

# Quantum<sup>®</sup>

## SANtricity 11.30 CLI/Script Commands Programming Guide

# StorNext QD7000

Firmware 8.30.xx.xx



SANtricity 11.30 CLI/Script Commands Programming Guide, 6-68578-01 Rev A, April 2017 Product of USA.

Quantum Corporation provides this publication "as is" without warranty of any kind, either express or implied, including but not limited to the implied warranties of merchantability or fitness for a particular purpose. Quantum Corporation may revise this publication from time to time without notice.


#### **COPYRIGHT STATEMENT**

© 2017 Quantum Corporation. All rights reserved.

Your right to copy this manual is limited by copyright law. Making copies or adaptations without prior written authorization of Quantum Corporation is prohibited by law and constitutes a punishable violation of the law.

#### **TRADEMARK STATEMENT**

Artico, Be Certain (and the Q brackets design), DLT, DXi, DXi Accent, DXi V1000, DXi V2000, DXi V4000, DXiV-Series, FlexTier, Lattus, the Q logo, the Q Quantum logo, Q-Cloud, Quantum (and the Q brackets design), the Quantum logo, Quantum Be Certain (and the Q brackets design), Quantum Vision, Scalar, StorageCare, StorNext, SuperLoader, Symform, the Symform logo (and design), vmPRO, and Xcellis are either registered trademarks or trademarks of Quantum Corporation and its affiliates in the United States and/or other countries. All other trademarks are the property of their respective owners. Products mentioned herein are for identification purposes only and may be registered trademarks or trademarks of their respective companies. All other brand names or trademarks are the property of their respective owners. Quantum specifications are subject to change.



# Preface

---

---

**Note:** The 8.30.xx.xx firmware (Lehigh) is used in the QD7000 (E5600, Titan RAID controller, only). Refer to the [NetApp to Quantum Naming Decoder](#) section for additional information.

---

This section provides the following information:

- [Audience](#)
- [Prerequisites](#)
- [NetApp to Quantum Naming Decoder](#)
- [Product Safety Statements](#)
- [Contacts](#)
- [Comments](#)
- [Quantum Global Services](#)

---

## Audience

---

This manual is intended for storage customers and technicians.

## Prerequisites

Prerequisites for installing and using this product include knowledge of:





- Servers and computer networks
- Network administration
- Storage system installation and configuration
- Storage area network (SAN) management and direct attach storage (DAS)
- Fibre Channel (FC) and Ethernet protocols

## NetApp to Quantum Naming Decoder

Use [Table 1](#) to correlate the NetApp product nomenclature to the equivalent Quantum-storage naming conventions.

Table 1 Product Nomenclature

<b>E-Series NetApp Product</b>	<b>Quantum-Storage</b>	<b>Description</b>
Controller-Drive Tray	Base System	Quantum uses Base System when referring to a drive tray with the RAID controllers.
Drive Tray	Expansion Unit	Quantum uses Expansion Unit when referring to a drive tray with the environmental services modules (ESMs).
E5600 (Code Name: Titan)	RAID Controller	Four 16Gb/s FC SFP+ host ports
E5500 (Code Name: Soyuz)	RAID Controller	Four 16Gb/s FC SFP+ host ports
E5400 (Code Name: Pikes Peak)	RAID Controller	Four 8Gb/s FC SFP+ host ports
DE6600 (Code Name: Wembley)	4U 60-drive enclosure	Sixty 3.5 inch disk drives

E-Series NetApp Product	Quantum-Storage	Description
<p>E5560 or E5660 (DE6600 4U drive enclosure with E5500 or E5600 RAID controllers)</p>	<p>Quantum StorNext QD7000</p>	
<p>E5460 (DE6600 4U drive enclosure with E5400 RAID controllers)</p>	<p>Quantum StorNext QD6000</p>	
<p>E5424 (DE5600 24-drive 2U drive enclosure (Code Name: Camden with E5400 RAID controllers)</p>	<p>Quantum StorNext QS2400</p>	
<p>E5412 (DE1600 12-drive 2U drive enclosure (Code Name: Ebbets with E5400 RAID controllers)</p>	<p>Quantum StorNext QS1200</p>	

---

## Product Safety Statements

---

Quantum will not be held liable for damage arising from unauthorized use of the product. The user assumes all risk in this aspect.

This unit is engineered and manufactured to meet all safety and regulatory requirements. Be aware that improper use may result in bodily injury, damage to the equipment, or interference with other equipment.

---

**WARNING:** Before operating this product, read all instructions and warnings in this document and in the system, safety, and regulatory guide.

---

---

**警告** 在使用本产品之前，请先阅读本文档及系统、安全和法规信息指南中所有的说明和警告信息。

---

---

**警告** 操作本產品前，請先閱讀本文件及系統、安全與法規資訊指南中的指示與警告說明。

---

---

**ADVERSAL** Læs alle instruktioner og advarsler i dette dokument og i *Vejledning om system-sikkerheds- og lovgivningsoplysninger*, før produktet betjenes.

---

---

**AVERTISSEMENT** Avant d'utiliser ce produit, lisez la totalité des instructions et avertissements de ce document et du *Guide d'informations sur le système, la sécurité et la réglementation*.

---

---

**HINWEIS** Lesen Sie vor der Verwendung dieses Produkts alle Anweisungen und Warnhinweise in diesem Dokument und im System, Safety, and Regulatory Information Guide (Info-Handbuch: System, Sicherheit und Richtlinien).

---

---

לפני ההפעלה של מוצר זה, קרא את כל ההוראות והאזהרות הכלולות במסמך זה וכן במדריך מידע בנושאי מערכת, בטיחות ותקינה

---

**אזהרה**

**警告**

この製品を使用する前に、本文書、および『システム、安全、規制に関する情報ガイド』に記載しているすべての警告と指示をお読みください。

**경고**

이 제품을 작동하기 전에 이 문서 및 시스템, 안전, 및 규제 정보 안내서에 수록된 모든 지침과 경고 표지를 숙지하십시오.

**ПРЕДУПРЕЖДЕНИЕ**

Перед началом эксплуатации данного устройства ознакомьтесь во всеми инструкциями и предупреждениями, приведенными в данном документе и в *Справочном руководстве по устройству, технике безопасности и действующим нормативам*.

**ADVERTENCIA**

Antes de utilizar este producto, lea todas las instrucciones y advertencias en este documento y en la Guía informativa sobre sistema, seguridad y normas.

**VARNING**

Läs alla anvisningar och varningar i detta dokument och i *System, säkerhet och krav från myndigheter - Informationshandbok* innan denna produkt tas i bruk.

---

**Contacts**


---

For information about contacting Quantum, including Quantum office locations, go to:

<http://www.quantum.com/aboutus/contactus/index.aspx>

---

**Comments**


---

To provide comments or feedback about this document, or about other Quantum technical publications, send e-mail to:

[doc-comments@quantum.com](mailto:doc-comments@quantum.com)

---

## Quantum Global Services

---

Accelerate service issue resolution with these exclusive Quantum StorageCare services:

- **Service and Support Website** - Register products, license software, browse Quantum Learning courses, check backup software and operating system support, and locate manuals, FAQs, firmware downloads, product updates and more in one convenient location. Benefit today at:

<http://www.quantum.com/serviceandsupport/get-help/index.aspx#contact-support>

- **eSupport** - Submit online service requests, update contact information, add attachments, and receive status updates via email. Online Service accounts are free from Quantum. That account can also be used to access Quantum's Knowledge Base, a comprehensive repository of product support information. Get started at:

<http://www.quantum.com/customercenter/>

For further assistance, or if training is desired, contact the Quantum Customer Support Center:

<b>North America</b>	1-800-284-5101 (toll free) +1-720-249-5700
<b>EMEA</b>	+800-7826-8888 (toll free) +49-6131-324-185
<b>APAC</b>	+800-7826-8887 (toll free) +603-7953-3010

For worldwide support:

<http://www.quantum.com/serviceandsupport/get-help/index.aspx#contact-support>

---





SANtricity® 11.30

# Command Line Interface and Script Commands Programming Guide

August 2016 | 215-11213\_A0  
doccomments@netapp.com

 **NetApp**®



# Contents

<b>Deciding whether to use this guide .....</b>	<b>13</b>
<b>Formatting the commands .....</b>	<b>14</b>
Structure of a CLI command .....	14
Interactive mode .....	15
CLI command wrapper syntax .....	15
Command line parameters .....	17
Structure of a script command .....	21
Synopsis of the script commands .....	22
Recurring syntax elements .....	24
Naming conventions .....	31
Formatting CLI commands .....	32
Formatting rules for script commands .....	32
Formatting CLI commands in Windows PowerShell .....	34
Usage guidelines .....	34
Detailed error reporting .....	35
Exit status .....	36
Adding comments to a script file .....	37
Firmware compatibility levels .....	38
<b>Using curl commands with the E2800 storage array .....</b>	<b>39</b>
<b>Command reference .....</b>	<b>45</b>
Commands listed by function .....	45
Asynchronous mirroring commands .....	45
Auto-load balancing commands .....	46
AutoSupport bundle collection commands .....	46
Consistency group commands .....	46
Controller commands .....	47
Controller health image commands .....	47
Database commands .....	48
Disk pool commands .....	48
Drive commands .....	48
Event alert commands .....	49
Feature management commands .....	49
Host topology commands .....	49
Input output controller (IOC) core dump commands .....	49
Internet storage name service (iSNS) commands .....	49
iSCSI commands .....	50
iSCSI extension for RDMA (iSER) commands .....	50
Repository volume analysis commands .....	50
Session command .....	50
Snapshot group commands .....	50
Snapshot image commands .....	50

Snapshot volume commands .....	51
SSD cache commands .....	51
Storage array commands .....	51
Synchronous Mirroring commands .....	53
Tray commands .....	54
Uncategorized commands .....	54
Volume commands .....	54
Volume commands for disk pools .....	54
Volume copy commands .....	55
Volume group commands .....	55
Storage management .....	55
Managing an E2800 storage array .....	56
Managing multiple E2800 storage arrays .....	56
Management of mixed storage arrays .....	57
Event Management .....	58
Event management for E2800 storage arrays .....	60
Event management for the E2700, E5600, and EF560 .....	61
AutoSupport Overview .....	62
About AutoSupport and AutoSupport OnDemand .....	62
AutoSupport for E2800 storage arrays .....	64
AutoSupport for E2700, E5600, and EF560 storage arrays .....	65
Management Methods in SANtricity 11.30 .....	66
Terminology differences in SANtricity 11.30 .....	70
Commands listed alphabetically .....	72
Activate asynchronous mirroring .....	72
Activate iSCSI initiator .....	72
Activate storage array firmware .....	73
Activate Synchronous Mirroring .....	74
Add drives to SSD cache .....	77
Add member to consistency group .....	78
Add volume to asynchronous mirror group .....	80
Autoconfigure storage array .....	82
Autoconfigure storage array hot spares .....	88
Cancel asynchronous mirror group role reversal .....	88
Capture or view an AutoSupport log .....	89
Change SSD cache application type .....	90
Check asynchronous mirror group consistency .....	91
Check repository consistency .....	92
Check storage array connectivity .....	94
Check synchronous mirroring status .....	95
Check volume parity .....	95
Clear asynchronous mirroring fault .....	96
Clear drive channel statistics .....	98
Clear storage array configuration .....	99
Clear storage array controller health image .....	100

Clear storage array event log .....	101
Clear storage array firmware pending area .....	101
Clear storage array recovery mode .....	101
Clear volume reservations .....	102
Clear volume unreadable sectors .....	103
Create asynchronous mirror group .....	104
Create consistency group .....	107
Create consistency group snapshot image .....	115
Create consistency group snapshot volume .....	116
Create consistency group snapshot volume mapping .....	120
Create disk pool .....	121
Create host .....	127
Create host group .....	128
Create host port .....	129
Create iSCSI initiator .....	131
Create RAID volume (automatic drive select) .....	132
Create RAID volume (free extent based select) .....	137
Create RAID volume (manual drive select) .....	140
Create read-only snapshot volume .....	145
Create snapshot group .....	146
Create snapshot image .....	155
Create snapshot volume .....	156
Create SSD cache .....	159
Create storage array security key .....	161
Create synchronous mirroring .....	162
Create volume copy .....	165
Create volume group .....	167
Create volume in disk pool .....	172
Deactivate asynchronous mirroring .....	176
Deactivate synchronous mirroring .....	176
Delete asynchronous mirror group .....	177
Delete consistency group .....	178
Delete consistency group snapshot image .....	179
Delete consistency group snapshot volume .....	180
Delete disk pool .....	181
Delete host .....	182
Delete host group .....	182
Delete host port .....	183
Delete iSCSI initiator .....	184
Delete snapshot group .....	185
Delete snapshot image .....	186
Delete snapshot volume .....	187
Delete SSD cache .....	188
Delete volume .....	189
Delete volume from disk pool .....	190

