Battery Enabled

Write caching without batteries enables write caching to continue if the controller batteries are completely discharged, not fully charged, or not present. If you set this parameter to `TRUE` without an uninterruptible power supply (UPS) or other backup power source, you can lose data if the power to the storage array fails. This parameter has no effect if write caching is disabled.

Modification Priority

Modification priority defines the amount of system resources that are used when modifying volume properties. If you select the highest priority level, the volume modification uses the most system resources, which decreases the performance for host data transfers.

Cache Read Prefetch

The `cacheReadPrefetch` parameter enables the controller to copy additional data blocks into cache while the controller reads and copies data blocks that are requested by the host from the drive into cache. This action increases the chance that a future request for data can be fulfilled from cache. Cache read prefetch is important for multimedia applications that use sequential data transfers. The configuration settings for the storage array that you use determine the number of additional data blocks that the controller reads into cache. Valid values for the `cacheReadPrefetch` parameter are `TRUE` or `FALSE`.

Segment Size

The size of a segment determines how many data blocks that the controller writes on a single drive in a volume before writing data on the next drive. Each data block stores 512 bytes of data. A data block is the smallest unit of storage. The size of a segment determines how many data blocks that it contains. For example, an 8-KB segment holds 16 data blocks. A 64-KB segment holds 128 data blocks.

When you enter a value for the segment size, the value is checked against the supported values that are provided by the controller at run time. If the value that you entered is not valid, the controller returns a list of valid values. Using a single drive for a single request leaves other drives available to simultaneously service other requests.

If the volume is in an environment where a single user is transferring large units of data (such as multimedia), performance is maximized when a single data transfer request is serviced with a single data stripe. (A data stripe is the segment size that is multiplied by the number of drives in the volume group that are used for data transfers.) In this case, multiple drives are used for the same request, but each drive is accessed only once.

For optimal performance in a multiuser database or file system storage environment, set your segment size to minimize the number of drives that are required to satisfy a data transfer request.

Minimum Firmware Level

7.83

Set Volume Copy

NOTE With firmware version 7.83 the `copyType=(online | offline)` parameter is no longer used.

This command defines the properties for a volume copy pair. This command is valid for both snapshot (legacy) volume copy pairs and new snapshot volume copy pairs.

Syntax

```
set volumeCopy target [targetName]  
[source [sourceName]]  
copyPriority=(highest | high | medium | low | lowest)  
targetReadOnlyEnabled=(TRUE | FALSE)]
```

Parameters

Parameter	Description
target	The name of the target volume for which you want to define properties. Enclose the target volume name in square brackets ([]). If the target volume name has special characters, you also must enclose the target volume name in double quotation marks (" ").
source	The name of the source volume for which you want to define properties. Enclose the source volume name in square brackets ([]). If the source volume name has special characters, you also must enclose the source volume name in double quotation marks (" ").
copyPriority	The priority that the volume copy has relative to host I/O activity. Valid values are highest, high, medium, low, or lowest.
targetReadOnlyEnabled	The setting so that you can write to the target volume or only read from the target volume. To write to the target volume, set this parameter to FALSE. To prevent writing to the target volume, set this parameter to TRUE.

Notes

When you use this command, you can specify one or more of the optional parameters.

Minimum Firmware Level

5.40

7.77 adds creating a volume copy with snapshot (legacy).

7.83 removes the `copyType=(online | offline)` parameter.

Set Volume Group

This command defines the properties for a volume group.

Syntax

```
set volumeGroup [volumeGroupName]
addDrives=(trayID1,drawerID1,slotID1 ...
trayIDn,drawerIDn,slotIDn)
raidLevel=(0 | 1 | 3 | 5 | 6)
owner=(a | b)
```

Parameters

Parameter	Description
<code>volumeGroup</code>	The alphanumeric identifier (including - and _) of the volume group for which you want to set properties. Enclose the volume group identifier in square brackets ([]).
<code>addDrives</code>	The location of the drive that you want to add to the volume group. For high-capacity drive trays, specify the tray ID value, the drawer ID value, and the slot ID value for the drive. For low-capacity drive trays, specify the tray ID value and the slot ID value for the drive. Tray ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the tray ID value, the drawer ID value, and the slot ID value in parentheses.
<code>raidLevel</code>	The RAID level for the volume group. Valid values are 0, 1, 3, 5, or 6.
<code>owner</code>	The controller that owns the volume group. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Use this parameter only if you want to change the volume group owner.

Notes

Host I/O errors might result in volume groups with more than 32 volumes. This operation also might result in internal controller reboots because the timeout period ends before the volume group definition is set. If you experience this issue, quiesce the host I/O operations, and try the command again.

When you use this command, you can specify one or more of the parameters.

NOTE Specifying the `addDrives` parameter or the `raidLevel` parameter starts a long-running operation that you cannot stop.

The `addDrives` parameter supports both high-capacity drive trays and low-capacity drive trays. A high-capacity drive tray has drawers that hold the drives. The drawers slide out of the drive tray to provide access to the drives. A low-capacity drive tray does not have drawers. For a high-capacity drive tray, you must specify the identifier (ID) of the drive tray, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive tray, you need only specify the ID of the drive tray and the ID of the slot in which a drive resides. For a low-capacity drive tray, an alternative method for identifying a location for a drive is to specify the ID of the drive tray, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

Troubleshooting

Attempting to expand large volume groups by adding drives, also called Dynamic Capacity Expansion (DCE), may fail with the following message:

Return code: Error 26 - The modification operation cannot complete because of the number of drives in the volume group and the segment size of the associated volumes. Reduce the segment size of all volumes in the volume group to 128 KB or below using the Change Segment Size option. Then, retry the operation.

Systems running 7.35.xx.xx firmware may fail with the following message instead of the one noted above:

Return code: Error 462 - A SYMbol procedure could not be carried out because the firmware could not allocate sufficient cache memory. Operation when error occurred: PROC_startVolum

In addition to the above messages, a Major Event Log (MEL) event indicating insufficient cache available to complete the DCE operation may occur.

Any controller running 7.xx firmware may encounter this problem.

DCE requires enough memory to buffer the data read from the original volume and the data to be written to the expanded volume. Some combination of number of drives in the expansion operation, stripe size, and whether mirror operations are enabled may result in not enough memory being available to complete the DCE operation.

If the above situation is encountered, possible workarounds are as follows:

- Create the desired size volume group using other unassigned drives.
- Delete the current volume group and then recreate the volume group with the desired number of drives.
- Reduce the segment size being used and then retry the operation.
- If possible, add additional memory to the controller and then retry the operation.

Minimum Firmware Level

6.10

7.10 adds RAID 6 capability.

7.30 removes the `availability` parameter.

7.60 adds the `drawerID` user input.

Set Volume Group Forced State

This command moves a volume group into a Forced state. Use this command if the `start volumeGroup import` command does not move the volume group to an Imported state or if the import operation does not work because of hardware errors. In a Forced state, the volume group can be imported, and you can then identify the hardware errors.

Syntax

```
set volumeGroup [volumeGroupName] forcedState
```

Parameter

Parameter	Description
volumeGroup	The alphanumeric identifier (including - and _) of the volume group that you want to place in a Forced state. Enclose the volume group identifier in square brackets ([]).

Notes

You can move the drives that comprise a volume group from one storage array to another storage array. The CLI provides three commands that let you move the drives. The commands are `start volumeGroup export`, `start volumeGroup import`, and `set volumeGroup forcedState`.

In the Forced state, you can perform an import operation on the volume group.

Minimum Firmware Level

7.10

Set Volume Mapping

This command defines the logical unit number (LUN) mapping between a volume a host or host group. This command is applicable to volumes in either a volume group or disk pools.

NOTE You cannot use this command for a snapshot volume that is used in online volume copy.

Syntax

```
set (volume ["volumeName"] | volume <wwID> | accessVolume)
logicalUnitNumber=LUN
(host="hostName" |
hostGroup=("hostGroupName" | defaultGroup)
```

Parameter

Parameter	Description
volume	The name of the specific volume for which you are retrieving information. Enclose the volume name in square brackets ([]). If the volume name has special characters, you also must enclose the volume name in double quotation marks (" ") inside square brackets. If the volume name consists only of numbers, such as "1002," you also must enclose the volume name in double quotation marks (" ") inside square brackets. The World Wide Identifier (WWID) of the volume for which you are setting properties. You can use the WWID instead of the volume name to identify the volume. Enclose the WWID in angle brackets (<>).
accessVolume	The logical unit number for the access volume. The logical unit number is the only property that you can set for the access volume.
logicalUnitNumber	The logical unit number that you want to use to map to a specific host. This parameter also assigns the host to a host group
host	The name of the host to which the volume is mapped. Enclose the host name in double quotation marks (" ").
hostGroup	The name of the host group to which the volume is mapped. Enclose the host group name in double quotation marks (" "). defaultGroup is the host group that contains the host to which the volume is mapped.

Notes

A host group is an optional topological element that you can define if you want to designate a collection of hosts that share access to the same volumes. The host group is a logical entity. Define a host group only if you have two or more hosts that can share access to the same volumes.

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

The access volume is the volume in a SAN environment that is used for in-band communication between the storage management software and the storage array controller. This volume uses a LUN address and consumes 20 MB of storage space that is not available for application data storage. An access volume is required only for in-band managed storage arrays. If you specify the `accessVolume` parameter, the only property you can set is the `logicalUnitNumber` parameter.

Minimum Firmware Level

5.20

7.83 adds snapshot volume for disk pools.

Show Alert Severities

NOTE This command is an **SMcli** command, not a script command. You must run this command from a command line. You cannot run this command from the script editor in the storage management software

This command shows all of the severities for which an alert is sent. This command cannot show information for a specific type of severity.

Syntax

```
SMcli -alertSeverities
```

Parameters

None.

Minimum Firmware Level

7.83

Show Asynchronous Mirror Group

This command displays configuration information for one or more asynchronous mirror groups. This command also displays the asynchronous mirrored pairs associated with each asynchronous mirror group, including incomplete asynchronous mirrored pairs.

You also can use this command to show the progress of periodic data synchronization on all of the mirrored pairs within the asynchronous mirror group.

Syntax

```
show (allAsyncMirrorGroups | asyncMirrorGroup  
["asyncMirrorGroupName"] )  
[summary]
```

Parameters

Parameter	Description
allAsyncMirrorGroups	Use this parameter if you want to display the properties for all asynchronous mirror groups.
asyncMirrorGroup	Use this parameter to display the properties for an asynchronous mirror group. Enclose the asynchronous mirror group name in square brackets ([]). If the asynchronous mirror group name has special characters, you also must enclose the asynchronous mirror group name in double quotation marks (" ") inside square brackets. If the asynchronous mirror group name consists only of numbers, such as "1002," you also must enclose the asynchronous mirror group name in double quotation marks (" ") inside square brackets.

Parameter	Description
summary	Use this optional parameter to show a concise list of information about the synchronization progress of one or more asynchronous mirror groups.

Minimum Firmware Level

7.840

Show Asynchronous Mirror Group Synchronization Progress

This command displays the progress of *periodic* synchronization of the asynchronous mirror group between the local and remote storage array. This command returns the progress of data synchronization on all of the mirrored pairs within the asynchronous mirror group. This command shows the progress as a percentage of data synchronization that has been completed.

NOTE There are two types of synchronization: initial synchronization and periodic synchronization. Initial asynchronous mirror group synchronization progress is displayed in the **Long Running Operations** dialog and by executing the `show storageArray longRunningOperations` command.

Syntax

```
show asyncMirrorGroup [ "asyncMirrorGroupName" ]
[synchronizationProgress]
```

Parameter

Parameter	Description
asyncMirrorGroup	The name of the asynchronous mirror group for which you want to check the synchronization progress. Enclose the asynchronous mirror group name in square brackets ([]). If the asynchronous mirror group name has special characters, you also must enclose the asynchronous mirror group name in double quotation marks (" ") inside square brackets. If the asynchronous mirror group name consists only of numbers, such as "1002," you also must enclose the asynchronous mirror group name in double quotation marks (" ") inside square brackets.
synchronizationProgress	Use this optional parameter to display the periodic synchronization progress of the asynchronous mirror group.

Minimum Firmware Level

7.84

Show Cache Backup Device Diagnostic Status

This command returns the status of backup device diagnostic tests started by the `start cacheBackupDevice diagnostic` command. If the diagnostics have finished, all of the results of the diagnostic tests are shown. If the diagnostics have not finished, only the results of the diagnostic tests that finished are shown. The results of the test are shown on the terminal, or you can write the results to a file.

Syntax

```
show cacheBackupDevice controller [(a | b)] diagnosticStatus  
[file="fileName"]
```

Parameters

Parameter	Description
controller	The controller that has the cache backup device on which you are running the diagnostic tests. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the storage management software returns a syntax error.
file	The name of the file that contains the result of the diagnostic tests. Enclose the file name in double quotation marks (" "). This command does not automatically append a file extension to the file name. You must add an extension when you enter the file name.

Minimum Firmware Level

7.60 adds the capability for cache backup device diagnostics.

Show Cache Memory Diagnostic Status

This command returns the status of cache memory diagnostics started by the `start controller diagnostic` command. If the diagnostics have finished, all of the results of the diagnostic tests are shown. If all of the diagnostics have not finished, only the results of the diagnostic tests that finished are shown.

Syntax

```
show cacheMemory controller [(a | b)] diagnosticStatus  
file="fileName"
```

Parameters

Parameter	Description
controller	The controller that has the cache memory on which you are running the diagnostic tests. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]).

Parameter	Description
file	The name of the file that contains the result of the diagnostic tests. Enclose the file name in double quotation marks (" "). This command does not automatically append a file extension to the file name. You must add an extension when you enter the file name.

Minimum Firmware Level

7.60 adds the capability for the cache memory diagnostics.

Show Consistency Group

This command returns information about one or more consistency groups.

Syntax

```
show (consistencyGroup [consistencyGroupName] |
consistencyGroups [consistencyGroupName1 ...
consistencyGroupNameN] |
allConsistencyGroups)
[(summary | schedule)]
```

Parameters

Parameter	Description
consistencyGroup or consistencyGroups	The name of the specific consistency group for which you are retrieving information. You can enter more than one consistency group name. Enclose the consistency group name in square brackets ([]). You can enter more than one snapshot image name or sequence number. Enclose all of the snapshot image names in one set of double quotation marks (" ") inside square brackets ([]). Separate each snapshot image name with a white space.
allConsistencyGroups	This setting returns information about all of the consistency groups in the storage array.
summary	This setting returns a concise list of information about the consistency groups.
schedul	This parameter returns information about any schedules for a consistency group.

Minimum Firmware Level

7.83

Show Consistency Group Snapshot Image

This command shows one snapshot image or several snapshot images that are in one or more snapshot consistency groups.

Syntax

```
show ((CGSnapImage [(CGSnapImageName |
CGSnapImageSequenceNumber)]) |
CGSnapImages [(CGSnapImageNumber1
... CGSnapImageNumbern |
CGSnapImageSequenceNumber1
... CGSnapImageSequenceNumbern)]) |
allCGSnapImages)
[summary]
```

Parameters

Parameter	Description
CGSnapImage or CGSnapImages	<p>The name of the snapshot image in a consistency group. The name of a snapshot image is comprised of two parts:</p> <ul style="list-style-type: none"> ■ The name of the consistency group ■ An identifier for the snapshot image in the consistency group. <p>The identifier for the snapshot image can be one of these:</p> <ul style="list-style-type: none"> ■ An integer value that is the sequence number of the snapshot in the consistency group. ■ NEWEST - Use this option when you want to show the latest snapshot image created in the consistency group. ■ OLDEST - Use this option when you want to show the earliest snapshot image created in the consistency group. <p>Enclose the snapshot image name in double quotation marks (" ") inside square brackets ([]).</p> <p>You can enter more than one snapshot image name or sequence number. Enclose all of the snapshot image names in one set of double quotation marks (" ") inside square brackets ([]). Separate each snapshot image name with a white space.</p>
allCGSnapImages	The setting to return all of the snapshot images from the consistency groups.
summary	The setting to return a concise list of information about all of the snapshot images in the storage array.

Notes

The name of a snapshot image has two parts separated by a colon (:):

- The identifier of the snapshot consistency group
- The identifier of the snapshot image

For example, if you want to show snapshot image 12345 in a snapshot consistency group that has the name snapCGroup1, you would use this command:

```
show CGsnapImage [ "snapCGroup1:12345" ];
```

To show the most recent snapshot image in a snapshot consistency group that has the name snapCGroup1, you would use this command:

```
show CGsnapImage [ "snapCGroup1:newest" ];
```

To show the snapshot images in several snapshot consistency groups that have the names snapCGroup1, snapCGroup2, and snapCGroup3, you would use this command:

```
show CGsnapImages [ "snapCGroup1:12345 snapCGroup2:newest  
snapCGroup3:oldest" ];
```

Note that in these examples the snapshot consistency group name is separated from the snapshot image identifier by a colon (:).

Minimum Firmware Level

7.83

Show Controller

For each controller in a storage array, this command returns the following information:

- The status (Online or Offline)
- The current firmware and NVSRAM configuration
- The pending firmware configuration and NVSRAM configuration (if any)
- The board ID
- The product ID
- The product revision
- The serial number
- The date of manufacture
- The cache size or the processor size
- The date and the time to which the controller is set
- The associated volumes (including the preferred owner)
- The Ethernet port
- The physical disk interface
- The host interface, which applies only to Fibre Channel host interfaces

Syntax

```
show (allControllers | controller [(a | b)]) [summary]
```

Parameters

Parameter	Description
allControllers	The setting to return information about both controllers in the storage array.
controller	The setting to return information about a specific controller in the storage array. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]).
summary	The setting to return a concise list of information about both controllers in the storage array.