Delete volume group .....	191
Diagnose controller .....	192
Diagnose controller iSCSI host cable .....	193
Diagnose synchronous mirroring .....	196
Disable AutoSupport at the EMW level SMcli version .....	197
Disable storage array feature .....	197
Display AutoSupport bundle collection settings .....	198
Display automatic support bundle collection configuration .....	200
Display AutoSupport message collection schedule .....	200
Display automatic support bundle collection schedule .....	202
Download drive firmware .....	202
Download environmental card firmware .....	204
Download storage array drive firmware .....	204
Download storage array firmware/NVSRAM .....	205
Download storage array NVSRAM .....	207
Download tray configuration settings .....	207
Enable controller data transfer .....	208
Enable disk pool security .....	209
Enable or disable ODX .....	209
Enable or disable AutoSupport at the EMW management domain level (SMcli only) .....	210
Enable or disable the AutoSupport OnDemand feature at the EMW management domain level (SMcli only) .....	211
Enable or disable the AutoSupport OnDemand Remote Diagnostics feature at the EMW management domain level (SMcli only) .....	213
Enable or disable VAAI .....	214
Enable storage array feature .....	214
Enable volume group security .....	216
Establish asynchronous mirrored pair .....	216
Export storage array security key .....	218
Import storage array security key .....	218
Increase capacity of volume in disk pool or volume group .....	219
Initialize thin volume .....	221
Load storage array DBM database .....	223
Recopy volume copy .....	224
Recover disabled drive ports .....	225
Recover RAID volume .....	225
Recover SAS port mis-wire .....	229
Re-create synchronous mirroring repository volume .....	230
Reduce disk pool capacity .....	233
Register SNMP community .....	234
Register SNMP trap destination .....	234
Remove drives from SSD cache .....	235
Remove incomplete asynchronous mirrored pair from asynchronous mirror group .....	236

Remove member volume from consistency group .....	237
Remove synchronous mirroring .....	238
Remove volume copy .....	239
Remove volume from asynchronous mirror group .....	239
Remove volume LUN mapping .....	240
Rename snapshot volume .....	242
Rename SSD cache .....	242
Repair volume parity .....	243
Replace drive .....	244
Reset asynchronous mirror group statistics .....	246
Reset AutoSupport message collection schedule .....	247
Reset controller .....	248
Reset drive .....	249
Reset iSCSI IP address .....	250
Reset storage array diagnostic data .....	251
Reset storage array Infiniband statistics baseline .....	252
Reset storage array iSCSI baseline .....	252
Reset storage array iSER baseline .....	253
Reset storage array RLS baseline .....	254
Reset storage array SAS PHY baseline .....	254
Reset storage array SOC baseline .....	254
Reset storage array volume distribution .....	255
Resume asynchronous mirror group .....	255
Resume consistency group snapshot volume .....	256
Resume snapshot image rollback .....	258
Resume snapshot volume .....	259
Resume SSD cache .....	260
Resume synchronous mirroring .....	261
Revive drive .....	262
Revive snapshot group .....	263
Revive snapshot volume .....	263
Revive volume group .....	264
Save asynchronous mirror group statistics .....	265
Save controller NVSRAM .....	267
Save drive channel fault isolation diagnostic status .....	267
Save drive log .....	268
Save input output controller (IOC) dump .....	269
Save auto-load balancing statistics .....	270
Save storage array configuration .....	270
Save storage array controller health image .....	272
Save storage array DBM database .....	273
Save storage array DBM validator information file .....	274
Save storage array diagnostic data .....	275
Save storage array events .....	276
Save storage array firmware inventory .....	278

Save storage array InfiniBand statistics .....	278
Save storage array iSCSI statistics .....	279
Save storage array iSER statistics .....	280
Save storage array performance statistics .....	281
Save storage array RLS counts .....	281
Save storage array SAS PHY counts .....	282
Save storage array SOC counts .....	282
Save storage array state capture .....	283
Save storage array support data .....	283
Save tray log .....	289
Schedule automatic support bundle collection configuration .....	289
Set asynchronous mirror group .....	294
Set consistency group attributes .....	296
Set consistency group snapshot volume .....	298
Set controller .....	299
Set controller DNS settings .....	304
Set controller NTP settings .....	305
Set controller service action allowed indicator .....	306
Set disk pool .....	307
Set disk pool (modify disk pool) .....	310
Set drawer service action allowed indicator .....	312
Set drive channel status .....	313
Set drive hot spare .....	314
Set drive service action allowed indicator .....	314
Set drive state .....	315
Set event alert filtering .....	316
Set FIPS drive security identifier .....	318
Set foreign drive to native .....	319
Set host .....	320
Set host channel .....	321
Set host group .....	322
Set host port label .....	323
Set iSCSI initiator .....	324
Set iSCSI target properties .....	326
Set iSER target .....	327
Set read-only snapshot volume to read/write volume .....	328
Set session .....	330
Set snapshot group attributes .....	331
Set snapshot group media scan .....	333
Set snapshot group repository volume capacity .....	334
Set snapshot group schedule .....	336
Set snapshot volume media scan .....	339
Set snapshot volume repository volume capacity .....	340
Set SSD cache for a volume .....	342
Set storage array .....	344



Set storage array AutoSupport bundle disable .....	348
Set storage array AutoSupport bundle enable .....	348
Set storage array controller health image allow overwrite .....	348
Set storage array to enable or disable Automatic Load Balancing .....	349
Set storage array ICMP response .....	350
Set storage array iSNS registration .....	351
Set storage array iSNS server IPv4 address .....	353
Set storage array iSNS server IPv6 address .....	354
Set storage array iSNS server listening port .....	354
Set storage array iSNS server refresh .....	355
Set storage array controller battery learn cycle .....	356
Set storage array redundancy mode .....	357
Set storage array security key .....	358
Set storage array time .....	358
Set storage array tray positions .....	359
Set storage array unnamed discovery session .....	359
Set synchronous mirroring .....	360
Set thin volume attributes .....	362
Set tray identification .....	364
Set tray service action allowed indicator .....	365
Set volume attributes for a volume in a disk pool .....	366
Set volume attributes for a volume in a volume group .....	371
Set volume copy .....	377
Set volume group .....	378
Set volume group forced state .....	380
Set volume mapping .....	380
Show asynchronous mirror groups .....	382
Show asynchronous mirror group synchronization progress .....	383
Show blocked events .....	384
Show consistency group .....	384
Show consistency group snapshot image .....	386
Show controller .....	387
Show controller NVSRAM .....	390
Show current iSCSI sessions .....	390
Show disk pool .....	392
Show drive .....	393
Show drive channel statistics .....	395
Show drive download progress .....	396
Show drive performance statistics .....	396
Show host ports .....	398
Show replaceable drives .....	399
Show snapshot group .....	399
Show snapshot image .....	401
Show snapshot volumes .....	403
Show SNMP communities .....	404

Show SNMP MIB II system group variables .....	405
Show SSD cache .....	406
Show SSD cache statistics .....	407
Show storage array .....	410
Show storage array auto configuration .....	414
Show storage array controller health image .....	419
Show storage array DBM database .....	420
Show storage array host topology .....	421
Show storage array LUN mappings .....	422
Show storage array negotiation defaults .....	423
Show storage array ODX setting .....	423
Show storage array power information .....	424
Show storage array unconfigured iSCSI initiators .....	424
Show storage array unreadable sectors .....	425
Show string .....	426
Show synchronous mirroring volume candidates .....	426
Show synchronous mirroring volume synchronization progress .....	427
Show thin volume .....	428
Show volume .....	429
Show volume action progress .....	431
Show volume copy .....	431
Show volume copy source candidates .....	432
Show volume copy target candidates .....	432
Show volume group .....	433
Show volume group export dependencies .....	434
Show volume group import dependencies .....	435
Show volume performance statistics .....	435
Show volume reservations .....	437
Specify the AutoSupport delivery method .....	438
Start asynchronous mirroring synchronization .....	440
Start consistency group snapshot rollback .....	443
Start controller trace .....	444
Start disk pool locate .....	446
Start drive channel fault isolation diagnostics .....	447
Start drive channel locate .....	449
Start drive initialize .....	449
Start drive locate .....	450
Start drive reconstruction .....	451
Start input output controller (IOC) dump .....	452
Start iSCSI DHCP refresh .....	453
Start FDE secure drive erase .....	455
Start snapshot image rollback .....	456
Start SSD cache locate .....	457
Start SSD cache performance modeling .....	458
Start storage array configuration database diagnostic .....	459

Start storage array controller health image .....	460
Start storage array iSNS server refresh .....	461
Start storage array locate .....	462
Start synchronous mirroring synchronization .....	462
Start tray locate .....	463
Start volume group defragment .....	463
Start volume group export .....	464
Start volume group import .....	464
Start volume group locate .....	465
Start volume initialization .....	466
Stop consistency group snapshot rollback .....	466
Stop consistency group snapshot volume .....	468
Stop disk pool locate .....	469
Stop drive channel fault isolation diagnostics .....	469
Stop drive channel locate .....	470
Stop drive locate .....	470
Stop drive replace .....	471
Stop pending snapshot images on consistency group .....	471
Stop snapshot group pending snapshot images .....	472
Stop snapshot image rollback .....	473
Stop snapshot volume .....	474
Stop SSD cache locate .....	475
Stop SSD cache performance modeling .....	476
Stop storage array configuration database diagnostic .....	478
Stop storage array drive firmware download .....	478
Stop storage array iSCSI session .....	479
Stop storage array locate .....	479
Stop tray locate .....	480
Stop volume copy .....	480
Stop volume group locate .....	481
Suspend asynchronous mirror group .....	481
Suspend SSD cache .....	482
Suspend synchronous mirroring .....	482
Test alerts .....	484
Test asynchronous mirror group connectivity .....	484
Test the AutoSupport configuration .....	485
Test SNMP trap destination .....	487
Unregister SNMP community .....	487
Unregister SNMP trap destination .....	488
Update SNMP community .....	488
Update SNMP MIB II system group variables .....	488
Update SNMP trap destination .....	489
Validate storage array security key .....	490
Changes in the CLI and script commands in this release .....	491
<b>Copyright information .....</b>	<b>496</b>

<b>Trademark information .....</b>	<b>497</b>
<b>How to send comments about documentation and receive update notifications .....</b>	<b>498</b>
<b>Changes in the CLI and script commands in this release .....</b>	<b>499</b>

## Deciding whether to use this guide

---

This guide is intended for system operators, system administrators, and service personnel who are responsible for configuring, managing, and maintaining a storage array running SANtricity Storage Manager 11.30. Use this guide if you install and manage storage arrays.

This guide does not provide extensive conceptual background or information about how to install and configure the SANtricity Storage Manager software using the command line interface. Refer to the [SANtricity 11.30 Configuring a Storage Array Using the Command Line Interface](#) to learn how to configure a storage array. Refer to the [SANtricity 11.30 Maintaining a Storage Array Using the Command Line Interface](#) to learn how to maintain a storage array.

This guide is based on the following assumptions:

- You have a knowledge of basic storage area network (SAN) hardware and installation skills.
- You understand storage technology, Redundant Array of Independent Disks (RAID) concepts, and networking technologies.
- You have a basic knowledge of SAN hardware functionality (controllers, drives, and hosts) and SAN cabling.

### Where to Find the Latest Information About the Product

You can find information about the latest version of the product, including new features and fixed issues, and a link to the latest documentation at the following link: [E-Series Documentation Center](#).

### Related information

[NetApp E-Series and EF-Series Systems Documentation Center](#)

[SANtricity 11.30 Maintaining a Storage Array Using the Command Line Interface](#)

[SANtricity 11.30 Configuring a Storage Array Using the Command Line Interface](#)

## Formatting the commands

---

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 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.

Some software features described in this document might not be available for your NetApp E-Series Storage System. For questions about available features, contact your NetApp account representative.

The CLI is a feature of the SANtricity Storage Manager software. To access this product, go to the NetApp Support Site at [support.netapp.com](http://support.netapp.com).

## 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` is a CLI value that defines 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 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>	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

**Attention:** To run all of the CLI commands you must have administrator privileges. Some CLI commands will run without administrator privileges. Many commands, however, will not run. If the CLI command does not run because you do not have correct privileges, the CLI returns an exit code of 12.

## Examples

The following examples demonstrate the command line parameters described in *Command line parameters* on page 17.

```
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]
```

**Restriction:** The `-a` command line option is not supported for the E2800 storage array. Use the RESTful API for script support and the SANtricity System Manager for a graphical user interface.

```
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]
```

**Restriction:** The `-x` command line option is not supported for the E2800 storage array. Use the RESTful API for script support and the SANtricity System Manager for a graphical user interface.

```
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]
```

**Restriction:** The `-a` and `-x` command line options are not supported for the E2800 storage array. Use the RESTful API for script support and the SANtricity System Manager for a graphical user interface.

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

**Restriction:** The `-s` command line option is not supported for the E2800 storage array. Use the RESTful API for script support and the SANtricity System Manager for a graphical user interface.

```
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 parameters

Parameter	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"> <li>• If you are managing a storage array by using a host through in-band storage management, you must use the <i>-n</i> parameter or the <i>-w</i> parameter if more than one storage array is connected to the host.</li> <li>• 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.</li> <li>• 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> parameter.</li> <li>• 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> parameter.</li> </ul>
<i>-A</i>	<p>Adds a storage array to the configuration file. If you do not follow the <i>-A</i> parameter 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"> <li>• 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.</li> <li>• 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.</li> </ul> <p><b>Restriction:</b> This command line option is obsolete for the E2800 storage array. Use the RESTful API, SANtricity System Manager, or see <a href="#">Using curl commands with the E2800 storage array</a> on page 39.</p>
<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 (;). You cannot place more than one <i>-c</i> parameter on the same command line. You can include more than one script command after the <i>-c</i> parameter.</p>
<i>-d</i>	<p>Shows the contents of the script configuration file. The file content has this format: <i>storage-system-name host-name1 host-name2</i></p>
<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>

Parameter	Definition
-f (lowercase)	<p>Specifies a file name that contains script commands that you want to run on the specified storage array. The -f parameter is similar to the -c parameter in that both parameters are intended for running script commands. The -c parameter runs individual script commands. The -f parameter 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 parameter 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 parameter with these parameters:</p> <ul style="list-style-type: none"> <li>• -a</li> <li>• -x</li> </ul>
-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"> <li>• <code>eventOnly</code> – Only the event information is included in the email.</li> <li>• <code>profile</code> – The event and array profile information is included in the email.</li> </ul> <p>You can specify the frequency for the email deliveries using the -q parameter.</p>
-i (lowercase)	<p>Shows the IP address of the known storage arrays. Use the -i parameter with the -d parameter. The file content 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> <p><b>Restriction:</b> This command line option is obsolete for the E2800 storage array. Use the RESTful API, SANtricity System Manager, or see <a href="#">Using curl commands with the E2800 storage array</a> on page 39.</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 parameter 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>

Parameter	Definition
-o	<p>Specifies a file name for all output text that is a result of running the script commands. Use the -o parameter with these parameters:</p> <ul style="list-style-type: none"> <li>• -c</li> <li>• -f</li> </ul> <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 parameter 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 parameter:</p> <ul style="list-style-type: none"> <li>• everyEvent – Information is returned with every email alert notification.</li> <li>• 2 – Information is returned no more than once every two hours.</li> <li>• 4 – Information is returned no more than once every four hours.</li> <li>• 8 – Information is returned no more than once every eight hours.</li> <li>• 12 – Information is returned no more than once every 12 hours.</li> <li>• 24 – Information is returned no more than once every 24 hours.</li> </ul> <p>Using the -I parameter you can specify the type of information in the email alert notifications.</p> <ul style="list-style-type: none"> <li>• If you set the -I parameter to eventOnly , the only valid value for the -q parameter is everyEvent .</li> <li>• If you set the -I parameter to either the profile value or the supportBundle value, this information is included with the emails with the frequency specified by the -q parameter.</li> </ul>
-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 recreate snapshot volume command. This parameter reduces time by not running background processes for the duration of the command.</p> <p>Do not use this parameter 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 parameter causes operations that depend on background information to fail.</p>

Parameter	Definition
-R (uppercase)	<p>Defines the user role for the password. The roles can be either:</p> <ul style="list-style-type: none"> <li>• <code>admin</code> – The user has privilege to change the storage array configuration.</li> <li>• <code>monitor</code> – The user has privilege to view the storage array configuration, but cannot make changes.</li> </ul> <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"> <li>• The dual password feature is not enabled on the storage array.</li> <li>• Only one <code>admin</code> role is set and the <code>monitor</code> role is not set for the storage array.</li> </ul>
-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 parameter suppresses these messages:</p> <ul style="list-style-type: none"> <li>• <code>Performing syntax check</code></li> <li>• <code>Syntax check complete</code></li> <li>• <code>Executing script</code></li> <li>• <code>Script execution complete</code></li> <li>• <code>SMcli completed successfully</code></li> </ul>
-s (lowercase)	<p>Shows the alert settings in the configuration file when used with the <code>-d</code> parameter.</p> <p><b>Restriction:</b> This command line option is obsolete for the E2800 storage array. Use the RESTful API, SANtricity System Manager, or see <a href="#">Using curl commands with the E2800 storage array</a> on page 39.</p>
-v	Shows the current global status of the known devices in a configuration file when used with the <code>-d</code> parameter.
-w	<p>Specifies the WWID of the storage array. This parameter is an alternate to the <code>-n</code> parameter. Use the <code>-w</code> parameter with the <code>-d</code> parameter to show the WWIDs of the known storage arrays. The file content has this format:</p> <pre>storage-system-name world-wide-ID IP-address1 IP-address2</pre>
-X (uppercase)	Deletes a storage array from a configuration.

Parameter	Definition
-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.  <b>Restriction:</b> This command line option is obsolete for the E2800 storage array. Use the RESTful API, SANtricity System Manager, or see <a href="#">Using curl commands with the E2800 storage array</a> on page 39.
-?	Shows usage information about the CLI commands.

## Structure of a script command

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.

Object Type	Identifier
controller	a or b
drive	For enclosures with drawers, use tray ID, drawer ID, and slot ID. Alternatively, just tray ID and slot ID.
replacementDrive	For enclosures with drawers, use tray ID, drawer ID, and slot ID. Alternatively, just tray ID and slot ID.
driveChannel	Drive channel identifier
host	User label

Object Type	Identifier
hostChannel	Host channel identifier
hostGroup	User label
hostPort	User label
iscsiInitiator	User label or iSCSI Qualified Name (IQN)
iscsiTarget	User label or IQN
storageArray	Not applicable
tray	Tray ID
volume	Volume user label or volume World Wide Identifier (WWID) (set 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.

Syntax	Description
<code>accept <i>object</i></code> <code>{<i>statement-data</i>}</code>	Performs the pending operation.
<code>activate <i>object</i></code> <code>{<i>statement-data</i>}</code>	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> <code>{<i>statement-data</i>}</code>	Automatically creates a configuration that is based on the parameters that are specified in the command.

Syntax	Description
<code>check object {statement-data}</code>	Starts an operation to report on errors in the object, which is a synchronous operation.
<code>clear object {statement-data}</code>	Discards the contents of some attributes of an object. This operation is destructive and cannot be reversed.
<code>create object {statement-data}</code>	Creates an object of the specified type.
<code>deactivate object {statement-data}</code>	Removes the environment for an operation.
<code>delete object</code>	Deletes a previously created object.
<code>diagnose object {statement-data}</code>	Runs a test and shows the results.
<code>disable object {statement- data}</code>	Prevents a feature from operating.
<code>download object {statement-data}</code>	Transfers data to the storage array or to the hardware that is associated with the storage array.
<code>enable object {statement-data}</code>	Sets a feature to operate.
<code>load object {statement-data}</code>	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 {statement-data}</code>	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 {statement-data}</code>	Re-creates an object from saved configuration data and the statement parameters. (This command is similar to the <code>create</code> command.)
<code>remove object {statement-data}</code>	Removes a relationship between objects.

Syntax	Description
<code>repair object {statement-data}</code>	Repairs errors found by the <code>check</code> command.
<code>replace object {statement-data}</code>	The specified object replaces an existing object in the storage array.
<code>reset object {statement-data}</code>	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 {statement-data}</code>	Writes information about the object to a file.
<code>set object {statement-data}</code>	Changes object attributes. All changes are completed when the command returns.
<code>show object {statement-data}</code>	Shows information about the object.
<code>start object {statement-data}</code>	Starts an asynchronous operation. You can stop some operations after they have started. You can query the progress of some operations.
<code>stop object {statement-data}</code>	Stops an asynchronous operation.
<code>suspend object {statement-data}</code>	Stops an operation. You can then restart the suspended operation, and it continues from the point where it was suspended.
<code>validate object {statement-data}</code>	Validates a security key.

## Recurring syntax elements

Recurring syntax elements are a general category of parameters and options that you can use in the script commands. The following table lists the conventions used in the recurring syntax elements.



Convention	Definition
a   b	Alternative ("a" or "b")
<i>italicized-words</i>	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

The following table lists the recurring syntax parameters and the values that you can use with the recurring syntax parameters.

Recurring Syntax	Syntax Value
<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> <sup>6</sup>   <i>hotSpareCount=integer-literal</i>   <i>segmentSize=segment-size-spec</i>   <i>cacheReadPrefetch=(TRUE   FALSE)</i> <i>securityType=(none   capable   enabled)</i> <sup>7</sup>   <i>dataAssurance=(none   enabled)</i> <sup>5</sup>
<i>boolean</i>	(TRUE   FALSE)
<i>cache-flush-modifier-setting</i>	immediate, 0, .25, .5, .75, 1, 1.5, 2, 5, 10, 20, 60, 120, 300, 1200, 3600, infinite
<i>capacity-spec</i>	<i>integer-literal</i> [KB   MB   GB   TB   Bytes]
<i>count-based-repository-spec</i>	<i>repositoryRAIDLevel =repository-raid-level</i> <i>repositoryDriveCount=integer-literal</i> [ <i>repositoryVolumeGroupUserLabel =user-label</i> ] [ <i>driveType=drive-type</i> <sup>4</sup> ] [ <i>trayLossProtect=(TRUE   FALSE)</i> <sup>1</sup> ]   [ <i>drawerLossProtect=(TRUE   FALSE)</i> <sup>2</sup> ]   [ <i>dataAssurance=(none   enabled)</i> <sup>5</sup> ]
<i>create-raid-vol-attr-value-list</i>	<i>create-raid-volume-attribute-value-pair</i> { <i>create-raid-volume-attribute-value-pair</i> }

Recurring Syntax	Syntax Value
<i>create-raid-volume-attribute-value-pair</i>	<pre>capacity=<i>capacity-spec</i>   owner=(a   b)   cacheReadPrefetch=(TRUE   FALSE)   segmentSize=<i>integer-literal</i>   usageHint=<i>usage-hint-spec</i></pre>
<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>	<pre>copyPriority=(highest   high   medium   low   lowest)   targetReadOnlyEnabled=(TRUE   FALSE)   copyType=(offline   online)   repositoryPercentOfBase=(20   40   60   120   default)   repositoryGroupPreference=(sameAsSource   otherThanSource   default)</pre>
<i>drive-media-type</i>	(HDD   SSD   unknown   allMedia) <i>HDD</i> means hard disk drive. <i>SSD</i> means solid state disk.
<i>drive-spec</i>	<i>trayID,slotID</i> or <i>trayID,drawerID,slotID</i> A drive is defined as two or three integer 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 drive-spec</i>
<i>drive-type</i>	(fibre   SATA   SAS)  <b>Note:</b> Only SAS drives are supported for firmware versions 7.86 and later.
<i>error-action</i>	(stop   continue)
<i>ethernet-port-options</i>	<pre>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)</pre>
<i>feature-identifier</i>	<b>Note:</b> Most features in SANtricity 11.30 are enabled by default.  <i>driveSecurity</i>
<i>filename</i>	<i>string-literal</i>

Recurring Syntax	Syntax Value
<i>gid</i>	<i>string-literal</i>
<i>hex-literal</i>	A literal in the range of 0x00 - 0xFF.
<i>host-card-identifier</i>	(1   2   3   4)
<i>host-type</i>	<i>string-literal</i>   <i>integer-literal</i>
<i>instance-based- repository-spec</i>	<pre>(repositoryRAIDLevel =repository-raid-level repositoryDrives= (drive-spec-list) [repositoryVolumeGroupUserLabel =user-label] [trayLossProtect=(TRUE   FALSE)<sup>1</sup>]   [drawerLossProtect=(TRUE   FALSE)<sup>2</sup>]   (repositoryVolumeGroup=user-label [freeCapacityArea=integer-literal<sup>3</sup>])</pre> <p>Specify the <code>repositoryRAIDLevel</code> parameter with the <code>repositoryDrives</code> parameter. Do not specify the RAID level or the drives with the volume group. Do not set a value for the <code>trayLossProtect</code> parameter when you specify a volume group.</p>
<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>iscsi-host-port</i>	<pre>(1   2   3   4)</pre> <p>The host port number might be 2, 3, or 4 depending on the type of controller you are using.</p>
<i>iscsi-host-port-options</i>	<pre>IPv4Address=ip-address   IPv6LocalAddress=ipv6-address   IPv6RoutableAddress=ipv6-address   IPv6RouterAddress=ipv6-address   enableIPv4=(TRUE   FALSE)   enableIPv6=(TRUE   FALSE)   enableIPv4Priority=(TRUE   FALSE)   enableIPv6Priority=(TRUE   FALSE)   IPv4ConfigurationMethod=(static   dhcp)   IPv6ConfigurationMethod=(static   auto)   IPv4GatewayIP=ip-address   IPv6HopLimit=integer   IPv6NdDetectDuplicateAddress=integer   IPv6NdReachableTime=time-interval   IPv6NdRetransmitTime=time-interval   IPv6NdTimeOut=time-interval   IPv4Priority=integer   IPv6Priority=integer   IPv4SubnetMask=ip-address   IPv4VlanId=integer   IPv6VlanId=integer   maxFramePayload=integer   tcpListeningPort=tcp-port-id   portSpeed=(autoNegotiate   1   10)</pre>

Recurring Syntax	Syntax Value
<i>iscsiSession</i>	[ <i>session-identifier</i> ]
<i>nvsram-offset</i>	<i>hex-literal</i>
<i>nvsramBitSetting</i>	<i>nvsram-mask</i> , <i>nvsram-value</i> = <i>0xhexadecimal</i> , <i>0xhexadecimal</i>   <i>integer-literal</i> The <i>0xhexadecimal</i> value is typically a value from 0x00 to 0xFF.
<i>nvsramByteSetting</i>	<i>nvsram-value</i> = <i>0xhexadecimal</i>   <i>integer-literal</i> The <i>0xhexadecimal</i> value is typically a value from 0x00 to 0xFF.
<i>portID</i>	(0-127)
<i>raid-level</i>	(0   1   3   5   6)
<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</i> =(a b)   <i>cacheReadPrefetch</i> =(TRUE   FALSE)   <i>dataAssurance</i> =(none   enabled)
<i>repository-raid-level</i>	(1   3   5   6)
<i>repository-spec</i>	<i>instance-based-repository-spec</i>   <i>count-based-repository-spec</i>
<i>segment-size-spec</i>	<i>integer-literal</i> - all capacities are in base-2.
<i>serial-number</i>	string-literal
<i>slotID</i>	(1-24)
<i>test-devices</i>	<i>controller</i> =(a b) <i>esms</i> =( <i>esm-spec-list</i> ) <i>drives</i> =( <i>drive-spec-list</i> )
<i>test-devices-list</i>	<i>test-devices</i> { <i>test-devices</i> }
<i>time-zone-spec</i>	(GMT+HH:MM   GMT-HH:MM) [ <i>dayLightSaving</i> =HH:MM]
<i>trayID-list</i>	<i>trayID</i> { <i>trayID</i> }

Recurring Syntax	Syntax Value
<i>usage-hint-spec</i>	<pre>usageHint=(multiMedia   database   fileSystem)</pre> <p>The usage hint, or expected I/O characteristics, of the volume are used by the controller to indicate an appropriate default volume segment size and dynamic cache read prefetch. For file system and database, a 128 KB segment size is used. For multimedia, a 256 KB segment size is used. All three usage hints enable dynamic cache read prefetch.</p>
<i>user-label</i>	<p><i>string-literal</i></p> <p>Valid characters are alphanumeric, the dash, and the underscore.</p>
<i>user-label-list</i>	<i>user-label</i> { <i>user-label</i> }
<i>volumeGroup-number</i>	<i>integer-literal</i>
<i>wwID</i>	<i>string-literal</i>

<sup>1</sup>For tray loss protection to work, your configuration must adhere to the following guidelines:

Level	Criteria for Tray Loss Protection	Minimum number of trays required
Disk Pool	The disk pool contains no more than two drives in a single tray.	6
RAID 6	The volume group contains no more than two drives in a single tray.	3
RAID 3 or RAID 5	Each drive in the volume group is located in a separate tray.	3
RAID 1	Each drive in a RAID 1 pair must be located in a separate tray.	2
RAID 0	Cannot achieve Tray Loss Protection.	Not applicable

<sup>2</sup>For drawer loss protection to work, your configuration must adhere to the following guidelines:

Level	Criteria for Drawer Loss Protection	Minimum number of drawers required
Disk Pool	The disk pool includes drives from all five drawers and there are an equal number of drives in each drawer. A 60-drive tray can achieve Drawer Loss Protection when the disk pool contains 15, 20, 25, 30, 35, 40, 45, 50, 55, or 60 drives.	5
RAID 6	The volume group contains no more than two drives in a single drawer.	3
RAID 3 or RAID 5	Each drive in the volume group is located in a separate drawer.	3
RAID 1	Each drive in a mirrored pair must be located in a separate drawer.	2

Level	Criteria for Drawer Loss Protection	Minimum number of drawers required
RAID 0	Cannot achieve Drawer Loss Protection.	Not applicable

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. You can have drawer loss protection without tray loss protection. You cannot have tray loss protection without drawer loss protection. If the `trayLossProtect` parameter and the `drawerLossProtect` parameter are not set to the same value, the storage array returns an error message and a storage array configuration will not be created.

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

<sup>4</sup> The default drive (drive type) is `SAS`.

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.

<sup>5</sup> The `dataAssurance` parameter relates to the Data Assurance (DA) feature.

The Data Assurance (DA) feature increases data integrity across the entire storage system. DA enables the storage array to check for errors that might occur when data is moved between the hosts and the drives. When this feature is enabled, the storage array appends error-checking codes (also known as cyclic redundancy checks or CRCs) to each block of data in the volume. After a data block is moved, the storage array uses these CRC codes to determine if any errors occurred during transmission. Potentially corrupted data is neither written to disk nor returned to the host.

If you want to use the DA feature, start with a volume group or disk pool that includes only drives that support DA. Then, create DA-capable volumes. Finally, map these DA-capable volumes to the host using an I/O interface that is capable of DA. I/O interfaces that are capable of DA include Fibre Channel, SAS, and iSER InfiniBand (iSCSI Extensions for RDMA/IB). DA is not supported by iSCSI over TCP/IP, or by the SRP InfiniBand.

**Note:** When all the required hardware and the I/O interface is DA-capable, you can set the `dataAssurance` parameter to `enabled` and then use DA with certain operations. For example, you can create a volume group that includes DA-capable drives, and then create a volume within that volume group that is DA-enabled. Other operations that use a DA-enabled volume have options to support the DA feature.

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

<sup>7</sup> The `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 backslashes (\\) 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.

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

**Note:** If you have any doubt as to the validity of a name, use both double quotation marks and square brackets. Using both makes sure that the name works and does not cause any processing issues.

## 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 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;
```