Notes

The following list is an example of the information that is returned by the `show controller` command. This example only shows how the information is presented and should not be considered to represent best practice for a storage array configuration.

```
Controller in slot A
  Status: Online
  Current configuration
    Firmware version: 96.10.21.00
    Appware version: 96.10.21.00
    Bootware version: 96.10.21.00
    NVSRAM version: N4884-610800-001
  Pending configuration
    Firmware version: Not applicable
    Appware version: Not applicable
    Bootware version: Not applicable
    NVSRAM version: Not applicable
    Transferred on: Not applicable
  Board ID: 4884
  Product ID: INF-01-00
  Product revision: 9610
  Serial number: 1T14148766
  Date of manufacture: October 14, 2006
  Cache/processor size (MB): 1024/128
  Date/Time: Wed Feb 18 13:55:53 MST 2008
  Associated Volumes (* = Preferred Owner):
    1*, 2*, CTL 0 Mirror Repository*, Mirror Repository
1*,
```

```
JCG_Remote_MirrorMenuTests*
Ethernet port: 1
  MAC address: 00:a0:b8:0c:c3:f5
  Host name: ausctlr9
  Network configuration: Static
  IP address: 172.22.4.249
  Subnet mask: 255.255.255.0
  Gateway: 172.22.4.1
  Remote login: Enabled
Drive interface: Fibre
  Channel: 1
  Current ID: 125/0x1
  Maximum data rate: 2 Gbps
  Current data rate: 1 Gbps
  Data rate control: Switch
  Link status: Up
Drive interface: Fibre
  Channel: 2
  Current ID: 125/0x1
  Maximum data rate: 2 Gbps
  Current data rate: 1 Gbps
  Data rate control: Switch
  Link status: Up
Drive interface: Fibre
  Channel: 3
  Current ID: 125/0x1
  Maximum data rate: 2 Gbps
  Current data rate: 1 Gbps
  Data rate control: Switch
  Link status: Up
Drive interface: Fibre
  Channel: 4
  Current ID: 125/0x1
  Maximum data rate: 2 Gbps
  Current data rate: 1 Gbps
  Data rate control: Switch
  Link status: Up
Host interface: Fibre
  Port: 1
  Current ID: Not applicable/0xFFFFFFFF
  Preferred ID: 126/0x0
  NL-Port ID: 0x011100
  Maximum data rate: 2 Gbps
  Current data rate: 1 Gbps
  Data rate control: Switch
  Link status: Up
```

```

Topology: Fabric Attach
World-wide port name: 20:2c:00:a0:b8:0c:c3:f6
World-wide node name: 20:2c:00:a0:b8:0c:c3:f5
Part type: HPFC-5200    revision 10
Host interface: Fibre
Port: 2
Current ID: Not applicable/0xFFFFFFFF
Preferred ID: 126/0x0
NL-Port ID: 0x011100
Maximum data rate: 2 Gbps
Current data rate: 1 Gbps
Data rate control: Switch
Link status: Up
Topology: Fabric Attach
World-wide port name: 20:2c:00:a0:b8:0c:c3:f7
World-wide node name: 20:2c:00:a0:b8:0c:c3:f5
Part type: HPFC-5200    revision 10

```

When you use the `summary` parameter, the command returns the list of information without the drive channel information and the host channel information.

The `show storageArray` command also returns detailed information about the controller.

Minimum Firmware Level

5.43 adds the `summary` parameter.

Show Controller Diagnostic Status

This command returns the status of controller diagnostics started by the `start controller diagnostic` command. If the diagnostics have finished, the entire results of the diagnostic tests are shown. If the diagnostic tests have not finished, only the results of the of the tests that are finished are shown. The results of the test are shown on the terminal, or you can write the results to a file.

Syntax

```
show controller [(a | b)] diagnosticStatus [file=filename]
```

Parameters

Parameter	Description
<code>controller</code>	The setting to return information about a specific controller in the storage array. Valid controller identifiers are <code>a</code> or <code>b</code> , where <code>a</code> is the controller in slot A, and <code>b</code> is the controller in slot B. Enclose the controller identifier in square brackets ([]).
<code>file</code>	The name of the file that contains the results of the diagnostic tests. This command does not automatically append a file extension to the file name. You must add an extension when you enter the file name.

Minimum Firmware Level

7.70 adds the capability for controller diagnostic status.

Show Controller NVSRAM

This command returns a list of the NVSRAM byte values for the specified host type. If you do not enter the optional parameters, this command returns a list of all of the NVSRAM byte values. To view an example of a table of NVSRAM values that are returned by this command, refer to the "Examples of Information Returned by the Show Command" appendix in *Configuring and Maintaining a Storage Array Using the Command Line Interface* document.

Syntax

```
show (allControllers | controller [(a | b)])  
NVSRAM [hostType=hostTypeIndexLabel | host="hostName"]
```

Parameters

Parameter	Description
allControllers	The setting to return information about both controllers in the storage array.
controller	The setting to return information about a specific controller in the storage array. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]).
hostType	The index label or number of the host type. Use the <code>show storageArray hostTypeTable</code> command to generate a list of available host type identifiers.
host	The name of the host that is connected to the controllers. Enclose the host name in double quotation marks (" ").

Notes

Use the `show controller NVSRAM` command to show parts of or all of the NVSRAM before using the `set controller` command to change the NVSRAM values. Before making any changes to the NVSRAM, contact your Technical Support Representative to learn what regions of the NVSRAM you can modify.

Minimum Firmware Level

6.10

Show Current iSCSI Sessions

This command returns information about an iSCSI session for either an iSCSI initiator or an iSCSI target.

Syntax

```
show iscsiInitiator ["initiatorName"] iscsiSessions  
show iscsiTarget ["targetName"] iscsiSessions
```

Parameters

Parameter	Description
<code>iscsiInitiator</code>	The name of the iSCSI initiator for which you want to obtain session information. Enclose the iSCSI initiator name in double quotation marks (" "). You must also enclose the name in either square brackets ([]) or angle brackets (<>).
<code>iscsiTarget</code>	The name of the iSCSI target for which you want to obtain session information. Enclose the iSCSI target name in double quotation marks (" "). You must also enclose the name in either square brackets ([]) or angle brackets (<>).

Notes

If you enter this command without defining any arguments, this command returns information about all of the iSCSI sessions that are currently running. The following command returns information about all of the current iSCSI sessions:

```
show iscsiSessions
```

To limit the information that is returned, enter a specific iSCSI initiator or a specific iSCSI target. This command then returns information about the session for only the iSCSI initiator or the iSCSI target that you named.

Minimum Firmware Level

7.10

Show Disk Pool

This command returns this information about a disk pool:

- The status (such as optimal, online, offline)
- The total capacity
- The preservation capacity, usable capacity and unusable capacity
- The current owner (the controller in slot A or the controller in slot B)
- The drive media (Fibre Channel, SATA, or SAS)
- The drive interface (Fibre Channel, SATA, or SAS)
- The associated volumes and free capacity
- The associated drives

Syntax

```
show diskPool [diskPoolName]
```

Parameter

Parameter	Description
diskPool	The name of the disk pool for which you are retrieving information. Enclose the disk pool name in square brackets ([]).

Notes

Use this command to show the disk pool content of the storage array profile.

Minimum Firmware Level

7.83

Show Drive

For each drive in the storage array, this command returns the following information:

- The total number of drives
- The type of drive (Fibre Channel, SATA, or SAS)
- Information about the basic drive:
 - The tray location and the slot location
 - The status
 - The capacity
 - The data transfer rate
 - The product ID
 - The firmware level
- Information about the drive channel:
 - The tray location and the slot location
 - The preferred channel
 - The redundant channel
- Hot spare coverage
- Details for each drive

Depending on the size of your storage array, this information can be several pages long. To view an example of the drive information that is returned by the `show drives` command, refer to the "Examples of Information Returned by the Show Commands" topic in "Configuring and Maintaining a Storage Array Using the Command Line Interface." In addition, the drive information is returned for the `show storageArray profile` command.

Syntax

```
show (allDrives
[driveMediaType=(HDD | SSD | unknown | allMedia)] |
[driveType=(fibre | SATA | SAS)]) |
drive [trayID,drawerID,slotID] |
drives [trayID1,drawerID1,slotID1 ...
trayIDn,drawerIDn,slotIDn)
summary
```

Parameters

Parameter	Description
allDrives	The setting to return information about all of the drives in the storage array.
driveMediaType	The type of drive media for which you want to retrieve information. The following values are valid types of drive media: <ul style="list-style-type: none"> ■ HDD – Use this option when you have hard drives in the drive tray. ■ SSD – Use this option when you have solid state drives in the drive tray. ■ unknown – Use this option if you are not sure what types of drive media are in the drive tray. ■ allMedia – Use this option when you want to use all types of drive media that are in the drive tray.
driveType	The type of drive for which you want to retrieve information. You cannot mix drive types. <p>Valid drive types are :</p> <ul style="list-style-type: none"> ■ fibre ■ SATA ■ SAS <p>If you do not specify a drive type, the command defaults to fibre.</p>
drive or drives	The location of the drive for which you want to retrieve information. For high-capacity drive trays, specify the tray ID value, the drawer ID value, and the slot ID value for each drive. For low-capacity drive trays, specify the tray ID value and the slot ID value for each drive. Tray ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the tray ID values, the drawer ID values, and the slot ID values in parentheses.
summary	The setting to return the status, the capacity, the data transfer rate, the product ID, and the firmware version for the specified drives.

Notes

To determine information about the type and location of all of the drives in the storage array, use the `allDrives` parameter.

To determine the information about the Fibre Channel, SATA, or SAS drives in the storage array, use the `driveType` parameter.

To determine the type of drive in a specific location, use the `drive` parameter, and enter the tray ID and the slot ID for the drive.

The `drive` parameter supports both high-capacity drive trays and low-capacity drive trays. A high-capacity drive tray has drawers that hold the drives. The drawers slide out of the drive tray to provide access to the drives. A low-capacity drive tray does not have drawers. For a high-capacity drive tray, you must specify the identifier (ID) of the drive tray, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive tray, you need only specify the ID of the drive tray and the ID of the slot in which a drive resides. For a low-capacity drive tray, an alternative method for identifying a location for a drive is to specify the ID of the drive tray, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

Minimum Firmware Level

5.43

7.60 adds the `drawerID` user input and the `driveMediaType` parameter.

Show Drive Channel Statistics

This command shows the cumulative data transfer for the drive channel and error information. If the controller has automatically degraded a drive channel, this command also shows interval statistics. When you use this command, you can show information about one specific drive channel, several drive channels, or all drive channels.

Syntax

```
show (driveChannel [(1 | 2 | 3 | 4 | 5 | 6 | 7 | 8)] |  
      driveChannels [1 2 3 4 5 6 7 8] |  
      allDriveChannels) stats
```

Parameters

Parameter	Description
<code>driveChannel</code>	The identifier number of the drive channel for which you want to show information. Valid drive channel values are 1, 2, 3, 4, 5, 6, 7, or 8. Enclose the drive channel in square brackets ([]). Use this parameter when you want to show the statistics for only one drive channel.

Parameter	Description
driveChannels	The identifier numbers of several drive channels for which you want to show information. Valid drive channel values are 1, 2, 3, 4, 5, 6, 7, or 8. Enclose the drive channels in square brackets ([]) with the drive channel value separated with a space. Use this parameter when you want to show the statistics for more than one drive channel.
allDriveChannels	The identifier that selects all of the drive channels.

Notes

None.

Minimum Firmware Level

6.10

7.15 adds an update to the drive channel identifier.

Show Drive Download Progress

This command returns the status of firmware downloads for the drives that are targeted by the `download drive firmware` command or the `download storageArray driveFirmware` command.

Syntax

```
show allDrives downloadProgress
```

Parameters

None.

Notes

When all of the firmware downloads have successfully completed, this command returns good status. If any firmware downloads fail, this command shows the firmware download status of each drive that was targeted. This command returns the statuses shown in this table.

Status	Definition
Successful	The downloads completed without errors.
Not Attempted	The downloads did not start.
Partial Download	The downloads are in progress.
Failed	The downloads completed with errors.

Minimum Firmware Level

6.10

Show Drive Performance Statistics

For each drive in the storage array, this command returns the following information:

- The storage array in which has the drive or drives
- The current I/O latency
- The maximum I/O latency
- The minimum I/O latency
- The average I/O latency

Syntax

```
show (allDrives
drive [trayID,drawerID,slotID] |
drives [trayID1,drawerID1,slotID1 ...
trayIDn,drawerIDn,slotIDn])
performanceStats
```

Parameters

Parameter	Description
allDrives	The setting to return information about all of the drives in the storage array.
drive or drives	The location of the drive for which you want to retrieve information. For high-capacity drive trays, specify the tray ID value, the drawer ID value, and the slot ID value for each drive. For low-capacity drive trays, specify the tray ID value and the slot ID value for each drive. Tray ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the tray ID values, the drawer ID values, and the slot ID values in parentheses.

Notes

The `drive` parameter supports both high-capacity drive trays and low-capacity drive trays. A high-capacity drive tray has drawers that hold the drives. The drawers slide out of the drive tray to provide access to the drives. A low-capacity drive tray does not have drawers. For a high-capacity drive tray, you must specify the identifier (ID) of the drive tray, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive tray, you need only specify the ID of the drive tray and the ID of the slot in which a drive resides. For a low-capacity drive tray, an alternative method for identifying a location for a drive is to specify the ID of the drive tray, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

The `show drive performanceStats` command returns drive performance statistics as shown in this example:

```

"Performance Monitor Statistics for Storage Array: remote_pp
-
Date/Time: 10/23/12 3:47:27 PM -
Polling interval in seconds: 5"

"Objects", "Current IO Latency", "Maximum IO Latency", "Minimum
IO Latency",
"Average IO Latency"

"Capture Iteration: 1", "", "", "", ""
"Date/Time: 10/23/12 3:47:27 PM", "", "", "", ""
"Drive Tray 0, Slot 1", "0.0", "0.0", "0.0", "0.0"

```

Minimum Firmware Level

7.86

Show Host Interface Card Diagnostic Status

This command returns the status of running, interrupted, or completed host interface card diagnostics started by the `start hostCard diagnostic` command. If the diagnostics have finished, the entire results of the diagnostic tests are shown. If the diagnostics have not finished, only the results of the tests that are finished are shown. The results of the test are shown on the terminal, or you can write the results to a file.

Syntax

```

show hostCard controller [(a | b)] diagnosticStatus
[progressOnly] [file=filename]

```

Parameters

Parameter	Description
<code>controller</code>	The controller that has the host interface card on which you are running the diagnostic tests. Valid controller identifiers are <code>a</code> or <code>b</code> , where <code>a</code> is the controller in slot A, and <code>b</code> is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the storage management software returns a syntax error.
<code>progressOnly</code>	The <code>progressOnly</code> parameter, shows the progress of the diagnostic test without waiting for the diagnostic tests to completely finish.
<code>file</code>	The name of the file that contains the results of the diagnostic tests. This command does not automatically append a file extension to the file name. You must add an extension when you enter the file name.

Notes

The `progressOnly` parameter is useful for seeing the progress of command scripts that need to sequentially complete operations.

Minimum Firmware Level

7.70 adds the capability for controller host interface card diagnostics.

Show Host Ports

For all of the host ports that are connected to a storage array, this command returns this information:

- The host port identifier
- The host port name
- The host type

Syntax

```
show allHostPorts
```

Parameters

None.

Notes

This command returns HBA host port information similar to this example.

HOST PORT IDENTIFIER	HOST PORT NAME	HOST TYPE
12:34:56:54:33:22:22:22	Jupiter1	Solaris
12:34:56:78:98:98:88:88	Pluto1	Windows 2000/Server
2003 Clustered		
54:32:12:34:34:55:65:66	Undefined	Undefined

Minimum Firmware Level

5.40

Show Snapshot Group

This command returns this information about one or more snapshot image groups.

Syntax

```
show (allSnapGroups | snapGroup [snapGroupName] |  
snapGroups ["snapGroupName1" ... "snapGroupName"])  
[summary | schedule]
```

Parameters

Parameter	Description
allSnapGroups	The parameter to return information about all of the snapshot groups in the storage array.

Parameter	Description
snapGroup or snapGroups	The name of the specific snapshot group for which you are retrieving information. Enclose the snapshot group name in square brackets ([]). If you enter more than one snapshot group name, separate each name with a space. Enclose each snapshot group name in double quotation marks (" "). Enclose all of the snapshot group names in only one set of square brackets ([]).
summary	The parameter to return a concise list of information about the snapshot groups.
schedule	The parameter to return a concise list of information about the schedules for the snapshot group copy operations.

Notes

You can use any combination of alphanumeric characters, underscore (_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

This command returns snapshot group information as shown in this example:

```

SNAPSHOT GROUP DETAILS

SNAPSHOT GROUPS-----

SUMMARY

Total Snapshot Groups: 1
Total Snapshot Images: 0
Enabled Snapshot Image Schedules: 0
Status: 1 Optimal, 0 Non Optimal

Name      Type      Status    Associated Base Volume
2_SG_01   Standard  Optimal   2
Total Repository Capacity Available Repository
Capacity Snapshot Image Limit
10.318 GB                10.318 GB
(100%)                   0

Snapshot Images Scheduled
0 No

DETAILS

Snapshot Group "2_SG_01"

Status: Optimal
Type: Standard
Associated base volume: 2

```

Cluster size: 65,536 bytes

Repository

Total repository volumes: 3
Aggregate repository status: Optimal
Total repository capacity: 10.318 GB
Used capacity: 0.000 MB (0%)
Available repository capacity: 10.318 GB (100%)
Repository full policy: Auto-purge Snapshot Images
Utilization alert threshold: 75%

Snapshot images

Total Snapshot images: 0
Auto-delete Snapshot images: Disabled
Snapshot image schedule: Not Applicable

Minimum Firmware Level

7.83

Show Snapshot Image

This command returns information about the snapshot images that a user had previously created.