## Formatting CLI commands in Windows PowerShell

The Windows PowerShell is an interactive and scripting shell that provides access to command-line tools. The Windows PowerShell improves upon the Windows Command Prompt with a more robust set of commands and scripting capability. You can run all of the CLI and script commands in the Windows PowerShell; however, the Windows PowerShell has some unique formatting requirements. The requirements are these:

- Begin all SMcli commands with a period and a forward slash (./)
- SMcli wrapper must be identified as an executable command with the .exe extension (SMcli.exe)
- Enclose the script command in single quotation marks ('')
- Double quotation marks that are part of a name, file path, or value must have a backslash before each double quotation mark character (\")

The following is an example of a CLI command to create a storage array name in the Windows Powershell. Note the use of the single quotation marks as delimiters for the script command and the backslash double quotation marks around the storage array name, identified as `userLabel` in the command syntax.

```
PS C:\...\StorageManager\client> ./SMcli.exe 123.45.67.88 123.45.67.89
-c 'set storageArray userLabel="\Engineering\";'
```

The following is an example of a CLI command to enable a premium feature in the Windows Powershell. Note the use of the backslash double quotation marks before the file path to the premium feature key.

```
PS C:\...\StorageManager\client> ./SMcli.exe 123.45.67.88 123.45.67.89
-c 'enable storageArray feature file="\C:\licenseKey.bin\";'
```

In the previous examples both upper case letters and lower case letters are used. This use is to help make clear how the commands are formatted. The Windows PowerShell is, however, not case sensitive and does not require the use specific cases.

## 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](#) on page 37.)

**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 `excp rpt . txt`. The CLI tries to place the `excp rpt . 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 `excp rpt . txt` file in the same directory from which the CLI is running. You cannot change the file name or the location. The `excp rpt . txt` file is overwritten every time that an exception occurs. If you want to save the information in the `excp rpt . 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. This exit status indicates that you do not have privileges for running a CLI command from the command line. You must have administrator privileges to run all of the CLI commands from a command line.
13	Invalid script syntax was found.
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.

Status Value	Meaning
29	The sender contact information file was not found.
30	The sender contact information file could not be read.
31	The <code>userdata.txt</code> file exists.
32	An invalid <code>-I</code> value in the email alert notification was specified.
33	An invalid <code>-f</code> value in the email alert notification was specified.
34	The <code>-r</code> option is not supported anymore.
35	Invalid alert severity specified.
36	The operation needs either the Administrator or Monitor password to be set.
37	The operation cannot be completed because an invalid Monitor password was entered.
38	The operation cannot be completed because an invalid Administrator password was entered.
39	The password provided is exceeding the character limit.
40	The <code>-R</code> monitor is not supported for this array. Use a valid role and retry the operation.
41	An error occurred while writing to or reading from the AutoSupport configuration file. Please retry this operation again.
42	Host address or mail server address is incorrect.

## 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:

- 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;
```

- 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;
```

- 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;
```

## Firmware compatibility levels

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

8.10

This level indicates that the entire script command runs under a minimum of controller firmware version 8.10.

The `create ssdCache` command has the following section.

#### Minimum firmware level

7.84

8.20.11 - Added the `securityType` parameter.

These notations indicate that the script command and all of the parameters except `securityType` run under a minimum of controller firmware version 7.84. The `securityType` parameter runs under a minimum of controller firmware version 8.20\_M3.

## Using curl commands with the E2800 storage array

---

The E2800 storage array contains an embedded web server. You can use the curl tool on the command line for getting or sending files from the array, using the URL. In cases where an SMcli command is not supported on the E2800 storage array, the curl command line is a convenient alternative.

Curl is available for download from [curl: The curl project](#).

**Note:** Because the E2800 storage array uses the TLSv2 protocol, curl defaults to this during protocol negotiations. If the version of curl being used does not support this protocol, the communication will fail.

**Important:** These commands operate on a single E2800 storage array. For the E2700, E5600 and EF560, an enabled Event Monitor on a management station performs alert and AutoSupport configuration tasks for every storage array in the management domain.

### Useful curl command line options

The following table describes a few useful options when sending commands to the E2800 storage array. For a complete list of curl options, see the curl documentation online using the link in the Related references section.

**Note:** If you want the parameter to contain spaces, you must enclose the entire parameter within double quotes ("). Within those quotes, you specify a quote as \".

**Note:** You must specify options and their arguments on the same line.

**Note:** Wherever you see angle brackets "<>" in the examples or descriptions below, substitute the appropriate value and do not include the brackets.

Option	Description
d	This option is used to pass data to the E2800 storage array. <b>Note:</b> POST data must be URL-encoded.

Option	Description
H	<p>There are several types of headers used in HTTP requests. Two types of headers used for managing the E2800 storage array are:</p> <ul style="list-style-type: none"> <li>• Accept header - the HTTP client uses this header to tell the E2800 storage array what content types are acceptable.</li> <li>• Content-Type header - the storage array sends a response that signifies the content type of the returned data.</li> </ul> <p>When you send data to the E2800 storage array (for example in a POST request), you need both headers. In this case the Content-Type header signifies the content type of the posted data.</p> <p>To tell the E2800 storage array which type of content to return, set the Content-Type header using the -H option. For example, to receive JSON data:</p> <pre data-bbox="456 695 1377 751">-H "Accept: application/json"</pre> <p>To tell the E2800 storage array which type of content to return, set the Content-Type header using the -H option. For example, to receive JSON data:</p> <pre data-bbox="456 867 1377 924">-H "Content-Type: application/json"</pre> <p><b>Note:</b> Because curl allows multiple -H options on a request, you can specify more than one header in different formats.</p>
k	<p>The k option is required to allow curl to run against a system that has a self-signed certificate and to suppress certificate verification.</p> <p><b>Note:</b> For systems that have a real CA signed certificate, this option is not required because curl performs certificate authentication against a provider.</p>
O	<p>This option allows results to be saved to a file. The file is saved to the current working directory with the same filename as the filename on the server.</p>
o	<p>This option allows results to be saved to a file with a custom filename. Follow the option with the filename of your choice.</p>
u	<p>Specify the user name and password as arguments to this option, as shown in the example below:</p> <pre data-bbox="456 1455 1377 1512">-u &lt;rw:password&gt;</pre> <p><b>Note:</b> The user name for all E-Series storage arrays is rw.</p>
v	<p>This option is the verbose option, useful for debugging.</p>



Option	Description
x	<p>This option changes the HTTP request method. Available methods for HTTP 1.1 are:</p> <ul style="list-style-type: none"> <li>• GET</li> <li>• HEAD</li> <li>• POST</li> <li>• PUT</li> <li>• DELETE</li> <li>• CONNECT</li> <li>• OPTIONS</li> <li>• TRACE</li> </ul>

### Retrieve alert configuration information

The following example command gets alert configuration data from the E2800 controller.

```
curl -k -u <rw:password> -H "Accept: application/json" "https://<array IP address>/devmgr/v2/storage-systems/1/device-alerts"
```

### Post new alert configuration data

The following example command passes alert configuration data to the E2800 controller as part of the request body. Note the use of the `-d` option and the JSON data sent using keyvalue pairs.

```
curl -k -u <rw:password> -H "Accept: application/json" -H "Content-Type: application/json" -d @- "https://<array IP address>/devmgr/v2/storage-systems/1/device-alerts" <<EOF
{
  "alertingEnabled": true,
  "emailServerAddress": "smtp.corp.mydomain.com",
  "emailSenderAddress": "test@mydomain.com",
  "sendAdditionalContactInformation": true,
  "additionalContactInformation": "Additional information",
  "recipientEmailAddresses": [
    "mytestaddr@somedomain.com"
  ]
}
EOF
```

### Perform a test against the email recipients configured for alerts

The following example command tests the alert configuration by sending email to configured recipients.

```
curl -X POST -k -u <rw:password> -H "Accept: application/json" -H "Content-Type: application/json" "https://<array IP address>/devmgr/v2/storage-systems/1/device-alerts/alert-email-test"
```

### Retrieve AutoSupport configuration data

The following example command gets AutoSupport configuration data from the E2800 storage array in JSON format.

```
curl -k -u <rw:password> -H "Accept: application/json" "https://<array IP address>/devmgr/v2/device-asup"
```

### Post new AutoSupport configuration data

The following example command posts new AutoSupport configuration data to the E2800 storage array in JSON format.

```
curl -k -u <rw:password> -H "Accept: application/json" -H "Content-Type: application/json" -d @- "https://<array IP address>/devmgr/v2/device-asup" <<EOF
{
  "asupCapable": true,
  "onDemandCapable": true,
  "asupEnabled": true,
  "onDemandEnabled": true,
  "remoteDiagsEnabled": true,
  "delivery": {
    "method": "https",
    "routingType": "direct",
    "proxyHost": null,
    "proxyPort": 0,
    "proxyrw": null,
    "proxyPassword": null,
    "proxyScript": null,
    "mailRelayServer": null,
    "mailSenderAddress": null
  },
  "destinationAddress": "https://testbed.netapp.com/put/AsupPut/",
  "schedule": {
    "dailyMinTime": 0,
    "dailyMaxTime": 1439,
    "weeklyMinTime": 0,
    "weeklyMaxTime": 1439,
    "daysOfWeek": []
  }
}
EOF
```

### Retrieve information about the web server

The following example command retrieves information about the web server running on an E2800 controller.

```
curl -k -u <rw:password> -H "Accept: application/json" "https://<array IP address>/devmgr/utlils/about"
```

### Restart the web server on an E2800 controller

The following example command restarts the web server on an E2800 controller.

```
curl -X POST -k -u <rw:password> "https://<array IP address>/devmgr/v2/restart"
```

### Retrieve the available AutoSupport log file list

The following example command gets the list of available AutoSupport log files from the E2800 controller.

```
curl -k -u <rw:password> -H "Accept: application/json" "https://<array IP address>/devmgr/v2/device-asup/logs"
```

### Retrieve a specific AutoSupport log file downloaded to the local host

The following example command downloads an AutoSupport log file from the E2800 controller to the local host.

```
curl -k -u <rw:password> -O "https://<array IP address>/devmgr/v2/device-asup/logs/ASUPMessages.0"
```

### Override to allow HTTP traffic to the E2800 storage array

The following example command overrides the HTTPS default to allow HTTP traffic.

**Note:** Run this command on both controllers.

**Note:** If the value in the payload is set to true all traffic will be redirected to HTTPS. If the value is false, HTTP traffic is allowed to the web server.

**Note:** Restart the web server after running this command.

**Note:** This command may return an error with an error code of 204. This is a no content reply which can be ignored.

```
curl -k -u <rw:password> -H "Accept: application/json" -H "Content-Type: application/json" -d @- "https://<array IP address>/devmgr/v2/override" <<EOF
{
  "key": "redirect.http",
  "value": "true"
}
EOF
```

### Retrieve the syslog configuration

The following example command gets the syslog configuration from the E2800 storage array.

```
curl -k -u <rw:password> -H "Accept: application/json" "https://<array IP address>/devmgr/v2/storage-systems/1/device-alerts/alert-syslog"
```

### Post a new syslog configuration

The following example command posts a new syslog configuration to the E2800 storage array in JSON format.

**Note:** To indicate to curl to read the format from stdin write@- .

```
curl -k -u <rw:password> -H "Accept: application/json" -H "Content-Type: application/json" -d @- "https://<array IP address>/devmgr/v2/storage-systems/1/device-alerts/alert-syslog" <<EOF
{
  "syslogReceivers":
  [
    {
```

```

        "serverName": "testserver.yourdomain.com",
        "portNumber": 514
    },
    {
        "defaultFacility": 3,
        "defaultTag": "StorageArray"
    }
EOF

```

### Test the syslog configuration

The following example command sends a test alert to the configured syslog server addresses.

```

curl -X POST -k -u <rw:password> -H "Accept: application/json" -H
"Content-Type: application/json" "https://<array IP address>/devmgr/v2/
storage-systems/1/device-alerts/alert-syslog-test"

```

### Get a list of failure log entries with details turned on

The following example command gets a list of failure log entries from the E2800 storage array with all available details.