Syntax for Showing a Specific Snapshot Image

```
show (snapImage ["snapImageName"] |  
snapImages ["snapImageName1" ... "snapImageNameN"] |  
allSnapImages)  
[summary]
```

Parameters

Parameter	Description
snapImage or snapImages	<p>The name of the snapshot image. The name of a snapshot image is comprised of two parts:</p> <ul style="list-style-type: none"> ■ The name of the snapshot group ■ An identifier for the snapshot image in the snapshot group <p>The identifier for the snapshot image can be one of these:</p> <ul style="list-style-type: none"> ■ An integer value that is the sequence number of the snapshot in the snapshot group. ■ NEWEST - Use this option when you want to show the latest snapshot image created in the snapshot group. ■ OLDEST - Use this option when you want to show the earliest snapshot image created in the snapshot group. <p>Enclose the snapshot image name in double quotation marks (" ") inside square brackets ([]).</p> <p>If you enter more than one snapshot image name, separate each name with a space. Enclose each snapshot image name in double quotation marks (" "). Enclose all of the snapshot group names in only one set of square brackets ([]).</p>
allSnapImages	The parameter to return information about all of the snapshot images in the storage array.
summary	This parameter returns a concise list of information about the snapshot images.

Notes

The name of a snapshot image has two parts separated by a colon (:):

- The identifier of the snapshot group
- The identifier of the snapshot image

For example, if you want to show snapshot image 12345 in a snapshot group that has the name snapGroup1, you would use this command:

```
show snapImage [ "snapGroup1:12345" ] ;
```

To show the most recent snapshot image in a snapshot group that has the name snapGroup1, you would use this command:

```
show snapImage [ "snapGroup1:newest" ] ;
```

To show the snapshot images in several snapshot consistency groups that has the names snapGroup1, snapGroup2, and snapGroup3, you would use this command:

```
show snapImages ["snapGroup1:12345 snapGroup2:newest
snapGroup3:oldest"];
```

Minimum Firmware Level

7.83

Show Snapshot Volumes

This command returns information about one or more snapshot volumes.

Syntax

```
show (allSnapVolumes | snapVolume ["snapVolumeName"] |
snapVolumes ["snapVolumeName1" ... "snapVolumeNameN"])
[summary]
```

Parameters

Parameter	Description
allSnapVolumes	The parameter to return information about all of the snapshot volumes in the storage array.
snapVolume or snapVolumes	The name of a specific snapshot volume about which you are retrieving information. Enclose the snapshot group name in double quotation marks (" ") inside of square brackets ([]). If you enter more than one snapshot image name, separate each name with a space. Enclose each snapshot image name in double quotation marks (" "). Enclose all of the snapshot group names in only one set of square brackets ([]).
summary	The parameter to return a concise list of information about the snapshot volumes.

Notes

You can use any combination of alphanumeric characters, underscore (_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

This command returns snapshot volume information as shown in this example:

```
SNAPSHOT VOLUME(SNAPSHOT-IMAGE BASED) SUMMARY
SUMMARY

Total Snapshot Volumes: 1
Most Recent Snapshot Volume: Day month date hh:mm:ss yyyy
Status: 1 Optimal, 0 Non Optimal

Name          Type      Status    Capacity    Associated
Base Volume
2_SV_0001     Standard  Optimal   3.000 GB    2
```

```

Snapshot Volume Timestamp  Snapshot Image
Timestamp    Mode
1/23/12 6:44:31 PM IST    1/23/12 6:27:36 PM
IST          Read Write

```

```

Total Repository Capacity  Available Repository Capacity
1.199 GB                  0.125 MB (0%)

```

The size of your monitor determines how the information wraps and will affect how the information appears.

Minimum Firmware Level

7.83

Show SSD Cache

This command displays information about the SSD cache.

Syntax

```
show ssdCache [ssdCacheName]
```

Parameter

Parameter	Description
ssdCache	The alphanumeric identifier (including - and _) of the SSD cache for which you want to get information. Enclose the identifier in square brackets ([]). If the SSD cache name contains special characters or consists only of numbers, you also must enclose the identifier in double quotation marks (" ") inside square brackets.

Notes

This command returns the SSD cache information similar to this example.

```

SSD Cache name:  my_cache

Status:          Optimal
Type:            Read Only
I/O characteristic type:  File System
Maximum capacity allowed:  1,862.645 GB
Current capacity:  557.792 GB
Additional capacity allowed  1,304.852 GB
Drive capacities:  All 278.896 GB
Quality of Service (QoS) Attributes
  Security capable:  No
  Secure:           No
  Data Assurance (DA) capable:  No
Associated drives:

Tray    Slot
0       4
0       11
Volumes using SSD cache:  volume_test

```

Minimum Firmware Level

7.84

Show SSD Cache Statistics

This command displays data about the SSD cache usage.

Syntax

```

show ssdCache [ssdCacheName] ssdCacheStatistics
[controller=(a|b|both)]
file="filename"

```

Parameters

Parameter	Description
ssdCache	The alphanumeric identifier (including - and _) of the SSD cache for which you want to get information. Enclose the identifier in square brackets ([]). If the SSD cache name contains special characters or consists only of numbers, you also must enclose the identifier in double quotation marks (" ") inside square brackets.
controller	Each controller stores SSD cache metadata for the volumes that it owns. Therefore, the SSD cache statistics are maintained and displayed per controller. Valid controller identifiers are a, b, or both, where a is the controller in slot A, b is the controller in slot B, and both is both controllers. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the default value is both.

Parameter	Description
file	<p>The file path and the file name to which you want to save the SSD cache statistics. Additional statistics are available when you save the statistics to a file.</p> <p>Enclose the file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\logs\statistics.csv".</p>

Notes

The following statistics are displayed on the screen:

- **Reads** – total number of host reads of SSD cache-enabled volumes

Compare the Reads relative to Writes. The Reads need to be greater than the Writes for effective SSD cache operation. The greater the ratio of Reads to Writes the better the operation of the cache.

- **Writes** – total number of host writes to SSD cache-enabled volumes
- **Cache Hits** – a count of the number of cache hits
- **Cache Hits (%)** – derived from Cache Hits/total reads

The Cache Hit percentage should be greater than 50% for effective SSD cache operation. A small number could be indicative of several things:

- Ratio of Reads to Writes is too small
- Reads are not repeated
- Cache capacity is too small

NOTE To help determine the ideal SSD cache capacity, you can run the Performance Modeling Tool using the `start ssdCache [ssdCacheName] performanceModeling` command.

- **Cache Allocation (%)** – The amount of SSD cache storage that is allocated, expressed as a percentage of the SSD cache storage that is available to this controller. Derived from allocated bytes / available bytes.

Cache Allocation percentage normally shows as 100%. If this number is less than 100%, it means either the cache has not been warmed or the SSD cache capacity is larger than all the data being accessed. In the latter case, a smaller SSD cache capacity could provide the same level of performance. Note that this does not indicate that cached data has been placed into the SSD cache, it is simply a preparation step before data can be placed in the SSD cache.

- **Cache Utilization (%)** – The amount of SSD cache storage that contains data from enabled volumes, expressed as a percentage of SSD cache storage that is allocated. This value represents the utilization or density of the SSD cache. Derived from user data bytes / allocated bytes.

Cache Utilization percentage normally is lower than 100%, perhaps much lower. This number shows the percent of SSD cache capacity that is filled with cache data. The reason this number is lower than 100% is that each allocation unit of the SSD cache, the SSD cache cache-block, is divided into smaller units called sub-blocks, which are filled somewhat independently. A higher number is generally better, but performance gains can be significant even with a smaller number.

These additional statistics are included when you save the data to a file:

- **Read Blocks** – Number of blocks in host reads
- **Write Blocks** – Number of blocks in host writes
- **Full Hit Blocks** – Number of block cache hits.

The full hit blocks indicate the number of blocks that have been read entirely from SSD cache. The SSD cache is only beneficial to performance for those operations that are full cache hits.

- **Partial Hits** – Number of host reads where at least one block, but not all blocks, were in the SSD cache. This is an SSD cache **miss** where the reads were satisfied from the base volume.

Partial cache hits and partial cache hit blocks result from an operation that has only a portion of its data in the SSD cache. In this case, the operation must get the data from the cached HDD volume. The SSD cache offers no performance benefit for this type of hit. If the partial cache hit blocks count is higher than the full cache hit blocks, it is possible that a different I/O characteristic type (file system, database, or web server) could improve the performance.

- **Partial Hits - Blocks** – Number of blocks in Partial Hits.

Partial cache hits and partial cache hit blocks result from an operation that has only a portion of its data in the SSD cache. In this case, the operation must get the data from the cached HDD volume. The SSD cache offers no performance benefit for this type of hit. If the partial cache hit blocks count is higher than the full cache hit blocks, it is possible that a different I/O characteristic type (file system, database, or web server) could improve the performance.

- **Misses** – Number of host reads where none of the blocks were in the SSD cache. This is an SSD cache miss where the reads were satisfied from the base volume.
- **Misses - Blocks** – Number of blocks in Misses
- **Populate Actions (Host Reads)** – Number of host reads where data was copied from the base volume to the SSD cache.
- **Populate Actions (Host Reads) - Blocks** – Number of blocks in Populate Actions (Host Reads).
- **Populate Actions (Host Writes)** – Number of host writes where data was copied from the base volume to the SSD cache.

The Populate Actions (Host Writes) count may be zero for the cache configuration settings that do not fill the cache as a result of a Write I/O operation.

- **Populate Actions (Host Writes) - Blocks** – Number of blocks in Populate Actions (Host Writes).
- **Invalidate Actions** – Number of times data was invalidated/removed from the SSD cache. A cache invalidate operation is performed for every host write request, every host read request with Forced Unit Access (FUA), every verify request, and in some other circumstances.
- **Recycle Actions** – Number of times that the SSD cache block has been re-used for another base volume and/or a different LBA range.

For effective cache operation, it is important that the number of recycles is small compared to the combined number of read and write operations. If the number of Recycle Actions is close to the combined number of Reads and Writes, then the SSD cache is thrashing. Either the cache capacity needs to be increased or the workload is not favorable for use with SSD cache.

- **Available Bytes** – Number of bytes available in the SSD cache for use by this controller.

The available bytes, allocated bytes, and user data bytes are used to compute the Cache Allocation % and the Cache Utilization %.

- **Allocated Bytes** – Number of bytes allocated from the SSD cache by this controller. Bytes allocated from the SSD cache may be empty or they may contain data from base volumes.

The available bytes, allocated bytes, and user data bytes are used to compute the Cache Allocation % and the Cache Utilization %.

- **User Data Bytes** – Number of allocated bytes in the SSD cache that contain data from base volumes.

The available bytes, allocated bytes, and user data bytes are used to compute the Cache Allocation % and the Cache Utilization %.

Minimum Firmware Level

7.84

Show Storage Array

This command returns configuration information about the storage array. The parameters return lists of values for the components and features in the storage array. You can enter the command with a single parameter or more than one parameter. If you enter the command without any parameters, the entire storage array profile is shown (which is the same information as if you entered the `profile` parameter).

Syntax

```
show storageArray | autoSupportConfig | profile |  
batteryAge | connections | defaultHostType | healthStatus |  
hostTypeTable | hotSpareCoverage | features | time |  
volumeDistribution | longRunningOperations | summary
```

Parameters

Parameter	Description
profile	The parameter to show all of the properties of the logical components and the physical components that comprise the storage array. The information appears in several screens.
autoSupportConfig	The parameter to return information about the current state of the operation to automatically collect support data. The following information is returned: <ul style="list-style-type: none">■ Whether the operation is enabled or disabled■ The location of the folder where the support data file is located
batteryAge	The parameter to show the status, the age of the battery in days, and the number of days until the battery needs to be replaced.
connections	The parameter to show a list of where the drive channel ports are located and where the drive channels are connected.
defaultHostType	The parameter to show the default host type and the host type index.
healthStatus	The parameter to show the health, logical properties, and physical component properties of the storage array.
hostTypeTable	The parameter to show a table of all of the host types that are known to the controller. Each row in the table shows a host type index and the platform that the host type index represents.
hotSpareCoverage	The parameter to show information about which volumes of the storage array have hot spare coverage and which volumes do not.
features	The parameter to show the feature configuration of the storage array.
time	The parameter to show the current time to which both controllers in the storage array are set.

Parameter	Description
volumeDistribution	The parameter to show the current controller owner for each volume in the storage array.
longRunningOperations	The parameter to show the long running operations for each volume group and each volume in the storage array. The longRunningOperation parameter returns this information: <ul style="list-style-type: none"> ■ Name of the volume group or volume ■ Long running operation ■ Status ■ % complete ■ Time left
summary	The parameter to show a concise list of information about the storage array configuration.

Notes

The `profile` parameter shows detailed information about the storage array. The information appears on several screens on a display monitor. You might need to increase the size of your display buffer to see all of the information. Because this information is so detailed, you might want to save the output of this parameter to a file. To save the output to a file, run the `show storageArray` command that looks like this example.

```
-c "show storageArray profile;" -o
  "c:\\folder\\storageArrayProfile.txt"
```

The previous command syntax is for a host that is running a Windows operating system. The actual syntax varies depending on your operating system.

The previous command syntax is for a host that is running a Windows operating system. The actual syntax varies depending on your operating system.

When you save information to a file, you can use the information as a record of your configuration and as an aid during recovery.

The `batteryAge` parameter returns information in this form.

```
Battery status: Optimal
  Age: 1 day(s)
  Days until replacement: 718 day(s)
```

The newer controller trays do not support the `batteryAge` parameter.

The `defaultHostType` parameter returns information in this form.

```
Default host type: Linux (Host type index 6)
```

The healthStatus parameter returns information in this form.

```
Storage array health status = optimal.
```

The hostTypeTable parameter returns information in this form.

```
NVSRAM HOST TYPE INDEX DEFINITIONS
INDEX AVT STATUS TYPE
0 Disabled Windows NT Non-Clustered (SP5 or higher)
1 (Default) Disabled Windows 2000/Server 2003 Non-Clustered
2 Disabled Solaris
3 Enabled HP-UX
4 Disabled AIX
5 Disabled Irix
6 Enabled Linux
7 Disabled Windows NT Clustered (SP5 or higher)
8 Disabled Windows 2000/Server 2003 Clustered
9 Enabled Netware Non-Failover
10 Enabled PTX
11 Enabled Netware Failover
12 Enabled Solaris (with Veritas DMP)
```

The hotSpareCoverage parameter returns information in this form.

```
The following volume groups are not protected: 2, 1
Total hot spare drives: 0
Standby: 0
In use: 0
```

The features parameter returns information that shows which features are enabled, disabled, evaluation, and available to be installed. This command returns the feature information in a format similar to this:

PREMIUM FEATURE	STATUS
asyncmirror	Trial available
syncMirror	Trial available/Deactivated
thinProvisioning	Trial available
driveSlotLimit	Enabled (12 of 192 used)
snapImage	Enabled (0 of 512 used) - Trial
version expires m/d/y	
snapshot	Enabled (1 of 4 used)
storagePartition	Enabled (0 of 2 used)
volumeCopy	Enabled (1 of 511 used)
SSDSupport	Disabled (0 of 192 used) - Feature
Key required	
driveSecurity	Disabled - Feature Key required
enterpriseSecurityKeyMgr	Disabled - Feature Key required
highPerformanceTier	Disabled - Feature Key required

The time parameter returns information in this form.

```
Controller in Slot A
Date/Time: Thu Jun 03 14:54:55 MDT 2004
Controller in Slot B
Date/Time: Thu Jun 03 14:54:55 MDT 2004
```

The longRunningOperations parameter returns information in this form:

```
LOGICAL
DEVICES OPERATION STATUS TIME REMAINING
Volume-2 Volume Disk Copy 10% COMPLETED 5 min
```

The information fields returned by the longRunningOperations parameter have these meanings:

- NAME is the name of a volume that is currently in a long running operation. The volume name must have the "Volume" as a prefix.
- OPERATION lists the operation being performed on the volume group or volume.
- % COMPLETE shows how much of the long running operation has been performed.
- STATUS can have one of these meanings:
 - Pending – The long running operation has not started but will start after the current operation is completed.
 - In Progress – The long running operation has started and will run until completed or stopped by user request.
- TIME LEFT indicates the duration remaining to completing the current long running operation. The time is in an "hours minute" format. If less than an hour remains, only the minutes are shown. If less than a minute remains, the message "less than a minute" is shown.

The volumeDistribution parameter returns information in this form.

```
volume name: 10
Current owner is controller in slot: A
volume name: CTL 0 Mirror Repository
Current owner is controller in slot: A
volume name: Mirror Repository 1
Current owner is controller in slot:A
volume name: 20
Current owner is controller in slot:A
volume name: JCG_Remote_MirrorMenuTests
Current owner is controller in slot:A
```

Minimum Firmware Level

5.00 adds the defaultHostType parameter.

5.43 adds the summary parameter.

6.10 adds the `volumeDistribution` parameter.

6.14 adds the `connections` parameter.

7.10 adds the `autoSupportConfig` parameter.

7.77 adds the `longRunningOperations` parameter.

7.83 returns information that includes the support for the new features released in the storage management software version 10.83. In addition, this information returned has been expanded to show the status of the features in the storage array.

Show Storage Array Auto Configure

This command shows the default auto-configuration that the storage array creates if you run the `autoConfigure storageArray` command. If you want to determine whether the storage array can support specific properties, enter the parameter for the properties when you run this command. You do not need to enter any parameters for this command to return configuration information.