```

curl -k -u <rw:password> -H "Accept: application/json" "https://<array
IP address>/devmgr/v2/storage-systems/1/failures?details=true"

```

## Command reference

---

**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](#) on page 31 lists the general rules for entering the names of storage array entities, such as volumes or drives, with the script commands.
- [Formatting CLI commands](#) on page 32 lists the general formatting rules that apply to the CLI command wrapper.
- [Formatting rules for script commands](#) on page 32 lists the general formatting rules that apply to the script command syntax.
- [Firmware compatibility levels](#) on page 38 explains how to interpret the firmware level information.
- [Commands listed by function](#) on page 45 lists the script commands organized into groups related to the physical features, the logical features, and the operational features of the storage array.

**Important: 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.

## Commands listed by function

### Asynchronous mirroring commands

[Activate asynchronous mirroring](#) on page 72

[Add volume to asynchronous mirror group](#) on page 80

[Cancel asynchronous mirror group role reversal](#) on page 88

[Check asynchronous mirror group consistency](#) on page 91

[Check storage array connectivity](#) on page 94

[Clear asynchronous mirroring fault](#) on page 96

[Create asynchronous mirror group](#) on page 104

[Deactivate asynchronous mirroring](#) on page 176

[Delete asynchronous mirror group](#) on page 177

[Establish asynchronous mirrored pair](#) on page 216

[Remove incomplete asynchronous mirrored pair from asynchronous mirror group](#) on page 236

*Remove volume from asynchronous mirror group* on page 239  
*Reset asynchronous mirror group statistics* on page 246  
*Reset iSCSI IP address* on page 250  
*Resume asynchronous mirror group* on page 255  
*Save asynchronous mirror group statistics* on page 265  
*Set asynchronous mirror group* on page 294  
*Show asynchronous mirror groups* on page 382  
*Show asynchronous mirror group synchronization progress* on page 383  
*Start asynchronous mirroring synchronization* on page 440  
*Suspend asynchronous mirror group* on page 481  
*Test asynchronous mirror group connectivity* on page 484

## Auto-load balancing commands

*Save auto-load balancing statistics* on page 270  
*Set storage array to enable or disable auto-load balancing* on page 349

## AutoSupport bundle collection commands

*Capture or view an AutoSupport log* on page 89  
*Display AutoSupport bundle collection settings* on page 198  
*Display AutoSupport message collection schedule* on page 200  
*Disable AutoSupport at the EMW level SMcli version* on page 197  
*Enable or disable AutoSupport at the EMW management domain level (SMcli only)* on page 210  
*Enable or disable the AutoSupport OnDemand feature at the EMW management domain level (SMcli only)* on page 211  
*Enable or disable the AutoSupport OnDemand Remote Diagnostics feature at the EMW management domain level (SMcli only)* on page 213  
*Reset AutoSupport message collection schedule* on page 247  
*Save storage array support data* on page 283  
*Specify the AutoSupport delivery method* on page 438  
*Test the AutoSupport configuration* on page 485

## Consistency group commands

*Add member to consistency group* on page 78  
*Create consistency group* on page 107  
*Create consistency group snapshot image* on page 115  
*Create consistency group snapshot volume* on page 116  
*Create consistency group snapshot volume mapping* on page 120  
*Delete consistency group* on page 178  
*Delete consistency group snapshot image* on page 179

*Delete consistency group snapshot volume* on page 180  
*Remove member volume from consistency group* on page 237  
*Resume consistency group snapshot volume* on page 256  
*Set consistency group attributes* on page 296  
*Set consistency group snapshot volume* on page 298  
*Show consistency group* on page 384  
*Show consistency group snapshot image* on page 386  
*Start consistency group snapshot rollback* on page 443  
*Stop consistency group snapshot rollback* on page 466  
*Stop consistency group snapshot volume* on page 468  
*Stop pending snapshot images on consistency group* on page 471

## Controller commands

*Clear drive channel statistics* on page 98  
*Enable controller data transfer* on page 208  
*Reset controller* on page 248  
*Save controller NVSRAM* on page 267  
*Save drive channel fault isolation diagnostic status* on page 267  
*Set controller* on page 299  
*Set controller DNS settings* on page 304  
*Set controller NTP settings* on page 305  
*Set controller service action allowed indicator* on page 306  
*Set drive channel status* on page 313  
*Set host channel* on page 321  
*Show controller* on page 387  
*Show controller NVSRAM* on page 390  
*Show drive channel statistics* on page 395  
*Start controller trace* on page 444  
*Start drive channel fault isolation diagnostics* on page 447  
*Start drive channel locate* on page 449  
*Start iSCSI DHCP refresh* on page 453  
*Stop drive channel fault isolation diagnostics* on page 469  
*Stop drive channel locate* on page 470

## Controller health image commands

*Save storage array controller health image* on page 272  
*Show storage array controller health image* on page 419  
*Start storage array controller health image* on page 460

## Database commands

- [Clear storage array configuration](#) on page 99
- [Clear storage array recovery mode](#) on page 101
- [Save storage array DBM database](#) on page 273
- [Save storage array DBM validator information file](#) on page 274
- [Show storage array DBM database](#) on page 420
- [Start storage array configuration database diagnostic](#) on page 459
- [Stop storage array configuration database diagnostic](#) on page 478

## Disk pool commands

- [Create disk pool](#) on page 121
- [Delete disk pool](#) on page 181
- [Reduce disk pool capacity](#) on page 233
- [Set disk pool](#) on page 307
- [Set disk pool \(modify disk pool\)](#) on page 310
- [Show disk pool](#) on page 392
- [Start disk pool locate](#) on page 446
- [Stop disk pool locate](#) on page 469

## Drive commands

- [Download drive firmware](#) on page 202
- [Replace drive](#) on page 244
- [Revive drive](#) on page 262
- [Save drive channel fault isolation diagnostic status](#) on page 267
- [Save drive log](#) on page 268
- [Set drive hot spare](#) on page 314
- [Set drive service action allowed indicator](#) on page 314
- [Set drive state](#) on page 315
- [Set FIPS drive security identifier](#) on page 318
- [Set foreign drive to native](#) on page 319
- [Show drive](#) on page 393
- [Show drive download progress](#) on page 396
- [Show drive performance statistics](#) on page 396
- [Show replaceable drives](#) on page 399
- [Start drive channel fault isolation diagnostics](#) on page 447
- [Start drive initialize](#) on page 449
- [Start drive locate](#) on page 450
- [Start drive reconstruction](#) on page 451



[Start FDE secure drive erase](#) on page 455

[Stop drive channel fault isolation diagnostics](#) on page 469

[Stop drive locate](#) on page 470

## Event alert commands

[Set event alert filtering](#) on page 316

[Show blocked events](#) on page 384

[Test alerts](#) on page 484

## Feature management commands

[Enable storage array feature](#) on page 214

[Disable storage array feature](#) on page 197

[Set storage array to enable or disable Automatic Load Balancing](#) on page 349

[Show storage array](#) on page 410

## Host topology commands

[Create host](#) on page 127

[Create host group](#) on page 128

[Create host port](#) on page 129

[Delete host](#) on page 182

[Delete host group](#) on page 182

[Delete host port](#) on page 183

[Set host](#) on page 320

[Set host channel](#) on page 321

[Set host group](#) on page 322

[Set host port label](#) on page 323

[Show host ports](#) on page 398

[Show storage array host topology](#) on page 421

## Input output controller (IOC) core dump commands

[Save input output controller \(IOC\) dump](#) on page 269

[Start input output controller \(IOC\) dump](#) on page 452

## Internet storage name service (iSNS) commands

[Set storage array](#) on page 344

[Set storage array iSNS registration](#) on page 351

[Set storage array iSNS server IPv6 address](#) on page 354

[Set storage array iSNS server listening port](#) on page 354

[Set storage array iSNS server refresh](#) on page 355

[Start storage array iSNS server refresh](#) on page 461

## ISCSI commands

- [Activate iSCSI initiator](#) on page 72
- [Create iSCSI initiator](#) on page 131
- [Delete iSCSI initiator](#) on page 184
- [Diagnose controller iSCSI host cable](#) on page 193
- [Reset storage array iSCSI baseline](#) on page 252
- [Save storage array iSCSI statistics](#) on page 279
- [Set iSCSI initiator](#) on page 324
- [Set iSCSI target properties](#) on page 326
- [Show current iSCSI sessions](#) on page 390
- [Show storage array negotiation defaults](#) on page 423
- [Show storage array unconfigured iSCSI initiators](#) on page 424
- [Start iSCSI DHCP refresh](#) on page 453
- [Stop storage array iSCSI session](#) on page 479

## ISCSI extension for RDMA (iSER) commands

- [Reset storage array iSER baseline](#) on page 253
- [Save storage array iSER statistics](#) on page 280
- [Set iSER target](#) on page 327

## Repository volume analysis commands

- [Check repository consistency](#) on page 92

## Session command

- [Set session](#) on page 330

## Snapshot group commands

- [Create snapshot group](#) on page 146
- [Delete snapshot group](#) on page 185
- [Revive snapshot group](#) on page 263
- [Set snapshot group attributes](#) on page 331
- [Set snapshot group media scan](#) on page 333
- [Set snapshot group repository volume capacity](#) on page 334
- [Set snapshot group schedule](#) on page 336
- [Set SSD cache for a volume](#) on page 342
- [Show snapshot group](#) on page 399
- [Stop snapshot group pending snapshot images](#) on page 472

## Snapshot image commands

- [Create snapshot image](#) on page 155

[Delete snapshot image](#) on page 186