Syntax

```
show storageArray autoConfiguration
[driveType=(fibre | SATA | SAS)
driveMediaType=(HDD | SSD | unknown | allMedia)
raidLevel=(0 | 1 | 3 | 5 | 6)
volumeGroupWidth=numberOfDrives
volumeGroupCount=numberOfVolumeGroups
volumesPerGroupCount=numberOfVolumesPerGroup
hotSpareCount=numberOfHotspares
segmentSize=segmentSizeValue
cacheReadPrefetch=(TRUE | FALSE)
securityType=(none | capable | enabled)]
```

Parameters

Parameter	Description
driveType	<p>The type of drives that you want to use for the storage array.</p> <p>The <code>driveType</code> parameter is not required if only one type of drive is in the storage array. You must use this parameter when you have more than one type of drive in your storage array.</p> <p>Valid drive types are :</p> <ul style="list-style-type: none"> ■ fibre ■ SATA ■ SAS <p>If you do not specify a drive type, the command defaults to <code>fibre</code>.</p>
driveMediaType	<p>The type of drive media that you want to use for the mirror repository volume group. Valid drive media are these:</p> <ul style="list-style-type: none"> ■ HDD – Use this option when you have hard drives in the drive tray. ■ SSD – Use this option when you have solid state drives in the drive tray. ■ unknown – Use if you are not sure what types of drive media are in the drive tray. ■ allMedia – Use this option when you want to use all types of drive media that are in the drive tray. <p>Use this parameter when you use the <code>repositoryDriveCount</code> parameter.</p> <p>You must use this parameter when you have more than one type of drive media in your storage array.</p>
raidLevel	<p>The RAID level of the volume group that contains the drives in the storage array. Valid RAID levels are 0, 1, 3, 5, or 6.</p>
volumeGroupWidth	<p>The number of drives in a volume group in the storage array, which depends on the capacity of the drives. Use integer values.</p>
volumeGroupCount	<p>The number of volume groups in the storage array. Use integer values.</p>
volumesPerGroupCount	<p>The number of equal-capacity volumes per volume group. Use integer values.</p>

Parameter	Description
hotSpareCount	The number of hot spares that you want in the storage array. Use integer values.
segmentSize	The amount of data (in KB) that the controller writes on a single drive in a volume before writing data on the next drive. Valid values are 8, 16, 32, 64, 128, 256, or 512.
cacheReadPrefetch	The setting to turn on or turn off cache read prefetch. To turn off cache read prefetch, set this parameter to <code>FALSE</code> . To turn on cache read prefetch, set this parameter to <code>TRUE</code> .
securityType	The setting to specify the security level when creating the volume groups and all associated volumes. These settings are valid: <ul style="list-style-type: none"> ■ <code>none</code> – The volume group and volumes are not secure. ■ <code>capable</code> – The volume group and volumes are capable of having security set, but security has not been enabled. ■ <code>enabled</code> – The volume group and volumes have security enabled.

Notes

If you do not specify any properties, this command returns the RAID Level 5 candidates for each drive type. If RAID Level 5 candidates are not available, this command returns candidates for RAID Level 6, RAID Level 3, RAID Level 1, or RAID Level 0. When you specify auto configuration properties, the controllers validate that the firmware can support the properties.

Drives and Volume Group

A volume group is a set of drives that are logically grouped together by the controllers in the storage array. The number of drives in a volume group is a limitation of the RAID level and the controller firmware. When you create a volume group, follow these guidelines:

- Beginning with firmware version 7.10, you can create an empty volume group so that you can reserve the capacity for later use.
- You cannot mix drive types, such as SAS and Fibre Channel, within a single volume group.
- The maximum number of drives in a volume group depends on these conditions:
 - The type of controller
 - The RAID level

- RAID levels include: 0, 1, 10, 3, 5, and 6.
 - In a CDE3992 or a CDE3994 storage array, a volume group with RAID level 0 and a volume group with RAID level 10 can have a maximum of 112 drives.
 - In a CE6998 storage array, a volume group with RAID level 0 and a volume group with RAID level 10 can have a maximum of 224 drives.
 - A volume group with RAID level 3, RAID level 5, or RAID level 6 cannot have more than 30 drives.
 - A volume group with RAID level 6 must have a minimum of five drives.
 - If a volume group with RAID level 1 has four or more drives, the storage management software automatically converts the volume group to a RAID level 10, which is RAID level 1 + RAID level 0.
- If a volume group contains drives that have different capacities, the overall capacity of the volume group is based on the smallest capacity drive.
- To enable tray loss protection, you must create a volume group that uses drives located in at least three drive trays.

Hot Spares

Hot spare drives can replace any failed drive in the storage array. A hot spare drive must have capacity greater than or equal to any drive that can fail. If a hot spare drive is smaller than a failed drive, you cannot use the hot spare drive to rebuild the data from the failed drive. Hot spare drives are available only for RAID Level 1, RAID Level 3, RAID Level 5, or RAID Level 6.

Segment Size

The size of a segment determines how many data blocks that the controller writes on a single drive in a volume before writing data on the next drive. Each data block stores 512 bytes of data. A data block is the smallest unit of storage. The size of a segment determines how many data blocks that it contains. For example, an 8-KB segment holds 16 data blocks. A 64-KB segment holds 128 data blocks.

When you enter a value for the segment size, the value is checked against the supported values that are provided by the controller at run time. If the value that you entered is not valid, the controller returns a list of valid values. Using a single drive for a single request leaves other drives available to simultaneously service other requests. If the volume is in an environment where a single user is transferring large units of data (such as multimedia), performance is maximized when a single data transfer request is serviced with a single data stripe. (A data stripe is the segment size that is multiplied by the number of drives in the volume group that are used for data transfers.) In this case, multiple drives are used for the same request, but each drive is accessed only once.

For optimal performance in a multiuser database or file system storage environment, set your segment size to minimize the number of drives that are required to satisfy a data transfer request.

Cache Read Prefetch

Cache read prefetch lets the controller copy additional data blocks into cache while the controller reads and copies data blocks that are requested by the host from the drive into cache. This action increases the chance that a future request for data can be fulfilled from cache. Cache read prefetch is important for multimedia applications that use sequential data transfers. The configuration settings for the storage array that you use determine the number of additional data blocks that the controller reads into cache. Valid values for the `cacheReadPrefetch` parameter are `TRUE` or `FALSE`.

Minimum Firmware Level

6.10

7.10 adds RAID Level 6 capability and removes hot spare limits.

Show Storage Array Core Dump

This command show details of the core dump on the controller cache, if a core dump is available.

Syntax

```
show storageArray coreDump
```

Parameters

None.

Minimum Firmware Level

7.83

Show Storage Array DBM Database

This command retrieves and shows metadata for the on-board backup locations of a storage array. When there are multiple backup locations, metadata is shown for each location.

Syntax

```
show storageArray dbmDatabase
```

Parameters

None.

Notes

This command returns the configuration database information in a format similar to this example.

Configuration Database MetadataLast Host Backup: <date stamp>						
CTL	LOCATION	REVISION	ID	GEN NUMBER	STATUS	ACCESS MODE
A	Cache	X.Y	999999	999999	R/W	Optimal
B	Cache	X.Y	999999	999999	R/W	Optimal
N/A	Disk	X.Y	999999	999999	R/W	Optimal

Minimum Firmware Level

7.83

Show Storage Array Host Topology

This command returns the storage partition topology, the host type labels, and the host type index for the host storage array.

Syntax

```
show storageArray hostTopology
```

Parameters

None.

Notes

This command returns the host topology information similar to this example.

```

TOPOLOGY DEFINITIONS
DEFAULT GROUP
Default type: Windows 2000/Server 2003 Non-Clustered
Host Group: scott
Host: scott1
Host Port: 28:37:48:55:55:55:55
Alias: scott11
Type: Windows 2000/Server 2003 Clustered
Host: scott2
Host Port: 98:77:66:55:44:33:21:23
Alias: scott21
Type: Windows 2000/Server 2003 Clustered
Host: Bill
Host Port: 12:34:55:67:89:88:88:88
Alias: Bill1
Type: Windows 2000/Server 2003 Non-Clustered
NVS RAM HOST TYPE INDEX DEFINITIONS
INDEX          AVT STATUS TYPE
0              Disabled  Windows NT Non-Clustered (SP5 or
higher)

```

1 (Default)	Disabled	Windows 2000/Server 2003 Non-Clustered
2	Disabled	Solaris
3	Enabled	HP-UX
4	Disabled	AIX
5	Disabled	Irix
6	Enabled	Linux
7	Disabled	Windows NT Clustered (SP5 or higher)
8	Disabled	Windows 2000/Server 2003 Clustered
9	Enabled	Netware Non-Failover
10	Enabled	PTX
11	Enabled	Netware Failover
12	Enabled	Solaris (with Veritas DMP)

Minimum Firmware Level

5.20

Show Storage Array LUN Mappings

This command returns information from the storage array profile about the logical unit number (LUN) mappings in the storage array. Default group LUN mappings are always shown. If you run this command without any parameters, this command returns all of the LUN mappings.

Syntax

```
show storageArray lunMappings [host ["hostName"] |
hostgroup ["hostGroupName"]]
```

Parameters

Parameter	Description
host	The name of a specific host for which you want to see the LUN mappings. Enclose the host name in double quotation marks (") inside of square brackets ([]).
hostGroup	The name of a specific host group for which you want to see the LUN mappings. Enclose the host group name in double quotation marks (") inside of square brackets ([]).

Notes

This command returns host topology information similar to this example.

```
MAPPINGS (Storage Partitioning - Enabled (0 of 16 used))
VOLUME NAME  LUN  CONTROLLER  ACCESSIBLE BY  VOLUME STATUS
Access Volume 7  A,B          Default Group  Optimal
21           21   B           Default Group  Optimal
22           22   B           Default Group  Optimal
```

Minimum Firmware Level

6.10

Show Storage Array Negotiation Defaults

This statement returns information about connection-level settings that are subject to initiator-target negotiation.

Syntax

```
show storageArray iscsiNegotiationDefaults
```

Parameters

None.

Notes

Information returned includes RAID controller tray default settings (that is, those settings that are the starting point for negotiation) and the current active settings.

Minimum Firmware Level

7.10

Show Storage Array Unconfigured iSCSI Initiators

This command returns a list of initiators that have been detected by the storage array but are not yet configured into the storage array topology.

Syntax

```
show storageArray unconfiguredIscsiInitiators
```

Parameters

None.

Minimum Firmware Level

7.10

Show Storage Array Unreadable Sectors

This command returns a table of the addresses of all of the sectors in the storage array that cannot be read. The table is organized with column headings for the following information:

- Volume user label
- Logical unit number (LUN)
- Accessible by (host or host group)
- Date/time
- Volume-relative logical block address (hexadecimal format – 0xxxxxxxx)
- Drive location (tray t, slot s)
- Drive-relative logical block address (hexadecimal format – 0xxxxxxxx)
- Failure type

The data is sorted first by the volume user label and second by the logical block address (LBA). Each entry in the table corresponds to a single sector.

Syntax

```
show storageArray unreadableSectors
```

Parameters

None.

Minimum Firmware Level

6.10

Show String

This command shows a string of text from a script file. This command is similar to the echo command in MS-DOS and UNIX.

Syntax

```
show "textString"
```

Parameters

None.

Notes

Enclose the string in double quotation marks (" ").

Minimum Firmware Level

6.10

**Show Synchronous
Mirroring Volume
Candidates**

This command returns information about the candidate volumes on a remote storage array that you can use as secondary volumes in a Synchronous Mirroring configuration.

NOTE In previous versions of this command the feature identifier was `remoteMirror`. This feature identifier is no longer valid and is replaced by `syncMirror`.

Syntax

```
show syncMirror candidates primary="volumeName"  
remoteStorageArrayName="storageArrayName"
```

Parameters

Parameter	Description
primary	The name of the local volume that you want for the primary volume in the remote-mirrored pair. Enclose the primary volume name in double quotation marks (" ").
remoteStorageArrayName	The remote storage array that contains possible volumes for a secondary volume. If the remote storage array name has special characters, you must also enclose the remote storage array name in double quotation marks (" ").

Minimum Firmware Level

5.40

Show Synchronous Mirroring Volume Synchronization Progress

This command returns the progress of data synchronization between the primary volume and the secondary volume in a Synchronous Mirroring configuration. This command shows the progress as a percentage of data synchronization that has been completed.

NOTE In previous versions of this command the feature identifier was `remoteMirror`. This feature identifier is no longer valid and is replaced by `syncMirror`.

Syntax

```
show syncMirror (localVolume ["volumeName"] |  
localVolumes ["volumeName1" ... "volumeNameN"])  
synchronizationProgress
```

Parameter

Parameter	Description
localVolume or localVolumes	The name of the primary volume of the remote mirrored pair for which you want to check synchronization progress. Enclose the primary volume name in double quotation marks (" ") inside of square brackets ([]).

Minimum Firmware Level

5.40

Show Thin Volume

This command returns the expansion history or the consumed capacity for the specified thin volume or volumes.

Syntax

```
show (allVolumes | volume [volumeName] |  
volumes ["volumeName1" ... "volumeNameN"]) (consumedCapacity  
| (expansionHistory [file=fileName]))
```

Parameters

Parameter	Description
allVolumes	The setting to return information about all of the volumes in the storage array.
volume or volumes	The name of the specific volume for which you are retrieving information. Enclose the volume name in square brackets ([]). If the volume name has special characters, you also must enclose the volume name in double quotation marks (" ") inside square brackets. If the volume name consists only of numbers, such as "1002," you also must enclose the volume name in double quotation marks (" ") inside square brackets. You can enter more than one volume name. Enclose all of the volume names in one set of square brackets ([]). Enclose each volume name in double quotation marks (" "). Separate each volume name with a white space.
consumedCapacity	The setting to return a concise list of information about the consumed capacity of the thin volumes.
expansionHistory	The setting to return a concise list of information about the expansion history of the thin volumes.
file	The file parameter specifies a file to log the output of the expansionHistory parameter. The file is valid only when used with the expansionHistory parameter. An invalid file name causes the command to fail.

Notes

With the `expansionHistory` parameter, the command returns information similar to the example shown below.

```
Thin volume name: volume-nameRepository volume Name: REPOS_NNNN
Logged Time      Expansion Typ      End Capacity
                  Start Capacity
MM/DD/YYYY      Manual|Automatic  NNNNNNNN bytes  NNNNNNNN bytes
HH:MM:SS
```

With the `consumedCapacity` parameter, the command returns information similar to the example shown below.

Volume	Provisioned Capacity	Consumed Capacity	Quota	% Prov. Consumed
<i>volumeName</i>	500.000 GB	230.000 GB	700.000 GB	46%

Minimum Firmware Level

7.83

Show Volume

For the volumes in a storage array, this command returns the following information:

- The number of volumes
- The name
- The status
- The capacity
- The RAID level
- The volume group where the volume is located
- Details:
 - The volume ID
 - The subsystem ID
 - The drive type (Fibre Channel, SATA, or SAS)
 - Tray loss protection
 - The preferred owner
 - The current owner
 - The segment size
 - The modification priority
 - The read cache status (enabled or disabled)
 - The write cache status (enabled or disabled)

- The write cache without batteries status (enabled or disabled)
- The write cache with mirror status (enabled or disabled)
- The flush write cache after time
- The cache read prefetch setting (TRUE or FALSE)
- The enable background media scan status (enabled or disabled)
- The media scan with redundancy check status (enabled or disabled)
- The snapshot (legacy) repository volumes
- The mirror repository volumes
- The snapshot (legacy) volumes
- The snapshot (legacy) copies

To view an example of the information returned by this command, refer to the topic "Examples of Information Returned by the Show Commands" in *Configuring and Maintaining a Storage Array Using the Command Line Interface*.

Syntax

```
show (allVolumes | volume [volumeName] |
volumes ["volumeName1" ... "volumeNameN"]) summary
```

Parameters

Parameter	Description
allVolumes	The setting to return information about all of the volumes in the storage array.
volume or volumes	The name of the specific volume for which you are retrieving information. Enclose the volume name in square brackets ([]). If the volume name has special characters, you also must enclose the volume name in double quotation marks (" ") inside square brackets. If the volume name consists only of numbers, such as "1002," you also must enclose the volume name in double quotation marks (" ") inside square brackets. You can enter more than one volume name. Enclose all of the volume names in one set of square brackets ([]). Enclose each volume name in double quotation marks (" "). Separate each volume name with a white space.
summary	The setting to return a concise list of information about the volumes.

Notes

For snapshot (legacy) volume copies, the `show volume` command returns information about the schedules for the snapshot (legacy) volume copies. The schedule information is in this form:

```
Schedule State: "Active" | "Disabled" | "Completed"
Last Run Time: <mm/dd/yyyy> <hh:mm a.m. | p.m.>
Next Run Time: <mm/dd/yyyy> <hh:mm a.m. | p.m.>
Start Date: <mm/dd/yyyy>End Date: <mm/dd/yyyy> | "No End
Date"
Days of Week: <Sunday - Saturday>, <Sunday - Saturday>, ...
Times for snapshot (legacy) recreate: <hh:mm a.m. | p.m.>,
<hh:mm a.m. | p.m.>
```

Minimum Firmware Level

5.00

5.43 adds the `summary` parameter.

7.77 adds the schedule information for the snapshot (legacy) volume copies.

Show Volume Action Progress

NOTE With firmware version 7.77, the `show volume actionProgress` command is deprecated. Replace this command with `show storageArray longRunningOperations`.

For a long-running operation that is currently running on a volume, this command returns information about the volume action and amount of the long-running operation that is completed. The amount of the long-running operation that is completed is shown as a percentage (for example, 25 means that 25 percent of the long-running operation is completed).

Syntax

```
show volume ["volumeName"] actionProgress
```

Parameter

Parameter	Description
volume	The name of the specific volume for which you are retrieving information. Enclose the volume name in square brackets ([]). If the volume name has special characters, you also must enclose the volume name in double quotation marks (" ") inside square brackets. If the volume name consists only of numbers, such as "1002," you also must enclose the volume name in double quotation marks (" ") inside square brackets.

Minimum Firmware Level

5.43

7.77 deprecates this command.

Show Volume Copy

This command returns this information about volume copy operations:

- The copy status
- The start time stamp
- The completion time stamp
- The copy priority
- The source volume World Wide Identifier (WWID) or the target volume WWID
- The target volume Read-Only attribute setting

You can retrieve information about a specific volume copy pair or all of the volume copy pairs in the storage array. This command is valid for both snapshot (legacy) volume copy pairs and new snapshot volume copy pairs.

Syntax

```
show volumeCopy (allVolumes | source ["sourceName"] |  
target ["targetName"])
```

Parameters

Parameter	Description
allVolumes	The setting to return information about volume copy operations for all of the volume copy pairs.
source	The name of the source volume about which you want to retrieve information. Enclose the source volume name in double quotation marks (" ") inside of square brackets ([]).
target	The name of the target volume about which you want to retrieve information. Enclose the target volume name in double quotation marks (" ") inside of square brackets ([]).