[Resume snapshot image rollback](#) on page 258

[Show snapshot image](#) on page 401

[Start snapshot image rollback](#) on page 456

[Stop snapshot image rollback](#) on page 473

## Snapshot volume commands

[Create read-only snapshot volume](#) on page 145

[Create snapshot volume](#) on page 156

[Delete snapshot volume](#) on page 187

[Rename snapshot volume](#) on page 242

[Resume snapshot volume](#) on page 259

[Revive snapshot volume](#) on page 263

[Set read-only snapshot volume to read/write volume](#) on page 328

[Set snapshot volume media scan](#) on page 339

[Set snapshot volume repository volume capacity](#) on page 340

[Set SSD cache for a volume](#) on page 342

[Show snapshot volumes](#) on page 403

[Stop snapshot volume](#) on page 474

## SSD cache commands

[Add drives to SSD cache](#) on page 77

[Change SSD cache application type](#) on page 90

[Create SSD cache](#) on page 159

[Delete SSD cache](#) on page 188

[Remove drives from SSD cache](#) on page 235

[Rename SSD cache](#) on page 242

[Resume SSD cache](#) on page 260

[Set SSD cache for a volume](#) on page 342

[Show SSD cache](#) on page 406

[Start SSD cache locate](#) on page 457

[Start SSD cache performance modeling](#) on page 458

[Stop SSD cache locate](#) on page 475

[Stop SSD cache performance modeling](#) on page 476

[Suspend SSD cache](#) on page 482

## Storage array commands

[Activate storage array firmware](#) on page 73

[Autoconfigure storage array](#) on page 82

*Autoconfigure storage array hot spares* on page 88

*Clear storage array configuration* on page 99

*Clear storage array event log* on page 101

*Clear storage array firmware pending area* on page 101

*Create storage array security key* on page 161

*Disable storage array feature* on page 197

*Download storage array drive firmware* on page 204

*Download storage array firmware/NVSRAM* on page 205

*Download storage array NVSRAM* on page 207

*Enable storage array feature* on page 214

*Export storage array security key* on page 218

*Import storage array security key* on page 218

*Load storage array DBM database* on page 223"

*Reset storage array diagnostic data* on page 251

*Reset storage array Infiniband statistics baseline* on page 252

*Reset storage array iSCSI baseline* on page 252

*Reset storage array RLS baseline* on page 254

*Reset storage array SAS PHY baseline* on page 254

*Reset storage array SOC baseline* on page 254

*Reset storage array volume distribution* on page 255

*Save storage array configuration* on page 270

*Save storage array DBM database* on page 273

*Save storage array DBM validator information file* on page 274

*Save storage array diagnostic data* on page 275

*Save storage array events* on page 276

*Save storage array firmware inventory* on page 278

*Save storage array InfiniBand statistics* on page 278

*Save storage array iSCSI statistics* on page 279

*Save storage array performance statistics* on page 281

*Save storage array RLS counts* on page 281

*Save storage array SAS PHY counts* on page 282

*Save storage array SOC counts* on page 282

*Save storage array state capture* on page 283

*Save storage array support data* on page 283

*Set storage array* on page 344

*Set storage array ICMP response* on page 350

*Set storage array iSNS server IPv4 address* on page 353

*Set storage array iSNS server IPv6 address* on page 354

*Set storage array iSNS server listening port* on page 354

*Set storage array iSNS server refresh* on page 355

*Set storage array controller battery learn cycle* on page 356

*Set storage array redundancy mode* on page 357

*Set storage array security key* on page 358

*Set storage array time* on page 358

*Set storage array tray positions* on page 359

*Show storage array* on page 410

*Show storage array auto configuration* on page 414

*Show storage array host topology* on page 421

*Show storage array LUN mappings* on page 422

*Show storage array negotiation defaults* on page 423

*Show storage array power information* on page 424

*Show storage array unreadable sectors* on page 425

*Show storage array unconfigured iSCSI initiators* on page 424

*Start FDE secure drive erase* on page 455

*Start storage array configuration database diagnostic* on page 459

*Start storage array iSNS server refresh* on page 461

*Start storage array locate* on page 462

*Stop storage array configuration database diagnostic* on page 478

*Stop storage array drive firmware download* on page 478

*Stop storage array iSCSI session* on page 479

*Stop storage array locate* on page 479

*Validate storage array security key* on page 490

## **Synchronous Mirroring commands**

*Activate Synchronous Mirroring* on page 74

*Check synchronous mirroring status* on page 95

*Create synchronous mirroring* on page 162

*Deactivate synchronous mirroring* on page 176

*Diagnose synchronous mirroring* on page 196

*Re-create synchronous mirroring repository volume* on page 230

*Remove synchronous mirroring* on page 238

*Resume synchronous mirroring* on page 261

*Set synchronous mirroring* on page 360

*Show synchronous mirroring volume candidates* on page 426

*Show synchronous mirroring volume synchronization progress* on page 427

*Start synchronous mirroring synchronization* on page 462

*Suspend synchronous mirroring* on page 482

## Tray commands

*Download tray configuration settings* on page 207

*Save tray log* on page 289

*Set Drawer Service Action Allowed Indicator* on page 312

*Set tray identification* on page 364

*Set tray service action allowed indicator* on page 365

*Start tray locate* on page 463

*Stop tray locate* on page 480

## Uncategorized commands

*Recover SAS port mis-wire* on page 229

*Show string* on page 426

## Volume commands

*Check volume parity* on page 95

*Clear volume reservations* on page 102

*Clear volume unreadable sectors* on page 103

*Create RAID volume (automatic drive select)* on page 132

*Create RAID volume (free extent based select)* on page 137

*Create RAID volume (manual drive select)* on page 140

*Delete volume* on page 189

*Recover RAID volume* on page 225

*Remove volume LUN mapping* on page 240

*Repair volume parity* on page 243

*Set SSD cache for a volume* on page 342

*Set volume attributes for a volume in a volume group* on page 371

*Show volume* on page 429

*Show volume action progress* on page 431

*Show volume performance statistics* on page 435

*Show volume reservations* on page 437

*Start volume initialization* on page 466

## Volume commands for disk pools

*Create volume in disk pool* on page 172

*Delete volume from disk pool* on page 190

*Increase capacity of volume in disk pool or volume group* on page 219

*Initialize thin volume* on page 221

*Set volume attributes for a volume in a disk pool* on page 366

[Set thin volume attributes](#) on page 362

[Set volume mapping](#) on page 380

[Show thin volume](#) on page 428

## Volume copy commands

[Create volume copy](#) on page 165

[Recopy volume copy](#) on page 224

[Remove volume copy](#) on page 239

[Set volume copy](#) on page 377

[Show volume copy](#) on page 431

[Show volume copy source candidates](#) on page 432

[Show volume copy target candidates](#) on page 432

[Stop volume copy](#) on page 480

## Volume group commands

[Create volume group](#) on page 167

[Delete volume group](#) on page 191

[Enable volume group security](#) on page 216

[Revive volume group](#) on page 264

[Set volume group](#) on page 378

[Set volume group forced state](#) on page 380

[Show volume group](#) on page 433

[Show volume group export dependencies](#) on page 434

[Show volume group import dependencies](#) on page 435

[Start volume group defragment](#) on page 463

[Start volume group export](#) on page 464

[Start volume group import](#) on page 464

[Start volume group locate](#) on page 465

[Stop volume group locate](#) on page 481

## Storage management

The SANtricity 11.30 release introduces a new kind of controller, the E2800. This section provides information on storage management for three different environments: an environment containing a single E2800, an environment containing multiple E2800s, and an environment containing one or more E2800s, in addition to one or more E2700s, E5600s, and/or EF560s.

## Managing an E2800 storage array

### What's new in an E2800 storage array

The E2800 is a new entry-level 12 GB SAS system with SANtricity 11.30 software. This system features the following new capabilities:

- Embedded web services
- Embedded SANtricity System Manager, with an easy to use graphical user interface
- The ability to view SSL information, get a Certificate Signing Request (CSR), and import a new certificate
- The ability to store and present up to 30 days of performance data, including I/O latency, IOPS, and throughput
- The ability to do application/workload tagging
- Easier alert management, including an embedded SNMP agent and MIB
- The ability to view, configure and test syslog receivers
- Embedded AutoSupport functionality

The new E2800 storage array can be managed using SANtricity 11.30, embedded directly on the controller, along with embedded web services. If you have a single E2800 and have no need to do mirroring or use a command line interface, a web browser is all you need to manage your system.



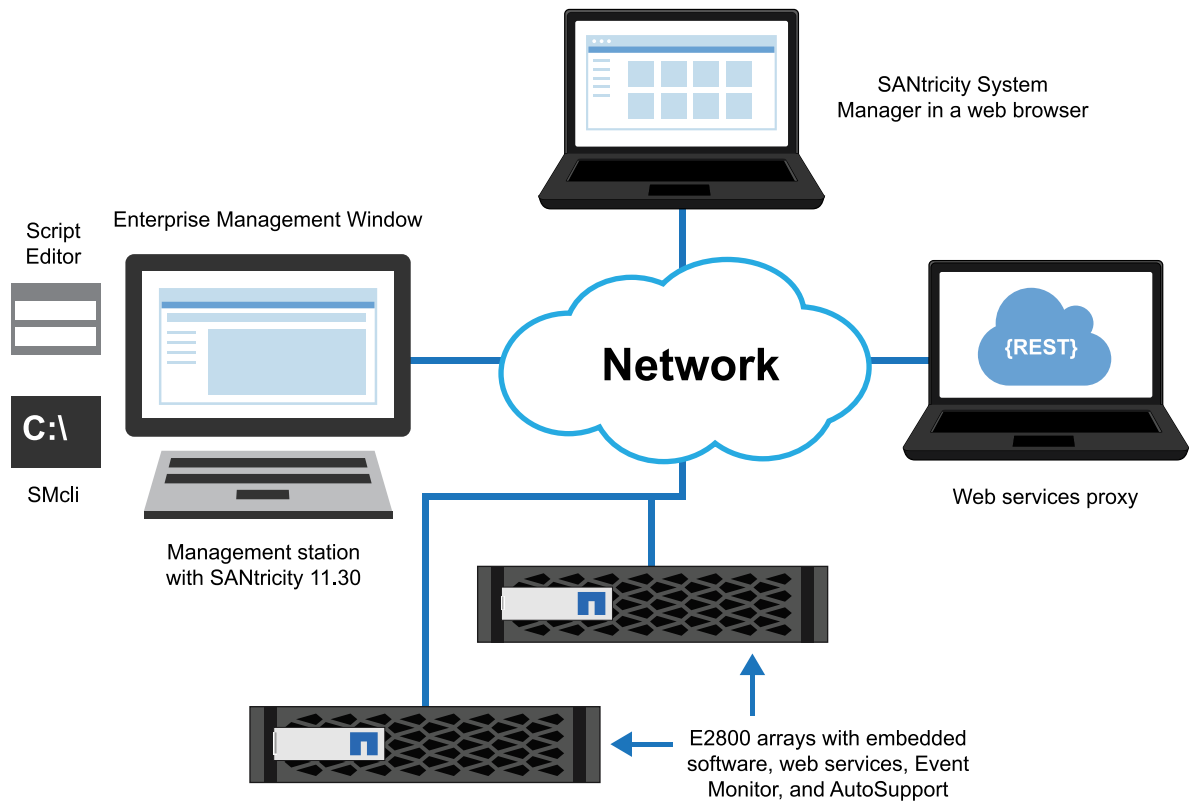
## Managing multiple E2800 storage arrays

Install SANtricity Storage Manager Enterprise Management Window for domain management, for mirroring operations between two storage arrays, and for using the SMcli or the Script Editor.

You also can use the RESTful API (Java or Python client libraries in the SDK or curl commands) to perform management for the E2800 storage array.

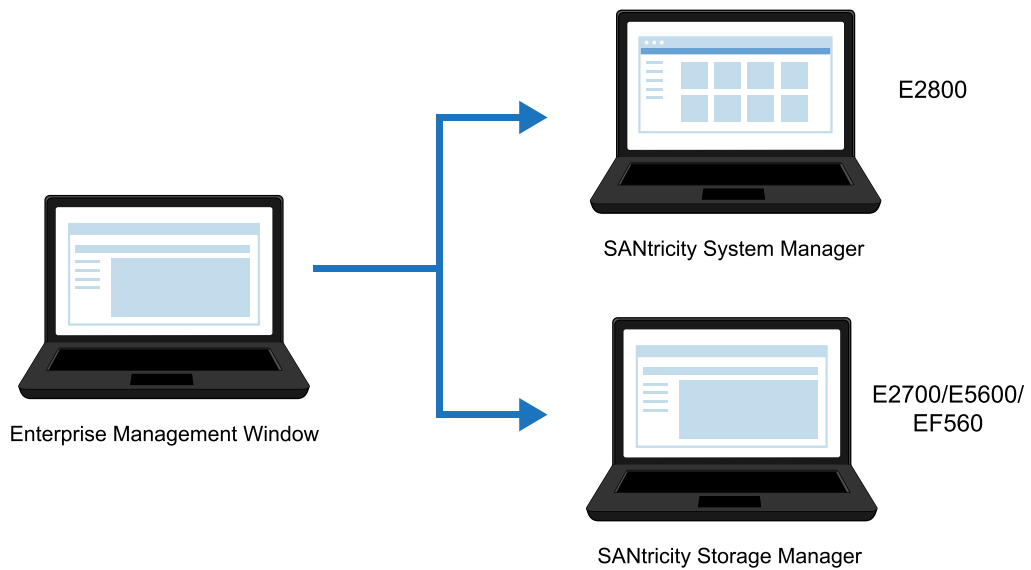
The following graphic shows the management components you can use to manage one or more E2800 storage arrays:





### Management of mixed storage arrays

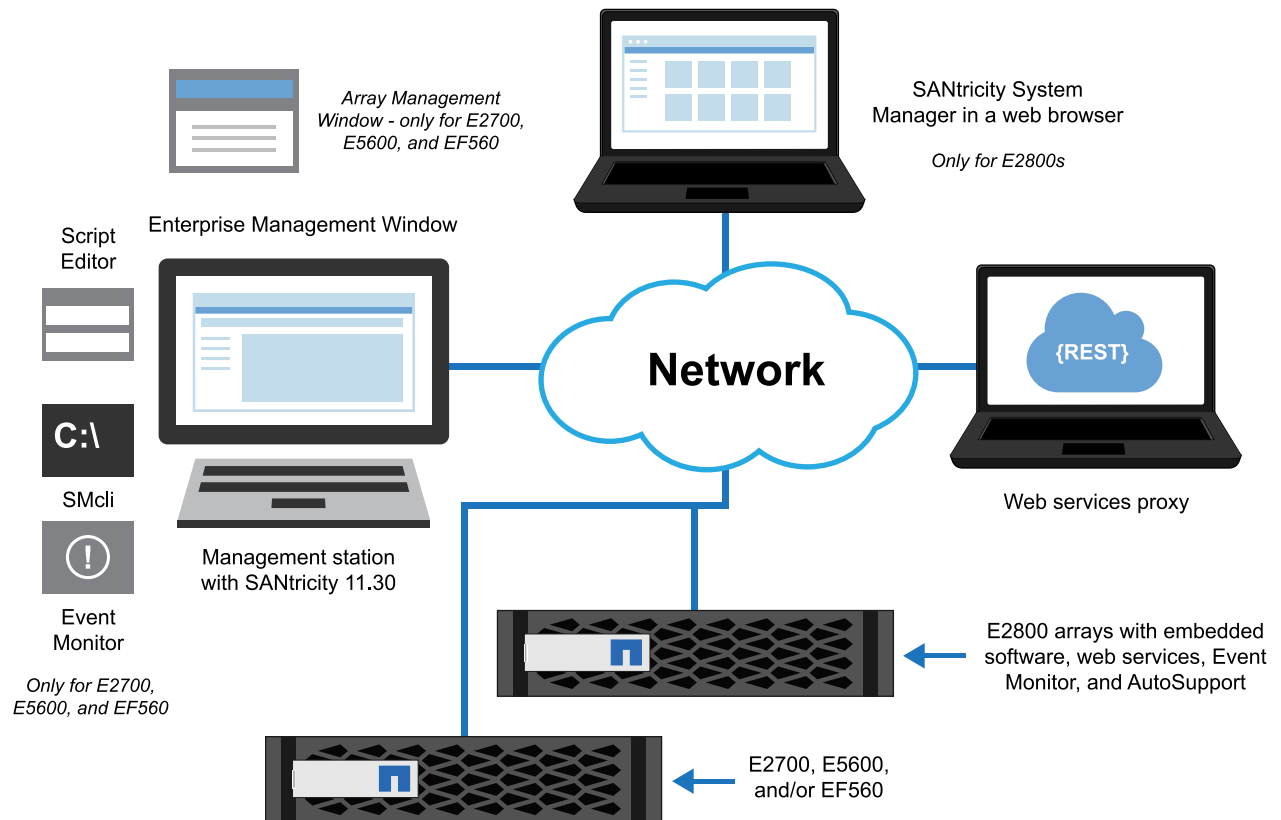