Minimum Firmware Level

5.40

Show Volume Copy Source Candidates

This command returns information about the candidate volumes that you can use as the source for a volume copy operation. This command is valid for both snapshot (legacy) volume copy pairs and new snapshot volume copy pairs.

Syntax

```
show volumeCopy sourceCandidates
```

Parameters

None.

Notes

This command returns volume copy source information as shown in this example.

```
Volume Name: finance
Capacity: 4.0 GB
Volume Group: 1
Volume Name: engineering
Capacity: 4.0 GB
Volume Group: 2
```

Minimum Firmware Level

6.10

Show Volume Copy Target Candidates

This command returns information about the candidate volumes that you can use as the target for a volume copy operation. This command is valid for both snapshot (legacy) volume copy pairs and new snapshot volume copy pairs.

Syntax

```
show volumeCopy source ["sourceName"] targetCandidates
```

Parameter

Parameter	Description
source	The name of the source volume for which you are trying to find a candidate target volume. Enclose the source volume name in double quotation marks (" ") inside of square brackets ([]).

Minimum Firmware Level

6.10

Show Volume Group

This command returns this information about a volume group:

- The status (Online or Offline)
- The drive type (Fibre Channel, SATA, or SAS)
- Tray loss protection (yes or no)
- The current owner (the controller in slot A or the controller in slot B)
- The associated volumes and free capacity
- The associated drives

Syntax

```
show volumeGroup [volumeGroupName]
```

Parameter

Parameter	Description
volumeGroup	The alphanumeric identifier of the volume group (including - and _) for which you want to show information. Enclose the volume group identifier in square brackets ([]).

Notes

This command returns volume group information as shown in this example:

```
Volume Group 1 (RAID 5)
  Status: Online
  Drive type: Fibre Channel
  Tray loss protection: No
  Current owner: Controller in slot A
  Associated volumes and free capacities:
    1 (1 GB), 1R1 (0.2 GB), Free Capacity (134.533 GB)
  Associated drives (in piece order):
    Drive at Tray 1, Slot 14
    Drive at Tray 1, Slot 13
    Drive at Tray 1, Slot 12
```

Minimum Firmware Level

6.10

Show Volume Group Export Dependencies

This command shows a list of dependencies for the drives in a volume group that you want to move from one storage array to a second storage array.

Syntax

```
show volumeGroup [volumeGroupName] exportDependencies
```

Parameter

Parameter	Description
volumeGroup	The alphanumeric identifier (including - and _) of the volume group for which you want to show export dependencies. Enclose the volume group identifier in square brackets ([]).

Notes

This command spins up the drives in a volume group, reads the DACstore, and shows a list of import dependencies for the volume group. The volume group must be in an Exported state or a Forced state.

Minimum Firmware Level

7.10

Show Volume Group Import Dependencies

This command shows a list of dependencies for the drives in a volume group that you want to move from one storage array to a second storage array.

Syntax

```
show volumeGroup [volumeGroupName] importDependencies  
[cancelImport=(TRUE | FALSE)]
```

Parameters

Parameter	Description
volumeGroup	The alphanumeric identifier (including - and _) of the volume group for which you want to show import dependencies. Enclose the volume group identifier in square brackets ([]).
cancelImport	The setting to spin the drives back down after the volume group dependencies have been read. To spin down the drives, set this parameter to TRUE. To let the drives stay spinning, set this parameter to FALSE.

Notes

This command returns the dependencies of a specific volume group, which must be in an Exported state or a Forced state. If a decision is made to retain the listed dependencies, then the `cancelImport` parameter can be enforced to spin the drives back down.

You must run the `show volumeGroup importDependencies` command before you run the `start volumeGroup import` command.

Minimum Firmware Level

7.10

Show Volume Performance Statistics

For each volume in the storage array, this command returns the following information:

- Storage Arrays
- Total IOs/s
- Read %
- Primary Read Cache Hit %
- Primary Write Cache Hit %
- SSD Cache Hit %
- Current MBs/s

- Maximum MBs/s
- Current IOs/s
- Maximum IOs/s
- Minimum IOs/s
- Average IOs/s
- Minimum MBs/s
- Average MBs/s
- Current IO Latency
- Maximum IO Latency
- Minimum IO Latency
- Average IO Latency

Syntax

```
show (allVolumes | volume ["volumeName"]
volumes ["volumeName1" ... "volumeNameN"]) performanceStats
```

Parameters

Parameter	Description
allVolumes	The parameter to return performance statistics about all of the volumes in the storage array.
volume or volumes	The name of the specific volume for which you are retrieving information. Enclose the volume name in square brackets ([]). If the volume name has special characters, you also must enclose the volume name in double quotation marks (" ") inside square brackets. If the volume name consists only of numbers, such as "1002," you also must enclose the volume name in double quotation marks (" ") inside square brackets. You can enter more than one volume name. Enclose all of the volume names in one set of square brackets ([]). Enclose each volume name in double quotation marks (" "). Separate each volume name with a white space.

Notes

Before you run the `show volume performanceStat` command, run the `set session performanceMonitorInterval` command and the `set session performanceMonitorIterations` command to define how often you collect the statistics.

The `show volume performanceStat` command returns volume performance statistics as shown in this example:

Performance Monitor Statistics for Storage Array: Tyler -
 Date/Time: 11/6/12 10:00: 34 AM - Polling interval in
 seconds: 5

"Storage Arrays", "Total IOs", "Read %", "Primary Read Cache
 Hit %",
 "Primary Write Cache Hit %", "SSD Read Cache Hit %", "Current
 MBs/sec",
 "Maximum MBs/sec", "Current IOs/sec", "Maximum
 IOs/sec", "Minimum IOs/sec",
 "Average IOs/sec", "Minimum MBs/sec", "Average
 MBs/sec", "Current IO Latency",
 "Maximum IO Latency", "Minimum IO Latency", "Average IO
 Latency"

"Capture Iteration:
 1", "", "", "", "", "", "", "", "", "", "", "", "", "", "", "", "", ""
 "Date/Time: 11/6/12 10:00:34
 AM", "", "", "", "", "", "", "", "", "", "", "", "", "", "", "", "", ""
 "", "", "" "Volume
 Unnamed", "0.0", "", "0.0", "0.0", "0.0", "0.0", "0.0", "0.0", "0.0",
 "0.0",
 "0.0", "0.0", "0.0", "0.0", "0.0", "0.0", "0.0"

Minimum Firmware Level

6.10

Show Volume Reservations

This command returns information about the volumes that have persistent reservations.

Syntax

```
show (allVolumes | volume [volumeName] |  

  volumes ["volumeName1" ... "volumeNameN"]) reservations
```

Parameters

Parameter	Description
allVolumes	The setting to return persistent reservation information about all of the volumes in the storage array.

Parameter	Description
volume or volumes	The name of the specific volume for which you are retrieving information. Enclose the volume name in square brackets ([]). If the volume name has special characters, you also must enclose the volume name in double quotation marks (" ") inside square brackets. If the volume name consists only of numbers, such as "1002," you also must enclose the volume name in double quotation marks (" ") inside square brackets. You can enter more than one volume name. Enclose all of the volume names in one set of square brackets ([]). Enclose each volume name in double quotation marks (" "). Separate each volume name with a white space.

Minimum Firmware Level

5.40

Start Asynchronous Mirroring Synchronization

This command starts Asynchronous Mirroring synchronization.

Syntax

```
start asyncMirrorGroup ["asyncMirrorGroupName"] synchronize
```

Parameter

Parameter	Description
asyncMirrorGroup	The name of the asynchronous mirror group for which you want to start synchronization. Enclose the asynchronous mirror group name in square brackets ([]). If the asynchronous mirror group name has special characters, you also must enclose the asynchronous mirror group name in double quotation marks (" ") inside square brackets. If the asynchronous mirror group name consists only of numbers, such as "1002," you also must enclose the asynchronous mirror group name in double quotation marks (" ") inside square brackets.

Minimum Firmware Level

7.84

Start Cache Backup Device Diagnostic

ATTENTION Before you run this diagnostic test, make sure that the cache backup device has a status of Optimal.

This command runs diagnostic tests to evaluate the functionality of the device that you use to backup the data in the cache if you lose power to the controller. The diagnostic tests are specific to the backup device that is in the controller. Before you run these tests, make these changes to the controller that has the backup device on which you want to run diagnostics:

- Place the controller into service mode (use the `set controller [(a | b)] availability=serviceMode` command).
- Attach the management client directly to the controller through the management Ethernet port.

NOTE In a dual-controller configuration, you must run these diagnostic tests through the controller that you want to evaluate. You cannot run these diagnostic tests through the partner controller.

Syntax

```
start cacheBackupDevice [(1 | n | all)]
controller [(a | b)]
diagnostic diagnosticType=(basic | extended)
[extendedTestID=(writePatterns | random)]
```

Parameters

Parameter	Description
cacheBackupDevice	<p>The identifier for the cache backup device on which you want to run the diagnostic tests. Valid cache backup device identifiers are 1, 2, 3, 4 or all.</p> <ul style="list-style-type: none">■ 1 for USB1 on the controller circuit board■ 2 for USB2 on the controller circuit board■ 3 for USB3 on the controller circuit board■ 4 for USB4 on the controller circuit board■ all for all of the USBs on the controller circuit board <p>NOTE – If you have only one cache backup device, the all identifier does not work.</p> <p>Enclose the identifier for the cache backup device in square brackets ([]).</p>

Parameter	Description
<code>controller</code>	The controller that has the cache backup device on which you want to run the diagnostic tests. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the storage management software returns a syntax error.
<code>diagnosticType</code>	The level of diagnostic testing that you want to run on the cache backup device. You can run one of these levels of testing: basic – This option validates the basic operation of the ability of the cache backup device to store cache data. This option determines these capabilities of the cache backup device: <ul style="list-style-type: none"> ■ Whether the cache backup device is write protected or the cache can write data to the device. ■ If the cache backup device is approaching its write cycle limit. extended – This option enables you to run more comprehensive diagnostic tests on the host interface card.
<code>extendedTestID</code>	This parameter selects the extended test option that you want to run. If you choose the <code>extended</code> parameter, you also must also use the <code>extendedTestID</code> parameter and one of the extended test options.

Extended Test Option	Description
writePatterns	This option writes a predefined pattern of data in blocks to the entire cache backup device. Each block that was written is then read back, and the data is verified for integrity and accuracy.
random	This option writes a random pattern to each flash block in the cache backup device.

Notes

- When an unexpected power loss occurs, cache memory can have data that has not been written to the drives. This data must be preserved so that it can be written to the drives when power is restored. The contents of the cache memory are backed up to a persistent storage device, such as a USB flash drive, a SATA drive, or a solid state device (SSD).
- The total storage capacity of the flash drives must be equal to the total cache memory, considering that all storage space in a flash drive is not always usable. For example, in a 1-GB flash drive, approximately 968 MB is usable. Also, in some flash drives, the Cyclic Redundancy Check (CRC) needs to be stored along with the data. Because the metadata region is persisted in these flash drives, the storage capacity for the flash drives must be greater than the size of the cache memory.
- You can run the diagnostic test on only one controller in the storage array at any one time.

Minimum Firmware Level

7.60 adds the capability for cache backup device diagnostics.

Start Cache Memory Diagnostic

This command runs extended diagnostic tests to evaluate the functionality of the cache memory in a controller. Before you run these tests, you must make these changes to the controller on which you want to run diagnostics:

- Place the controller into Service mode (use the `set controller [(a | b)] availability=serviceMode` command).
- Attach the management client directly to the controller through the management Ethernet port.

NOTE In a dual controller configuration, you must run these diagnostic tests through the controller that you want to evaluate. You cannot run these diagnostic tests through the partner controller.

Syntax

```
start cacheMemory controller [(a | b)] diagnostic
diagnosticType=(basic | extended)
[extendedTestID=(marchC | patterns | psuedoRndm| DMACopy)]
```

Parameters

Parameter	Description
<code>controller</code>	The controller that has the cache memory on which you want to run the diagnostic tests. Valid controller identifiers are <code>a</code> or <code>b</code> , where <code>a</code> is the controller in slot A, and <code>b</code> is the controller in slot B. Enclose the controller identifier in square brackets (<code>[]</code>). If you do not specify a controller, the storage management software returns a syntax error.
<code>diagnosticType</code>	The level of diagnostic testing that you want to run on the host interface card. You can run one of these levels of testing: <code>basic</code> – This option validates the ability of the cache memory to address and access data. <code>extended</code> – This option enables you to run more comprehensive diagnostic tests on the host interface card.
<code>extendedTestID</code>	This parameter selects the extended test option that you want to run. If you choose the <code>extended</code> parameter, you also must use the <code>extendedTestID</code> parameter and one of the extended test options.

Extended Test Option	Description
<code>marchC</code>	This option performs a March C test on specific regions of the Reconfigurable Processor Assembly (RPA) memory. This option tests for only one set of inverse patterns.
<code>patterns</code>	This option performs a word pattern test where the test sequence proceeds with a series of read/write operations for all locations in the specified memory region. The test uses a set of special patterns. The test writes and verifies several patterns at 32-bit widths.
<code>pseudoRndm</code>	This option generates a non-repetitive pattern for double word length, writes the pattern to the entire region, and reads back the pattern for verification.
<code>DMAcopy</code>	This option tests the capability of Direct Memory Access (DMA) copy operations across regions in the cache memory. This options uses the RPA hardware capabilities to move the data from one region to another region.

Start Consistency Group Snapshot Rollback

Notes

- You can run the diagnostic test on only one controller in the storage array at any one time.

Minimum Firmware Level

7.60 adds the capability for cache memory diagnostics.

This command starts a rollback operation to the member base volumes in a snapshot consistency group. The content of the base volumes changes immediately to match the point-in-time content of the consistency group snapshot volume. The base volumes immediately becomes available for read/write requests after the rollback operation has successfully completed.

The repository volume that is associated with the consistency group snapshot volume continues to track any new changes between the base volume and the consistency group snapshot volume that occur after the rollback operation is completed.

To stop a rollback operation to the member base volumes use the `stop cgSnapImage rollback` command.

Syntax

```
start cgSnapImage ["snapCGID:imageID"] rollback
[memberVolumeSet ["memberVolumeName1" ...
"memberVolumeNameN" ]]
```

Parameter

Parameter	Description
cgSnapImage	<p>The name of the consistency group snapshot image for which you want to start a rollback operation. The name of a snapshot image is comprised of two parts:</p> <ul style="list-style-type: none"> ■ The name of the snapshot group ■ An identifier for the snapshot image in the snapshot group. <p>The identifier for the snapshot image can be one of these:</p> <ul style="list-style-type: none"> ■ An integer value that is the sequence number of the snapshot in the snapshot group. ■ NEWEST - Use this option when you want to show the latest snapshot image created in the snapshot group. ■ OLDEST - Use this option when you want to show the earliest snapshot image created in the snapshot group. <p>Enclose the snapshot image name in double quotation marks (" ") inside square brackets ([]).</p>

Parameter	Description
memberVolumeSet	<p>The name of one or more member base volumes in a consistency group that you want to rollback. Enclose each member base volume name in double quotation marks (" ") inside parentheses (()).</p> <p>You can enter more than one volume name. Enclose all of the volume names in one set of square brackets ([]). Enclose each volume name in double quotation marks (" "). Separate each volume name with a white space.</p> <p>When the memberVolumeSet parameter is not used the rollback process applies to all member volumes of the consistency group.</p>

Notes

The name of a snapshot image has two parts separated by a colon (:):

- The identifier of the snapshot group
- The identifier of the snapshot image

For example, if you want to start a roll back operation for the newest snapshot image in an entire consistency group that has the name CG1, you would use this command:

```
start cgSnapImage ["CG1:newest"] rollback;
```

To start a roll back operation for the snapshot image 12345 for base volume members memVol1, memVol2, and memVol3 in a consistency group that has the name CG2, you would use this command:

```
start cgSnapImage ["CG2:12345"] rollback
memberVolumeset=("memVol1 memVol2 memVol3");
```

Minimum Firmware Level

7.83

Start Controller Diagnostic

This command runs diagnostic tests to evaluate the functionality of the controller card. Before you run these tests, you must make these changes to the controller on which you want to run diagnostics:

- Place the controller into Service Mode (use the `set controller [(a | b)] availability=serviceMode` command).
- Attach the management client directly to the controller through the management Ethernet port.

NOTE In a dual controller configuration, you must run these diagnostic tests through the controller that you want to evaluate. You cannot run these diagnostic tests through the partner controller.

Syntax

```
start controller [(a | b)] diagnostic diagnosticType=(basic
| extended)
[extendedTestID=(SRAM | FIFO | dataCopy| RAID5Parity |
RAID6Parity)]
```

Parameters

Parameter	Description
controller	The controller on which you want to run the diagnostic tests. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the storage management software returns a syntax error.
diagnostic	The level of diagnostic testing that you want to run on the host interface card. You can run one of these levels of testing: basic – This option validates the ability of the base controller to address and access data. extended – This option enables you to run more comprehensive diagnostic tests on the base controller card.
extendedTestID	This parameter selects the extended test option that you want to run. If you choose the extended parameter, you must also use the extendedTestID parameter and one of the extended test options.

Extended Test Option	Description
SRAM	This option tests for address, data, and data retention. The address test attempts to write to specific address offsets. The data test attempts to write several data patterns to the address offsets. The data retention test attempts to write a data pattern and then read the data pattern back after a delay. The purpose of the SRAM option is to find memory parity or error correcting code (ECC) errors.
FIFO	This option tests the active processor chip (APC) first in, first out (FIFO) data transmission of the Zip chip. The APC FIFO channels are tested concurrently by writing and verifying different patterns to each channel.

Extended Test Option	Description
dataCopy	This option tests the ability of the Zip chip to support data copy operations that can copy data from one area of the Zip SDRAM to another area of the Zip SDRAM. This test is performed on any available section of the Zip chip that is not busy.
RAID5Parity	This option tests the ability of the Zip APC to generate and verify RAID 5 parity data. Data buffers are set up in processor memory and parity is generated in processor memory. Some data buffers are set up in parallel architecture (RPA) memory and parity is generated for the data within the RPA memory. The parity that is generated within processor memory is then compared with the parity in the Zip APC.
RAID6Parity	This option tests the ability of the Zip APC to generate and verify RAID 6 parity data. Data buffers are set up in processor memory and parity is generated in processor memory. Some data buffers are set up in redundant parallel architecture (RPA) memory and parity is generated for the data within the RPA memory. The parity that is generated within processor memory is then compared with the parity in the Zip APC.

Notes

You can run the diagnostic test on only one controller in the storage array at any one time.

Minimum Firmware Level

7.60 adds the capability for controller card diagnostics.

Start Controller Trace

This command starts an operation that saves debug trace information to a compressed file. The debug trace information can be used by a Technical Support Representative to help analyze how well a storage array is running.

Syntax

```
start controller [(a | b | both)] trace
dataType=(current | flushed | currentFlushed | all)
[forceFlush=(TRUE | FALSE)]
```

Parameters

Parameter	Description
controller	The controller for which you want to collect the trace debug information. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. You can also simultaneously collect debug for both controllers by entering both. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the storage management software returns a syntax error.
dataType	The type of data that you want to collect: <ul style="list-style-type: none"> ■ current – Retrieves the current DQ traces ■ flushed – Retrieves all flushed DQ traces ■ currentFlushed – Retrieves both the current DQ trace and the flushed DQ trace ■ all – Retrieves the current DQ trace, flushed DQ trace, and all platform DQ traces <p>NOTE – If dataType=flushed and forceFlush=True, an error message is returned indicating that only active traces can be flushed to the buffer on retrieval.</p>
forceFlush	The setting to move the DQ information in the current buffer to the flushed buffer when the DQ trace information defined by the dataType parameter is retrieved. To enable force flush, set this parameter to TRUE. To disable force flush, set this parameter to FALSE. <p>NOTE – If dataType=flushed and forceFlush=True, an error message is returned indicating that only active traces can be flushed to the buffer on retrieval.</p>
file	The file path and the file name to which you want to save the DQ trace information. Refer to the Notes section for information about naming the files.

Notes

The DQ trace information is written to a compressed file with an extension of .zip. The file name is a combination of a user-defined file name and the storage array identifier (SAID). A constant of "dq" is also added to the file name. The complete file name has this form:

user_defined_file_name-SAID-dq.zip

The compressed file contains the information listed in this table.

File Name	Directory	Comments
user_provided_file_name-SAID-A.dq	SAID/timestamp/	The DQ trace data retrieved from controller A.
user_provided_file_name-SAID-B.dq	SAID/timestamp/	The DQ trace data retrieved from controller B.
user_provided_file_name-SAID-trace_description.xml	SAID/timestamp/	The description file in an xml format that describes the DQ file attributes for future data mining.

Minimum Firmware Level

7.75

Start Disk Pool Locate

This command identifies the drives that are logically grouped together to form the specified disk pool by blinking the indicator lights on the drives. (Use the `stop diskPool locate` command to turn off the indicator lights on the drives.)

Syntax

```
start diskPool [diskPoolName] locate
```

Parameter

Parameter	Description
diskPool	The name of the disk pool that you want to locate. Enclose the disk pool name in square brackets ([]).

Minimum Firmware Level

7.83

Start Drive Channel Fault Isolation Diagnostics

This command runs the drive channel fault isolation diagnostics and stores the results.

Syntax

```
start driveChannel [(1 | 2 | 3 | 4 | 5 | 6 | 7 | 8)]
controller [(a | b)] faultDiagnostics
(testDevices=[all |
controller=(a | b) |
esms=[trayID1 (left | right), ... , trayIDN (left | right)] |
drives=[trayID1,slotID1, ... , trayIDn,slotIDN]] |
dataPattern=(fixed | pseudoRandom) |
patternNumber=[(0xhexadecimal | number)] |
maxErrorCount=integer |
testIterations=integer |
timeout=timeInterval)
```

Parameters

Parameter	Description
driveChannel	The identifier number of the drive channel that you want to locate. Valid values for the identifier number for the drive channel are 1, 2, 3, 4, 5, 6, 7, or 8. Enclose the drive channel identifier number in square brackets ([]).
controller	The identifier letter of the controller that you want to test. Valid controller identifier values are a or b, where a is the controller in slot A, and b is the controller in slot B.
testDevices	The identifiers of the devices (controllers, environmental services module [ESMs], or drives) that you want to test. You can specify all or enter the specific identifiers for the devices that you want to diagnose. <ul style="list-style-type: none"> ■ controllers – identifiers are a or b, where a is the RAID controller module in slot A, and b is the RAID controller module in slot B ■ ESMs – identifiers are tray ID and left or right, where tray ID is a value from 0 through 99, and left or right are determined when viewing the drive tray from the rear ■ drives – identifiers are tray ID and slot ID, where tray ID values are 0 to 99 and slot ID values are 0 to 31
dataPattern	The method of repeatability that you want to test.
patternNumber	The hexadecimal data pattern that you want to use to run the test. This number can be any hexadecimal number between 0000 to FFFF. You must place 0x in front to indicate a hexadecimal number.

Parameter	Description
maxErrorCount	The number of errors that you want to accept before terminating the test.
testIterations	The number of times that you want to repeat the test.
timeout	The length of time in minutes that you want to run the test.

Notes

Use the `save driveChannel faultDiagnostics` command and the `stop driveChannel faultDiagnostics` command with the `start driveChannel faultDiagnostics` command. These commands are needed to save diagnostic test results to a file and to stop the diagnostic test.

Examples of valid `patternNumber` entries are `0xA5A5`, `0x3C3C`, `8787`, and `1234`.

You can also stop this command at any time by pressing **Ctrl+C**.

Minimum Firmware Level

7.15

Start Drive Channel Locate

This command identifies the drive trays that are connected to a specific drive channel by turning on the indicator lights for the drive tray that is connected to the drive channel. (Use the `stop driveChannel locate` command to turn off the indicator lights on the drive tray.)

Syntax

```
start driveChannel [(1 | 2 | 3 | 4 | 5 | 6 | 7 | 8)] locate
```

Parameter

Parameter	Description
driveChannel	The identifier number of the drive channel that you want to locate. Valid values for the identifier number for the drive channel are 1, 2, 3, 4, 5, 6, 7, or 8. Enclose the drive channel identifier number in square brackets ([]).

Minimum Firmware Level

6.10

7.15 adds an update to the drive channel identifier.

Start Drive Initialize

This command starts drive initialization.

ATTENTION Possible damage to the storage array configuration – As soon as you enter this command, all user data is destroyed.

Syntax

```
start drive [trayID,drawerID,slotID] initialize
```

Parameter

Parameter	Description
drive	The location of the drive that you want to reconstruct. For high-capacity drive trays, specify the tray ID value, the drawer ID value, and the slot ID value of the drive that you want to revive. For low-capacity drive trays, specify the tray ID value and the slot ID value of the drive that you want to revive. Tray ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the tray ID value, the drawer ID value, and the slot ID value in square brackets ([]).

Notes

The `drive` parameter supports both high-capacity drive trays and low-capacity drive trays. A high-capacity drive tray has drawers that hold the drives. The drawers slide out of the drive tray to provide access to the drives. A low-capacity drive tray does not have drawers. For a high-capacity drive tray, you must specify the identifier (ID) of the drive tray, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive tray, you need only specify the ID of the drive tray and the ID of the slot in which a drive resides. For a low-capacity drive tray, an alternative method for identifying a location for a drive is to specify the ID of the drive tray, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

Minimum Firmware Level

6.10

7.60 adds the `drawerID` user input.

Start Drive Locate

This command locates a drive by turning on an indicator light on the drive. (Run the `stop drive locate` command to turn off the indicator light on the drive.)

Syntax

```
start drive [trayID,drawerID,slotID] locate
```

Parameter

Parameter	Description
drive	The location of the drive that you want to reconstruct. For high-capacity drive trays, specify the tray ID value, the drawer ID value, and the slot ID value of the drive that you want to revive. For low-capacity drive trays, specify the tray ID value and the slot ID value of the drive that you want to revive. Tray ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the tray ID value, the drawer ID value, and the slot ID value in square brackets ([]).

Notes

The `drive` parameter supports both high-capacity drive trays and low-capacity drive trays. A high-capacity drive tray has drawers that hold the drives. The drawers slide out of the drive tray to provide access to the drives. A low-capacity drive tray does not have drawers. For a high-capacity drive tray, you must specify the identifier (ID) of the drive tray, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive tray, you need only specify the ID of the drive tray and the ID of the slot in which a drive resides. For a low-capacity drive tray, an alternative method for identifying a location for a drive is to specify the ID of the drive tray, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

Minimum Firmware Level

6.10

7.60 adds the `drawerID` user input.

Start Drive Reconstruction

This command starts reconstructing a drive.

Syntax

```
start drive [trayID,drawerID,slotID] reconstruct
```

Parameter

Parameter	Description
drive	The location of the drive that you want to reconstruct. For high-capacity drive trays, specify the tray ID value, the drawer ID value, and the slot ID value of the drive that you want to revive. For low-capacity drive trays, specify the tray ID value and the slot ID value of the drive that you want to revive. Tray ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the tray ID value, the drawer ID value, and the slot ID value in square brackets ([]).

Notes

The `drive` parameter supports both high-capacity drive trays and low-capacity drive trays. A high-capacity drive tray has drawers that hold the drives. The drawers slide out of the drive tray to provide access to the drives. A low-capacity drive tray does not have drawers. For a high-capacity drive tray, you must specify the identifier (ID) of the drive tray, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive tray, you need only specify the ID of the drive tray and the ID of the slot in which a drive resides. For a low-capacity drive tray, an alternative method for identifying a location for a drive is to specify the ID of the drive tray, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

Minimum Firmware Level

5.43

7.60 adds the `drawerID` user input.

Start Host Interface Card Diagnostic

This command runs diagnostic tests to evaluate the functionality of the controller host interface card. The diagnostic tests that this command runs are specific to the host interface card that is in the controller. Before you run these tests, you must make these changes to the controller that has the host interface card on which you want to run diagnostics:

- Place the controller into service mode (use the `set controller [(a | b)] availability=serviceMode` command).
- Attach the management client directly to the controller through the management Ethernet port.

NOTE In a dual controller configuration, you must run these diagnostic tests through the controller that you want to evaluate. You cannot run these diagnostic tests through the partner controller.

Syntax

```
start hostCard [(1 | 2 | 3 | 4)] controller [(a | b)]  
diagnostic  
diagnosticType=(basic | extended)  
[extendedTestID=(EDC | DMA | RAM | internalLoopback)]
```

Parameters

Parameter	Description
hostCard	The identifier for host interface card on which you want to run the diagnostic tests. Valid host interface card identifiers are 1, 2, 3, or 4. The value of the identifier is for the position of the host interface card in the controller tray or controller-drive tray. The position of the host interface card depends on the type of controller tray or controller-drive tray in your storage array. See the Notes for more information about the host interface card identifier and the position of the host interface cards in a controller tray. Enclose the controller identifier in square brackets ([]).
controller	The controller that has the host interface card on which you want to run the diagnostic tests. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the storage management software returns a syntax error.
diagnosticType	The level of diagnostic testing that you want to run on the host interface card. You can run one of these levels of testing: basic – This option validates the ability of the host interface card to transport I/O data. This option takes approximately 30 seconds to complete. extended – This option enables you to run more comprehensive diagnostic tests on the host interface card.
extendedTestID	This parameter selects the extended test option that you want to run. If you choose the extended parameter, you also must use the extendedTestID parameter and one of the extended test options.

Extended Test Option for Fibre Channel	Description
EDC	This option tests the Error Detection and Correction (EDC) generation, verification, and deletion functionality of the QE4 chip. This option tests all modes of the EDC operation, such as, insert, verify, and delete EDC data.
DMA	This option tests the capability of the QE4 chip to take part in a Direct Memory Access (DMA) operation. The DMA can be internal to the chip or can be performed using the services of the raw pool within the Reconfigurable Processor Assembly (RPA) memory.

Extended Test Option for iSCSI	Description
RAM	This option performs a read/write test for the local RAM, the SRAM, and also performs a checksum test for the NVRAM. This option performs the read/write test for the RAM and SRAM by writing data to the memory, reading back the data, and comparing the read data to the written data.
internalLoopBack	<p>This option tests the ability of the physical layer (PHY) to transmit data packets over the physical link. For this test, the PHY is set to an internal loopback mode. Data is then transmitted, received, and compared with the original data. The test is run in two passes:</p> <ul style="list-style-type: none"> ■ For the first pass, the data is predefined by the firmware. ■ For the second pass, the data is generated externally and then transmitted.

Notes

You can run the diagnostic test on only one controller in the storage array at any one time.

A controller can have either one or two host interface cards.

- If a controller has one host interface card, the value for the position of each host interface card depends on the position of the controller in the controller tray. The host interface card in the controller in controller tray slot A has a position value of 1. The host interface card in the controller in controller tray slot B has a position value of 2.

- If a controller has two host interface cards, the value for the position of each host interface card depends on the position of the host interface card in the controller and the position of the controller in the controller tray. In most cases the position of the host interface card is identified with labels such as Host Card 1 and Host Card 2 on each controller. The position value of the host interface cards are listed in this table.

Controller	Host Card Label	Position
A	Host Card 1 C	1
	Host Card 2	2
B	Host Card 1	3
	Host Card 2	4

You cannot use a loopback connection for the host interface card that you are testing.

Minimum Firmware Level

7.70 adds the capability for controller host interface card diagnostics.

Start iSCSI DHCP Refresh

This command initiates a refresh of the DHCP parameters for the iSCSI interface. If the configuration method for the interface is not set to DHCP, the procedure returns an error.

Syntax

```
start controller [(a | b)] iscsiHostPort [(1 | 2 | 3 | 4)]
dhcpRefresh
```

Parameter

Parameter	Description
controller	The identifier letter of the controller that has the iSCSI host ports. Valid controller identifier values are a or b, where a is the controller in slot A, and b is the controller in slot B.
iscsiHostPort	The identifier of the iSCSI port for which you want to refresh the DHCP parameters. Enclose the iSCSI host port identifier in square brackets ([]).

Notes

This operation ends the iSCSI connections for the portal and temporarily brings down the portal.

Minimum Firmware Level

7.10

Start Secure Drive Erase

This command erases all of the data from one or more full disk encryption (FDE) drives so that they can be reused as FDE drives. Run this command only when the FDE drives are no longer part of a secure volume group or disk pool, or when the security key is unknown.

Syntax

```
start secureErase (drive [trayID,slotID] |
drives [trayID1,slotID1 ... trayIDn,slotIDn])
```

Parameters

Parameter	Description
drive or drives	The tray and the slot where the drive resides. Tray ID values are 0 to 99. Slot ID values are 1 to 32. Enclose the tray ID values and the slot ID values in square brackets ([]).

Notes

The controller firmware creates a lock that restricts access to the FDE drives. FDE drives have a state called Security Capable. When you create a security key, the state is set to Security Enabled, which restricts access to all FDE drives that exist within the storage array.

Minimum Firmware Level

7.40

Start Snapshot (Legacy) Rollback

NOTE With firmware version 7.83 the order of the terms in the syntax is changed to be consistent with other commands. Replace the deprecated command syntax with the new command syntax.

This command starts a rollback operation for one or more snapshot (legacy) volumes. The content of the base volume changes immediately to match the point-in-time content of the selected snapshot (legacy) volume. The base volume immediately becomes available for read/write requests after the rollback operation has successfully completed. To stop a snapshot (legacy) rollback operation, use the `stop rollback volume` command.

The repository volume that is associated with the snapshot (legacy) continues to track any new changes between the base volume and the snapshot (legacy) volume that occur after the rollback operation is completed.

Syntax (Deprecated)

```
start rollback (volume [snapshotVolumeName |  
volumes ["snapshotVolumeName1" ... "snapshotVolumeNameN"])
```

Syntax (New)

```
start (volume [snapshotVolumeName |  
volumes ["snapshotVolumeName1" ... "snapshotVolumeNameN"])  
rollback
```

Parameter

Parameter	Description
volume or volumes	The name of the specific snapshot (legacy) volume or snapshot (legacy) volumes for which you want to start a rollback operation. Enclose the snapshot (legacy) volume name in square brackets ([]). If the snapshot (legacy) volume name has special characters, you must also enclose the snapshot (legacy) volume name in double quotation marks (" "). You can enter more than one volume name. Enclose all of the volume names in one set of square brackets ([]). Enclose each volume name in double quotation marks (" "). Separate each volume name with a white space.

Minimum Firmware Level

7.80

7.83 changes the order of the terms in the syntax.

Start Snapshot Image Rollback

This command starts a rollback operation for a set of snapshot images. The content of the base volume changes immediately to match the point-in-time content of the selected snapshot image volume. The base volume immediately becomes available for read/write requests after the rollback operation has successfully completed. To stop a snapshot image rollback operation, use the `stop rollback snapImage` command.

The repository volume that is associated with the snapshot image continues to track any new changes between the base volume and the snapshot image volume that occur after the rollback operation is completed.

NOTE You cannot use this command for snapshot images involved in online volume copy.

Syntax

```
start snapImage ["snapImageName"] rollback
```

Parameter

Parameter	Description
snapImage	<p>The name of the snapshot image. The name of a snapshot image is comprised of two parts:</p> <ul style="list-style-type: none"> ■ The name of the snapshot group ■ An identifier for the snapshot image in the snapshot group <p>The identifier for the snapshot image can be one of these:</p> <ul style="list-style-type: none"> ■ An integer value that is the sequence number of the snapshot in the snapshot group. ■ NEWEST - Use this option when you want to show the latest snapshot image created in the snapshot group. ■ OLDEST - Use this option when you want to show the earliest snapshot image created in the snapshot group. <p>Enclose the snapshot image name in double quotation marks (" ") inside square brackets ([]).</p>

Notes

The name of a snapshot image has two parts separated by a colon (:):

- The identifier of the snapshot group
- The identifier of the snapshot image

For example, if you want to start a rollback operation for snapshot image 12345 in a snapshot group that has the name snapGroup1, you would use this command:

```
start snapImage ["snapGroup1:12345"] rollback;
```

To start a rollback operation for the most recent snapshot image in a snapshot group that has the name snapGroup1, you would use this command:

```
start snapImage ["snapGroup1:newest"]rollback;
```

Minimum Firmware Level

7.83

Start SSD Cache Performance Modeling

This command starts performance modeling for the SSD cache. Performance modeling monitors and measures I/O activity for a period of time and estimates performance for various SSD cache sizes. Performance is estimated using two metrics: cache hit percentage and average response time. The performance modeling data is not available until you stop the operation using the `stop ssdCache performanceModeling` command.

Syntax

```
start ssdCache [ssdCacheName] performanceModeling
```

Parameters

Parameter	Description
ssdCache	The alphanumeric identifier (including - and _) of the SSD cache that you want to locate. Enclose the identifier in square brackets ([]). If the SSD cache name contains special characters or consists only of numbers, you also must enclose the identifier in double quotation marks (" ") inside square brackets.

Notes

Performance modeling ends and the performance modeling data is available when one of the following conditions occurs:

- Run the `stop ssdCache performanceModeling` command.
- Retrieve the performance modeling data using SANtricity ES Storage Manager.

Performance modeling ends, but no data is available when one of the following conditions occurs:

- You reboot the controller.
- You make any changes to the SSD cache configuration.
- The state of the SSD cache changes.

Minimum Firmware Level

7.84

Start Storage Array Configuration Database Diagnostic

This command runs a consistency check against a configuration database.

Syntax

```
start storageArray configDbDiagnostic  
[sourceLocation=(disk | onboard) |  
diagnosticType=(fileSystem | mirror) |  
controller[(a | b)]]
```

Parameters

Parameter	Description
sourceLocation	<p>This parameter specifies the location of the database.</p> <ul style="list-style-type: none"> ■ <code>disk</code> indicates that data comes directly from the database on the drive ■ <code>onboard</code> indicates that data comes from the RPA memory location <p>The default location is <code>disk</code>.</p>
diagnosticType	<p>The level of diagnostic testing that you want to run on the database. You can run one of these levels of testing:</p> <p><code>fileSystem</code> – This option check the the structural integrity of the database .</p> <p><code>mirror</code></p> <ul style="list-style-type: none"> ■ When the <code>sourceLocation</code> parameter is set to <code>disk</code>, the peer controller initiates a block check. ■ When the <code>sourceLocation</code> parameter is set to <code>onboard</code>, the peer controller initiates a record check. <hr/> <p>NOTE You can run the <code>mirror</code> option only from a command line or from the Script editor. This option is not available through the storage management software GUI. The <code>mirror</code> option starts a long running operation that you cannot stop.</p> <hr/> <p>The default value is <code>fileSystem</code>.</p>
controller	<p>The controller that has the database on which you want to run the diagnostic tests. Valid controller identifiers are <code>a</code> or <code>b</code>, where <code>a</code> is the controller in slot A, and <code>b</code> is the controller in slot B. Enclose the controller identifier in square brackets ([]).</p>

Notes

This command runs a consistency check against a configuration database. All database records are checked. Error data is written to a file in the data folder on disk automatically. You do not need to specify an output file.

NOTE Running a consistency check with the `diagnosticType` parameter set to `mirror` and with the `sourceLocation` parameter set to `onboard` can cause the operation to run for a long time. This can have adverse effects on host I/O processing. This operation should be done only under direction from the support organization.

Upon completion of the diagnostic test, the controller firmware returns one of these results:

- Diagnosis completed without errors. No ZIP file created.
- Diagnosis completed with errors. Refer to the ZIP file created at:

```
... \Install_dir\data\FirmwareUpgradeReports\timestamp_buildNo.zip
```

If the diagnostic test detects an inconsistency in the configuration database, the controller firmware performs these actions:

- Returns a description of the inconsistency
- Saves a ZIP file containing raw binary data

The controller firmware saves the ZIP file to this location:

```
... \Install_dir\data\FirmwareUpgradeReports\timestamp_buildNo.zip
```

You can use the binary data to help determine the cause of the problem, or you can send the file containing the binary data to a Technical Support Representative.

To stop the database configuration diagnostic test, use the `stop storageArray configDbDiagnostic` command.

In addition, you can start the database configuration diagnostic test through the storage management software GUI; however, you cannot stop the database configuration diagnostic test through the storage management software GUI. If you want to stop a running diagnostic test, you must use the `stop storageArray configDbDiagnostic` command.

Minimum Firmware Level

7.75

7.83 adds these parameters:

- `sourceLocation`
- `diagnosticType`
- `controller`

Start Storage Array Core Dump

This command produces a storage array core dump.

Syntax

```
start storageArray coreDump controller [(a|b)]
```

Parameters

Parameter	Description
controller	This parameter specifies the controller from which to produce the core dump. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the storage management software returns a syntax error.

Notes

This command forces the selected controller to dump its core data to cache. Use the `save storageArray coreDump` command to save a copy of the cache data to a host file.

Minimum Firmware Level

7.83

Start Storage Array iSNS Server Refresh

This command initiates a refresh of the network address information for the iSNS server. If the DHCP server is marginal or unresponsive, the refresh operation can take from two to three minutes to complete.

NOTE This command is for IPv4 only.

Syntax

```
start storageArray isnsServerRefresh
```

Parameter

None.

Notes

If you used the `set storageArray isnsIPv4ConfigurationMethod` command to set the configuration but did not set the configuration to DHCP, running the `start storageArray isnsServerRefresh` returns an error.

Minimum Firmware Level

7.10

Start Storage Array Locate

This command locates a storagearray by turning on the indicator lights for the storage array. (Use the `stop storageArray locate` command to turn off the indicator lights for the storage array.)

Syntax

```
start storageArray locate
```

Parameters

None.

Minimum Firmware Level

6.10

Start Synchronous Mirroring Synchronization

This command starts Synchronous Mirroring synchronization.

NOTE In previous versions of this command the feature identifier was `remoteMirror`. This feature identifier is no longer valid and is replaced by `syncMirror`.

Syntax

```
start syncMirror primary ["volumeName"] synchronize
```

Parameter

Parameter	Description
primary	The name of the primary volume for which you want to start synchronization. Enclose the primary volume name in double quotation marks (" ") inside of square brackets ([]).

Minimum Firmware Level

6.10

Start Tray Locate

This command locates a tray by turning on the indicator light. (Use the `stop tray locate` command to turn off the indicator light for the tray.)

Syntax

```
start tray [trayID] locate
```

Parameter

Parameter	Description
tray	The tray that you want to locate. Tray ID values are 0 to 99. Enclose the tray ID value in square brackets ([]).

Minimum Firmware Level

6.10

Start Volume Group Defragment

This command starts a defragment operation on the specified volume group.

NOTE Defragmenting a volume group starts a long-running operation that you cannot stop.

Syntax

```
start volumeGroup [volumeGroupName] defragment
```

Parameter

Parameter	Description
volumeGroup	The alphanumeric identifier of the volume group (including - and _) that you want to defragment. Enclose the volume group identifier in square brackets ([]).

Notes

Host I/O errors might result in the volume groups with more than 32 volumes. This operation also might result in internal controller reboots because the timeout period ends before the volume group definition is set. If you experience this issue, quiesce the host I/O operations, and try the command again.

Minimum Firmware Level

6.10

Start Volume Group Export

This command moves a volume group into an Exported state. Then you can remove the drives that comprise the volume group and reinstall the drives in a different storage array.

NOTE Within the volume group, you cannot move volumes that are associated with the premium features from one storage array to another storage array.

Syntax

```
start volumeGroup [volumeGroupName] export
```

Parameter

Parameter	Description
volumeGroup	The alphanumeric identifier of the volume group (including - and _) that you want to export. Enclose the volume group identifier in square brackets ([]).

Notes

When this command is successful, you can run the `start volumeGroup import` command to finish moving the volume group to a Complete state, which makes the volume group available to the new storage array.

If this command is unsuccessful because hardware problems prevented the completion of the export, use the `set volumeGroup forceState` command. The `set volumeGroup forceState` command lets you use the `start volumeGroup import` command to import a volume group.

After the volume group is in an Exported state or a Forced state, you can remove the drives that comprise the volume group from the storage array. You can reinstall the drives in a different storage array.

Minimum Firmware Level

7.10

Start Volume Group Import

This command moves a volume group into a Complete state to make a newly introduced volume group available to its new storage array. The volume group must be in an Exported state or a Forced state before you run this command. Upon successfully running the command, the volume group is operational.

NOTE Within the volume group, you cannot move volumes that are associated with the premium features from one storage array to another storage array.

Syntax

```
start volumeGroup ["volumeGroupName"] import
```

Parameter

Parameter	Description
volumeGroup	The alphanumeric identifier of the volume group (including - and _) that you want to import. Enclose the volume group identifier in double quotation marks (" ") inside square brackets ([]).

Notes

Higher-level volumes that are specifically related to premium features (Snapshot (Legacy), Synchronous Mirroring, Volume Copy, mapping, and persistent reservations) are removed as part of the import operation.

You must run the `show volumeGroup importDependencies` command before you run the `start volumeGroup import` command.

Minimum Firmware Level

7.10

Start Volume Group Locate

This command identifies the drives that are logically grouped together to form the specified volume group by blinking the indicator lights on the drives. (Use the `stop volumeGroup locate` command to turn off the indicator lights on the drives.)

Syntax

```
start volumeGroup [volumeGroupName] locate
```

Parameter

Parameter	Description
volumeGroup	The alphanumeric identifier of the volume group (including - and _) for which you want to locate the drives that belong to that volume group. Enclose the volume group identifier in square brackets ([]).

Minimum Firmware Level

6.16

Start Volume Initialization

This command starts the formatting of a volume in a storage array.

NOTE Formatting a volume starts a long-running operation that you cannot stop.

Syntax

```
start volume [volumeName] initialize
```

Parameter

Parameter	Description
volume	The name of the specific volume for which you are retrieving information. Enclose the volume name in square brackets ([]). If the volume name has special characters, you also must enclose the volume name in double quotation marks (" ") inside square brackets. If the volume name consists only of numbers, such as "1002," you also must enclose the volume name in double quotation marks (" ") inside square brackets.

Minimum Firmware Level

6.10

Stop Cache Backup Device Diagnostic

This command stops the cache backup device diagnostic tests that were started by the `start cacheBackupDevice diagnostic` command.

Syntax

```
stop cacheBackupDevice controller [(a | b)] diagnostic
```

Parameters

Parameter	Description
controller	The controller that has the cache backup device on which you are running the diagnostic tests. Valid controller identifiers are <code>a</code> or <code>b</code> , where <code>a</code> is the controller in slot A, and <code>b</code> is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the storage management software returns a syntax error.

Minimum Firmware Level

7.60 adds the capability for cache backup device diagnostics.

Stop Cache Memory Diagnostic

This command stops the cache memory diagnostic tests that were started by the `start cacheMemory diagnostic` command.

Syntax

```
stop cacheMemory controller [(a | b)] diagnostic
```

Parameter

Parameter	Description
controller	The controller that has the cache memory on which you are running the diagnostic tests. Valid controller identifiers are <code>a</code> or <code>b</code> , where <code>a</code> is the controller in slot A, and <code>b</code> is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the storage management software returns a syntax error.

Minimum Firmware Level

7.60 adds the capability for cache memory diagnostics.

Stop Consistency Group Snapshot Rollback

This command stops a rollback operation to the member base volumes in a snapshot consistency group that was initiated by the `start cgSnapImage rollback` command.

NOTE Canceling a consistency group snapshot rollback operation leaves the base volume in an indeterminate state with potentially invalid or inconsistent data. The related consistency group snapshot volume becomes disabled and unusable.

Syntax

```
stop cgSnapImage[ "snapCGID:imageID" ] | rollback  
[memberVolumeSet [ "memberVolumeName1" ...  
"memberVolumeNameN" ]]
```

Parameter

Parameter	Description
cgSnapImage	<p>The name of the consistency group snapshot image for which you want to stop a rollback operation. The name of a snapshot image is comprised of two parts:</p> <ul style="list-style-type: none">■ The name of the snapshot group■ An identifier for the snapshot image in the snapshot group <p>The identifier for the snapshot image can be one of these:</p> <ul style="list-style-type: none">■ An integer value that is the sequence number of the snapshot in the snapshot group.■ NEWEST - Use this option when you want to show the latest snapshot image created in the snapshot group.■ OLDEST - Use this option when you want to show the earliest snapshot image created in the snapshot group. <p>Enclose the snapshot image name in double quotation marks (" ") inside square brackets ([]).</p>
memberVolumeSet	<p>The name of one or more member base volumes in a consistency group that you want to stop a rollback operation. Enclose each member base volume name in double quotation marks (" ") inside parentheses (()).</p> <p>You can enter more than one member base volume name. Enclose all of the member base volume names in one set of square brackets ([]). Enclose each member base volume name in double quotation marks (" "). Separate each member base volume name with a white space.</p> <p>When the memberVolumeSet parameter is not used stopping the rollback process applies to all member volumes of the consistency group</p>

Notes

The name of a snapshot image has two parts separated by a colon (:):

- The identifier of the snapshot group
- The identifier of the snapshot image

For example, if you want to stop a roll back operation for the newest snapshot image in an entire consistency group that has the name CG1, you would use this command:

```
stop cgSnapImage ["CG1:newest"] rollback;
```

To stop a roll back operation for the snapshot image 12345 for base volume members memVol1, memVol2, and memVol3 in a consistency group that has the name CG2, you would use this command:

```
stop cgSnapImage ["CG2:12345"] rollback  
memberVolumeset=( "memVol1 memVol2 memVol3" );
```

Minimum Firmware Level

7.83

Stop Consistency Group Snapshot Volume

This command stops the copy-on-write operation for creating a consistency group snapshot volume. To restart the copy-on-write operation use the `resume cgSnapVolume` command.

NOTE This command does not delete a consistency group snapshot volume. To delete a consistency group snapshot volume use the `delete cgSnapVolume` command.

Syntax

```
stop cgSnapVolume [snapVolumeName]
```

Parameter

Parameter	Description
<code>cgSnapVolume</code>	The name of the consistency group snapshot volume that you want to stop creating. Enclose the name of the consistency group snapshot volume in square brackets ([]). The name of the consistency group snapshot volume that you want to delete. Enclose the consistency group snapshot volume name in square brackets ([]). If the consistency group snapshot volume name has special characters, you also must enclose the snapshot volume name in double quotation marks (" ").

Minimum Firmware Level

7.83

Stop Controller Diagnostic

This command stops the controller diagnostic tests that were started by the `start controller diagnostic` command.

Syntax

```
stop controller [(a | b)] diagnostic
```

Parameters

Parameter	Description
controller	The setting to return information about a specific controller in the storage array. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the storage management software returns a syntax error.

Minimum Firmware Level

7.70

Stop Disk Pool Locate

This command turns off the indicator lights on the drives that were turned on by the `start diskPool locate` command.

Syntax

```
stop diskPool locate
```

Parameters

None.

Minimum Firmware Level

7.83

Stop Drive Channel Fault Isolation Diagnostics

This command stops the drive channel fault isolation diagnostics, which stops the `start driveChannel fault isolation diagnostics` command before it completes.

Syntax

```
stop driveChannel faultDiagnostics
```

Parameters

None.

Notes

Use the `start driveChannel faultDiagnostics` command and the `save driveChannel faultDiagnostics` command with the `stop driveChannel faultDiagnostics` command. These commands are needed to start the diagnostic test and save diagnostic test results to a file.

You can also stop the `start driveChannel faultDiagnostics` command at any time by pressing **Ctrl+C**.

Minimum Firmware Level

7.15

Stop Drive Channel Locate

This command turns off the indicator lights on the drive trays that were turned on by the `start driveChannel locate` command.

Syntax

```
stop driveChannel locate
```

Parameters

None.

Minimum Firmware Level

6.10

Stop Drive Locate

This command turns off the indicator light on the drive that was turned on by the `start drive locate` command.

Syntax

```
stop drive locate
```

Parameters

None.

Minimum Firmware Level

6.10

Stop Host Interface Card Diagnostic

This command stops the host interface card diagnostic tests that were started by the `start host card diagnostic` command.

Syntax

```
stop host card controller [(a | b)] diagnostic
```

Parameters

Parameter	Description
<code>controller</code>	The controller that has the host interface card on which you are running the diagnostic tests. Valid controller identifiers are <code>a</code> or <code>b</code> , where <code>a</code> is the controller in slot A, and <code>b</code> is the controller in slot B. Enclose the controller identifier in square brackets (<code>[]</code>). If you do not specify a controller, the storage management software returns a syntax error.

Minimum Firmware Level

7.70 adds the capability for controller host interface card diagnostics.

Stop Pending Snapshot Images on Consistency Group

This command stops all of the pending snapshot images that are to be created on a snapshot consistency group. If the snapshot consistency group does not have any pending snapshot images, the storage management software displays an error message and does not run the command.

Syntax

```
stop consistencyGroup [consistencyGroupName]  
pendingSnapImageCreation
```

Parameters

Parameter	Description
<code>consistencyGroupName</code>	The name of the consistency group for which you want to stop any pending snapshot operation. Enclose the name of the consistency group in square brackets ([]).

Minimum Firmware Level

7.83

Stop Snapshot Group Pending Snapshot Images

This command cancels all of the pending snapshot images that are to be created on a snapshot group. If the snapshot group does not have any pending snapshot images, the firmware displays an error message and does not run the command. You can run this command on a snapshot group or a snapshot consistency group.

Syntax

```
stop (snapGroup [snapGroupName] |  
consistencyGroup [snapConsistencyGroupName])  
pendingSnapImageCreation
```

Parameters

Parameter	Description
<code>snapGroup</code>	The alphanumeric identifier (including - and _) of the snapshot group for which you want to stop pending snapshot images. Enclose the snapshot group identifier in square brackets ([]).
<code>consistencyGroup</code>	The alphanumeric identifier (including - and _) of the snapshot consistency group for which you want to stop pending snapshot images. Enclose the snapshot group identifier in square brackets ([]).

Minimum Firmware Level

7.83

Stop Snapshot (Legacy)

This command stops a copy-on-write operation.

Syntax

```
stop snapshot (volume [volumeName] |  
volumes ["volumeName1" ... "volumeNameN"])
```

Parameter

Parameter	Description
volume or volumes	The name of the specific volume for which you want to stop a copy-on-write operation. You can enter more than one volume name. You can enter more than one volume name. Enclose all of the volume names in one set of square brackets ([]). Enclose each volume name in double quotation marks (" "). Separate each volume name with a white space.

Notes

Names can be any combination of alphanumeric characters, underscore (_), hyphen (-), and pound (#). Names can have a maximum of 30 characters.

One technique for naming the snapshot (legacy) volume and the snapshot (legacy) repository volume is to add a hyphenated suffix to the original base volume name. The suffix distinguishes between the snapshot (legacy) volume and the snapshot (legacy) repository volume. For example, if you have a base volume with a name of Engineering Data, the snapshot (legacy) volume can have a name of Engineering Data-S1, and the snapshot (legacy) repository volume can have a name of EngineeringData-R1.

If you do not choose a name for either the snapshot (legacy) volume or the snapshot (legacy) repository volume, the storage management software creates a default name by using the base volume name. An example of the snapshot (legacy) volume name that the controllers might create is, if the base volume name is aaa and does not have a snapshot (legacy) volume, the default snapshot (legacy) volume name is aaa-1. If the base volume already has $n-1$ number of snapshot (legacy) volumes, the default name is aaa- n . An example of the snapshot (legacy) repository volume name that the controller might create is, if the base volume name is aaa and does not have a snapshot (legacy) repository volume, the default snapshot (legacy) repository volume name is aaa-R1. If the base volume already has $n-1$ number of snapshot (legacy) repository volumes, the default name is aaa-R n .

Enclose the volume names using one of these forms:

- On a Windows command line: `"volumeName"`
- In a Windows script engine window: `["volumeName"]`
- On a Linux command line: `\`volumeName\``
- In a Linux script engine window: `["volumeName"]`

Minimum Firmware Level

6.10

Stop Snapshot (Legacy) Rollback

This command stops a snapshot (legacy) rollback operation that was initiated by the `start rollback volume` command.

NOTE Canceling a rollback operation leaves the base volume in an indeterminate state with potentially invalid or inconsistent data. The related snapshot (legacy) volume becomes disabled and unusable.

Syntax

```
stop rollback volume [snapshotVolumeName]
```

Parameter

Parameter	Description
volume	The name of the specific snapshot (legacy) volume for which you want to stop a rollback operation. Enclose the snapshot (legacy) volume name in square brackets ([]). If the snapshot (legacy) volume name has special characters, you must also enclose the snapshot (legacy) volume name in double quotation marks (" ").

Minimum Firmware Level

7.80

Stop Snapshot Image Rollback

This command stops a snapshot image rollback operation that was initiated by the `start snapImage rollback` command.

NOTE Canceling a snapshot image rollback operation leaves the base volume in an indeterminate state with potentially invalid or inconsistent data. The related snapshot image volume becomes disabled and unusable.

Syntax

```
stop snapImage [snapCGID: imageID] rollback
```

Parameter

Parameter	Description
snapImage	<p>The name of the snapshot image for which you want to stop a rollback operation. The name of a snapshot image is comprised of two parts:</p> <ul style="list-style-type: none">■ The name of the snapshot group■ An identifier for the snapshot image in the snapshot group <p>The identifier for the snapshot image can be one of these:</p> <ul style="list-style-type: none">■ An integer value that is the sequence number of the snapshot in the snapshot group.■ NEWEST - Use this option when you want to show the latest snapshot image created in the snapshot group.■ OLDEST - Use this option when you want to show the earliest snapshot image created in the snapshot group. <p>Enclose the snapshot image name in double quotation marks (" ") inside square brackets ([]).</p>

Notes

The name of a snapshot image has two parts separated by a colon (:):

- The identifier of the snapshot group
- The identifier of the snapshot image

For example, if you want to stop a rollback operation for snapshot image 12345 in a snapshot group that has the name snapGroup1, you would use this command:

```
stop snapImage [ "snapGroup1:12345" ] rollback;
```

To stop a rollback operation for the most recent snapshot image in a snapshot group that has the name snapGroup1, you would use this command:

```
stop snapImage [ "snapGroup1:newest" ] rollback;
```

Minimum Firmware Level

7.83

Stop Snapshot Volume

This command stops a snapshot volume operation.

Syntax

```
stop snapVolume [ "snapVolumeName" ]
```

Parameters

Parameter	Description
snapVolume	The name of the snapshot volume that you want to stop. Enclose the snapshot volume name in double quotation marks (" ") inside square brackets ([]).

Notes

You can use any combination of alphanumeric characters, underscore (_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

Minimum Firmware Level

7.83

Stop SSD Cache Performance Modeling

This command stops the performance modeling operation and displays the performance modeling data for the SSD cache. Before running this command, you must start the performance modeling operation with the `start ssdCache performanceModeling` command. Optionally, you can save the data to a file.

Syntax

```
stop ssdCache [ssdCacheName] performanceModeling
[file="filename"]
```

Parameters

Parameter	Description
ssdCache	The alphanumeric identifier (including - and _) of the SSD cache that you want to locate. Enclose the identifier in square brackets ([]). If the SSD cache name contains special characters or consists only of numbers, you also must enclose the identifier in double quotation marks (" ") inside square brackets.
file	The file path and the file name to which you want to save the performance modeling data. Enclose the file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\logs\performance.csv" The default name of the file that contains the performance modeling data is <code>readLinkStatus.csv</code> . You can use any file name, but you must use the <code>.csv</code> extension.

Notes

This command returns the performance modeling information similar to this example. The size of your monitor determines how the information wraps and will affect how the information appears.

```

SSD Cache Name: my_cache
Start time: 4/18/12 2:38:26 PM IST
Stop time: 4/18/12 2:38:45 PM IST
Duration : 00:00:19

```

SSD Cache Performance Modeling Data (Response Time):

Reads	SSD Reads		HDD		A
	HDD Writes				
Overall	Avg.		Avg.		
vg.	Response	Response	% of	Response	%
Cache	Response	Time	I/Os	Time	I/Os
of	Time	% of			
Capacity	I/Os				
186 GB	0 ms	0 ms	0.0 %	0 ms	0.0 %
ms	0.0 %				
372 GB	0 ms	0 ms	0.0 %	0 ms	0.0 %
ms	0.0 %				
557 GB *	0 ms	0 ms	0.0 %	0 ms	0.0 %
ms	0.0 %				
558 GB	0 ms	0 ms	0.0 %	0 ms	0.0 %
ms	0.0 %				
744 GB	0 ms	0 ms	0.0 %	0 ms	0.0 %
ms	0.0 %				
931 GB	0 ms	0 ms	0.0 %	0 ms	0.0 %
ms	0.0 %				
1117 GB	0 ms	0 ms	0.0 %	0 ms	0.0 %
ms	0.0 %				
1303 GB	0 ms	0 ms	0.0 %	0 ms	0.0 %
ms	0.0 %				
1489 GB	0 ms	0 ms	0.0 %	0 ms	0.0 %
ms	0.0 %				
1675 GB	0 ms	0 ms	0.0 %	0 ms	0.0 %
ms	0.0 %				
1862 GB	0 ms	0 ms	0.0 %	0 ms	0.0 %
ms	0.0 %				

* = Current SSD cache physical capacity.

SSD Cache Performance Modeling Data (Cache Hit %):

Cache Capacity	Cache Hit %
186 GB	0 %
372 GB	0 %

557 GB *	0 %
558 GB	0 %
744 GB	0 %
931 GB	0 %
1117 GB	0 %
1303 GB	0 %
1489 GB	0 %
1675 GB	0 %
1862 GB	0 %

* = Current SSD cache physical capacity.

Minimum Firmware Level

7.84

Stop Storage Array Configuration Database Diagnostic

This command stops the diagnostic test to validate the configuration database in the controller firmware that was started by the `start storageArray configDbDiagnostic` command.

Syntax

`stop storageArray configDbDiagnostic`

Parameters

None.

Notes

The controller firmware returns a confirmation that the diagnostic test was cancelled.

In addition, you can start the database configuration diagnostic test through the storage management software GUI; however, you cannot stop the database configuration diagnostic test through the storage management software GUI. If you want to stop a running diagnostic test, you must use the `stop storageArray configDbDiagnostic` command.

If you try to use the `stop storageArray configDbDiagnostic` command after validation of the storage array configuration has finished, you do not receive any message that the validation has finished. This behavior is expected.

Minimum Firmware Level

7.75

7.77 refines usage.

Stop Storage Array Drive Firmware Download

This command stops a firmware download to the drives in a storage array that was started with the `download storageArray driveFirmware` command. This command does not stop a firmware download that is already in progress to a drive. This command stops all firmware downloads to drives that are waiting for the download.

Syntax

```
stop storageArray driveFirmwareDownload
```

Parameters

None.

Minimum Firmware Level

6.10

Stop Storage Array iSCSI Session

This command forces the termination of a storage array iSCSI session.

Syntax

```
stop storageArray iscsiSession [sessionNumber]
```

Parameter

Parameter	Description
iscsiSession	The identifier number of the iSCSI session. Enclose the identifier number of the iSCSI session in square brackets ([]).

Minimum Firmware Level

7.10

Stop Storage Array Locate

This command turns off the indicator lights on the storage array that were turned on by the `start storageArray locate` command.

Syntax

```
stop storageArray locate
```

Parameters

None.

Minimum Firmware Level

6.10

Stop Tray Locate

This command turns off the indicator light on the tray that was turned on by the `start tray locate` command.

Syntax

```
stop tray locate
```

Parameters

None.

Minimum Firmware Level

6.10

Stop Volume Copy

This command stops a volume copy operation. This command is valid for both snapshot (legacy) volume copy pairs and new snapshot volume copy pairs.

Syntax

```
stop volumeCopy target [targetName] source [sourceName]
```

Parameters

Parameter	Description
target	The name of the target volume for which you want to stop a volume copy operation. Enclose the target volume name in square brackets ([]). If the target volume name has special characters, you also must enclose the target volume name in double quotation marks (" ").
source	The name of the source volume for which you want to stop a volume copy operation. Enclose the source volume name in square brackets ([]). If the source volume name has special characters, you also must enclose the source volume name in double quotation marks (" ").

Minimum Firmware Level

5.40

Stop Volume Group Locate

This command turns off the indicator lights on the drives that were turned on by the `start volumeGroup locate` command.

Syntax

```
stop volumeGroup locate
```

Parameters

None.

Minimum Firmware Level

6.16

Suspend Asynchronous Mirror Group

This command suspends the synchronization of data on all mirrored pairs at the asynchronous mirror group level. This `suspend` command helps to reduce any performance impact to the host application that might occur while any changed data on the local storage array is copied to the remote storage array.

Syntax

```
suspend asyncMirrorGroup ["asyncMirrorGroupName"]
```

Parameters

Parameter	Description
asyncMirrorGroup	The name of the asynchronous mirror group that you want to suspend. Enclose the asynchronous mirror group name in square brackets ([]). If the asynchronous mirror group name has special characters, you also must enclose the asynchronous mirror group name in double quotation marks (" ") inside square brackets. If the asynchronous mirror group name consists only of numbers, such as "1002," you also must enclose the asynchronous mirror group name in double quotation marks (" ") inside square brackets.

Notes

When an asynchronous mirror group is in a suspended state, no attempt is made to copy data from the primary volumes to the secondary volumes of the mirrored pairs. Any writes to the primary side of the asynchronous mirror group are persistently logged in its associated mirror repository volumes. After the asynchronous mirror group is resumed, only the modified regions of the primary volumes are written to the secondary volumes.

Minimum Firmware Level

7.84

Suspend SSD Cache

This command temporarily stops caching for all of the volumes that are using the SSD cache. While caching is stopped, host reads are serviced from the base volumes instead of from the SSD cache.

Syntax

```
suspend ssdCache [ssdCacheName]
```

Parameter

Parameter	Description
ssdCache	The alphanumeric identifier (including - and _) of the SSD cache that you want to suspend. Enclose the identifier in square brackets ([]). If the SSD cache name contains special characters or consists only of numbers, you also must enclose the identifier in double quotation marks (" ") inside square brackets.

Notes

To restart caching, use the `resume ssdCache` command.

Suspend Synchronous Mirroring

Minimum Firmware Level

7.84

This command suspends a Synchronous Mirroring operation.

NOTE In previous versions of this command the feature identifier was `remoteMirror`. This feature identifier is no longer valid and is replaced by `syncMirror`.

Syntax

```
suspend syncMirror (primary [primaryVolumeName]  
primaries ["primaryVolumeName1" ... "primaryVolumeNameN"])  
writeConsistency=(TRUE | FALSE)
```

Parameters

Parameter	Description
<code>primary</code> or <code>primaries</code>	The name of the primary volume for which you want to suspend operation. Enclose the volume name in square brackets ([]). If the volume name has special characters, you must also enclose the volume name in double quotation marks (" "). You can enter more than one primary volume name. Enclose all of the primary volume names in one set of square brackets ([]). Enclose each primary volume name in double quotation marks (" "). Separate each primary volume name with a white space.
<code>writeConsistency</code>	This parameter defines whether the volumes identified in this command are in a write-consistency group or are separate. For the volumes in the same write-consistency group, set this parameter to TRUE. For the volumes that are separate, set this parameter to FALSE.

Notes

If you set the `writeConsistency` parameter to TRUE, the volumes must be in a write-consistency group (or groups). This command suspends all write-consistency groups that contain the volumes. For example, if volumes A, B, and C are in a write-consistency group and they have remote counterparts A', B', and C', the command:

```
suspend syncMirror volume ["A"] writeConsistency=TRUE
```

suspends A-A', B-B', and C-C'. If you have a write-consistency group 1={A, B, C} and write-consistency group 2={D, E, F}, the command:

```
suspend syncMirror volumes ["A" "D"] writeConsistency=TRUE
```

suspends both write-consistency groups.

Minimum Firmware Level

6.10

Test Alert Severities

NOTE This command is an SMcli command, not a script command. You must run this command from a command line. You cannot run this command from the script editor in the storage management software

This command sends out a test alert to the Windows Event Log and all configured syslog receivers.

Syntax

```
SMcli -alertTest
```

Parameters

None.

Minimum Firmware Level

7.83

Test Asynchronous Mirror Group Connectivity

This command tests possible communication problems between the local storage array and the remote the storage array associated with an asynchronous mirror group.

Syntax

```
diagnose asyncMirrorGroup ["asyncMirrorGroupName"]
testID=(all | connectivity | latency | bandwidth |
portConnections)
```

Parameter

Parameter	Description
asyncMirrorGroup	The name of an existing asynchronous mirror group that you want to test. Enclose the asynchronous mirror group name in square brackets ([]). If the asynchronous mirror group name has special characters, you also must enclose the asynchronous mirror group name in double quotation marks (" ").

Parameter	Description
testID	<p>The identifier for the diagnostic test you want to run. The identifier and corresponding tests are as follow:</p> <ul style="list-style-type: none"> ■ All— Performs all the tests associated with this command. ■ Connectivity — Verifies that the two controllers have a communication path. The connectivity test sends an inter-controller message between the storage arrays, and then validates that the corresponding asynchronous mirror group on the remote storage array exists. It also validates that the volume members of the asynchronous mirror group on the remote system match the volume members of the asynchronous mirror group on the local system. ■ Latency — Sends a SCSI test unit command to each volume on the remote storage array associated with the asynchronous mirror group to test the minimum, average, and maximum latency. ■ Bandwidth — Sends two inter-controller messages to the remote storage array to test the minimum, average, and maximum bandwidth as well as the negotiated link speed of the port on the controller performing the test. ■ Port connections — Shows the port that is being used for mirroring on the local storage array and the port that is receiving the mirrored data on the remote storage array.

Minimum Firmware Level

7.84

Validate Storage Array Security Key

This command validates the security key for a storage array that has full disk encryption (FDE) drives to make sure that the security key is not corrupt.

Syntax

```
validate storageArray securityKey
file="fileName"
passPhrase="passPhraseString"
```

Parameters

Parameter	Description
file	The file path and the file name that has the security key. Enclose file path and the file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\sup\seckey.slk" IMPORTANT – The file name must have an extension of .slk.
passPhrase	A character string that encrypts the security key so that you can store the security key in an external file. Enclose the pass phrase in double quotation marks (" ").

Notes

Your pass phrase must meet these criteria:

- The pass phrase must be between eight and 32 characters long.
- The pass phrase must contain at least one uppercase letter.
- The pass phrase must contain at least one lowercase letter.
- The pass phrase must contain at least one number.
- The pass phrase must contain at least one non-alphanumeric character, for example, < > @ +.

NOTE If your pass phrase does not meet these criteria, you will receive an error message.

Minimum Firmware Level

7.70

Deprecated Commands and Parameters

This appendix lists the commands, the command formats, and the parameters that are no longer supported by this level of software. The information is presented in two tables. The table in the section "[Deprecated Commands](#)" lists commands that are no longer supported in this level of software and the new commands that replaced them. The table in the section "[Deprecated Parameters](#)" lists the parameters that are no longer supported in this level of software and the new parameters that replaced them.

Deprecated Commands

Table 4 Commands Deprecated in Firmware Release 10.86

Deprecated Command	New Command
No commands were deprecated with this release of software.	

Deprecated Parameters

Table 5 Parameters Deprecated in Firmware Release 10.86

Deprecated Parameter	Revision
modificationPriority	This parameter is removed from this command: <ul style="list-style-type: none"> ■ Set Volume Attributes for a Disk Pool

Copyright © 2012 NetApp, Inc. All rights reserved.