The Enterprise Management Window (EMW) on the management station discovers E2800, E2700, E5600, and EF560 storage arrays in your domain. A mixed environment is easy to manage, because the EMW launches the appropriate software for managing each individual array. For the E2800 storage array, the EMW launches SANtricity System Manager in a web browser. For the E2700, E5600, or EF560 storage arrays, the EMW launches an Array Management Window.



In addition to using the Enterprise Management Window for domain management, you also use it for mirroring operations between two storage arrays, and for using the SMcli or the Script Editor.

For managing individual arrays, you can use the graphical user interfaces (GUID) described earlier, the SMcli or Script Editor (with some exceptions noted below), or the RESTful API.

The following graphic shows the management components in a mixed environment:



## Event Management

Critical events indicate a problem with the storage array. If you resolve the critical event immediately, you might prevent loss of data access.

When a critical event occurs, it is logged in the Event Log. All critical events are sent to the SNMP management console or to the email recipient that you have configured to receive alerts.

When you receive a critical event, refer to the Recovery Guru procedure for a detailed description of the critical event. Complete the Recovery Guru procedure to correct the critical event. To correct certain critical events, you might need to contact technical support.

There are differences in the way events and SNMP trap messages are managed between the new E2800 and the E2700, E5600 and EF560 storage arrays. These differences are described below. Because of the differences, there are different management methods available.

- For the E2800 storage array, you can configure event and SNMP capabilities and settings with the SANtricity System Manager, or you can use the RESTful API. You also can write scripts directly accessing the embedded web services on the controller, using the RESTful API.
- For the E2700, E5600 and EF560 storage arrays, you can configure event and SNMP capabilities via the SANtricity Storage Manager, or you can use the SMcli or the script editor in the

Enterprise Management Window. You must run the Event Monitor on at least one management station on the network.

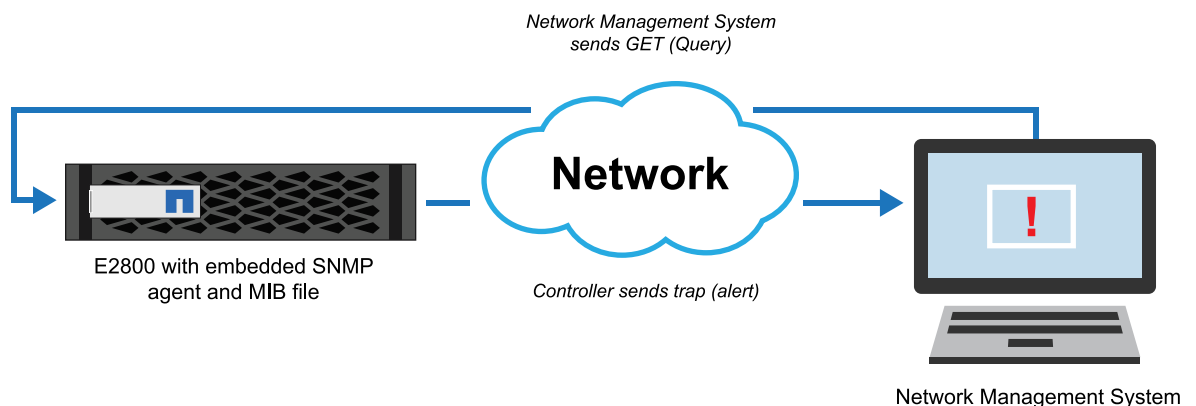
### About SNMP monitoring of the storage array

The following terminology definitions are useful in understanding the SNMP capabilities of the E-Series storage arrays:

Term	Description
Storage array	<p>This is a managed device that can be managed and/or monitored through SNMP.</p> <ul style="list-style-type: none"> <li>• With the E2700, E5600, and EF560 storage arrays, the Event Monitor enables an SNMP management application to provide monitoring capabilities for the array.</li> <li>• With the E2800 storage array, an SNMP agent is embedded on the controller, allowing an SNMP management application to provide simple monitoring capabilities for the array. The embedded SNMP agent is compliant with the SNMP V2C standard and RFC 1213 (MIB-II).</li> </ul>
SNMP agent	<p>Software that generates outgoing trap messages whenever critical MEL events are logged on the array.</p> <ul style="list-style-type: none"> <li>• With the E2700, E5600, and EF560 storage arrays, the Event Monitor contains the SNMP agent, and only generates outgoing trap messages.</li> <li>• With the E2800 storage array, an SNMP agent is embedded on the controller, allowing the array to provide basic query support and also handle GET requests from an SNMP management application.</li> </ul>
Management Information Base (MIB)	<p>A Management Information Base describes the structure of the data for a specific managed device.</p> <ul style="list-style-type: none"> <li>• With the E2700, E5600, and EF560 storage array, the Event Monitor on the storage management station contains the MIB file.</li> <li>• With the E2800 storage array, the MIB file is embedded on the controller.</li> </ul>
Trap message	<p>This is an asynchronous notification from the SNMP agent to the SNMP manager that a specific event has occurred on the managed device.</p>

Term	Description
Network Management System	This is an SNMP management application acting as an administration and monitoring console for SNMP managed devices, including E-Series storage arrays.
SNMP community	An SNMP community comprises both managed devices and SNMP management applications. The SNMP community string is included in the data packets that are passed over the network. <ul style="list-style-type: none"> <li>The community string is included in all trap notifications that are passed from an SNMP agent to an SNMP management application.</li> <li>The SNMP agent on the E2800 storage array rejects queries from an SNMP management system that do not contain the appropriate community string.</li> </ul>
SNMP trap destination	SNMP TRAP messages will only be sent to destinations you specify. Each trap destination must have a unique IP address. The embedded SNMP Agent supports configuration of both IPV4 and IPV6 destination addresses. A previously registered community must also be specified by the user when configuring a trap destination so that the embedded SNMP agent can place the correct community name in each outgoing trap message.

## Event management for E2800 storage arrays



## RESTful API sample code in Python and Java client libraries for new E2800 capabilities

Sample programs that perform the following tasks are available in the Python and Java client libraries, part of the REST SDK:

- View SSL configuration, get a Certificate Signing Request (CSR), and import a new certificate

- View, configure and test syslog receivers
- View historical data and rolled up statistics for I/O latency, IOPS, and throughput.
- Create volume sets based on application types (application tagging)
- Collect and view support bundle artifacts specific to the E2800 storage array

#### Related references

[Using curl commands with the E2800 storage array](#) on page 39

[Set event alert filtering](#) on page 316

[Clear storage array event log](#) on page 101

[Show blocked events](#) on page 384

[Save storage array events](#) on page 276

[Set event alert filtering](#) on page 316

[Test alerts](#) on page 484

[Register SNMP community](#) on page 234

[Register SNMP trap destination](#) on page 234

[Show SNMP communities](#) on page 404

[Show SNMP MIB II system group variables](#) on page 405

[Test SNMP trap destination](#) on page 487

[Unregister SNMP community](#) on page 487

[Unregister SNMP trap destination](#) on page 488

[Update SNMP community](#) on page 488

[Update SNMP MIB II system group variables](#) on page 488

[Update SNMP trap destination](#) on page 489

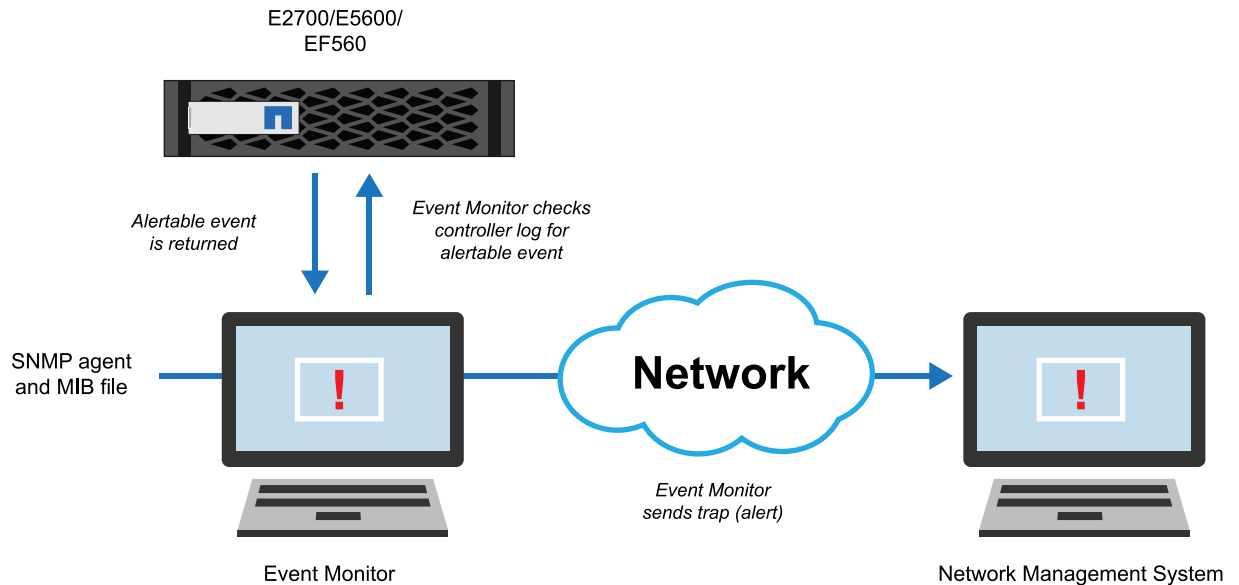
#### Related information

[curl: The curl project](#)

## Event management for the E2700, E5600, and EF560

### Event management for the E2700, E5600, and EF560 storage arrays

With the E2700, E5600, and EF560 storage arrays, the Event Monitor on the storage management station periodically queries the storage array for alertable events. The SNMP agent running in the Event Monitor then generates a trap message for each event and forwards to the SNMP management application.



### Related references

- [Set event alert filtering](#) on page 316
- [Show blocked events](#) on page 384
- [Set event alert filtering](#) on page 316
- [Clear storage array event log](#) on page 101
- [Save storage array events](#) on page 276
- [Test alerts](#) on page 484

## AutoSupport Overview

AutoSupport (ASUP) is a feature that enables storage arrays to automatically collect support data into a customer support bundle and send the data to NetApp technical support. Technical support can then perform remote troubleshooting and problem analysis with the storage management software.

### About AutoSupport and AutoSupport OnDemand

ASUP collects support data to report configuration, subsystem status, and exceptions in near-real time. ASUP OnDemand (AOD) is an optional extension to ASUP that introduces two-way communication between the E-Series array and the ASUP backend. ASUP RemoteDiagnostics provides the ability for NetApp technical support to remotely ask the given E-Series array to create a new ASUP message, or request retransmission of a particular ASUP message.

**Important:** Under no circumstances is business data from the storage array ever sent to NetApp support.

ASUP messages typically include a collection of system log 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).

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

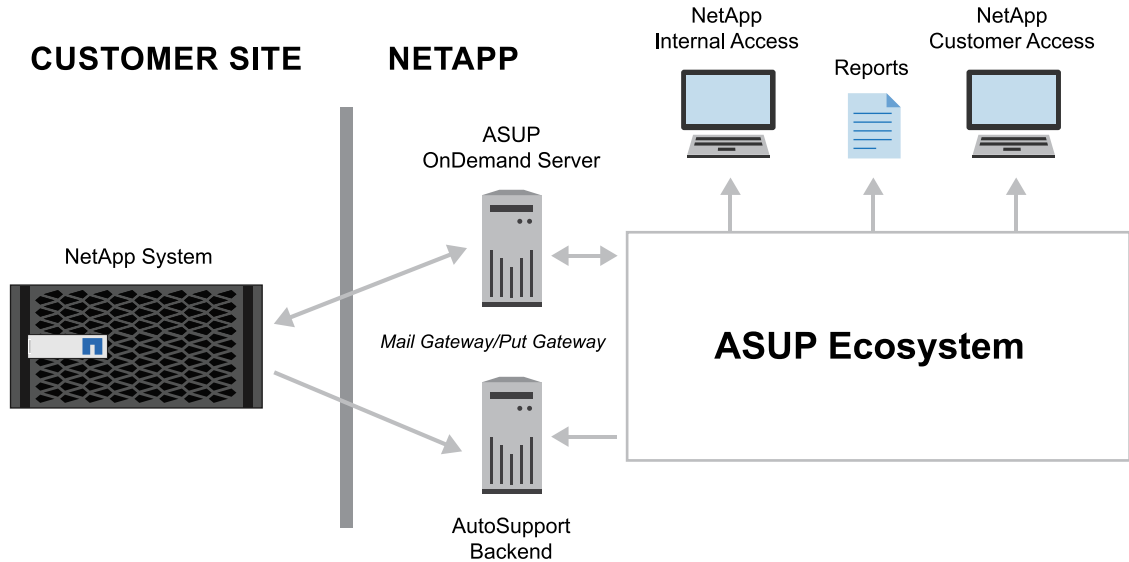
**Methods for collecting support data**

With the implementation of ASUP and AOD, you have three possible methods for collecting support data in a storage array:

- Standard ASUP collection - Data is automatically collected and sent to technical support, triggered by a daily or weekly schedule, or by an event occurrence that is specified to automatically send a support bundle.
- ASUP OnDemand - With AOD enabled, before an E-Series system can send any standard ASUP message to the ASUP backend, it must first obtain permission from the ASUP OnDemand Server (AODS). ASUP OnDemand also introduces the ability for a given E-Series array to check in periodically with the AODS to see if there are any pending actions for the particular E-Series array.
- Legacy support bundle collection (E2700, E5600, or EF560 storage arrays only) - Collection of legacy support bundle data is user-configurable at intervals you schedule. You can then manually send the support bundles to technical support.

For the E2700, E5600, or EF560 storage arrays only, ASUP operations and legacy support bundle operations are mutually exclusive on a given storage array. When you turn on ASUP, with or without AOD, you automatically disable legacy support bundle collection. If you want to run legacy support bundle collection, you must turn off ASUP. Because ASUP speeds up troubleshooting and problem analysis, ASUP is the preferred data collection method to use if available on the storage array.

**AutoSupport components**



AutoSupport component	Description
NetApp System	<p>The NetApp System is an array on the customer site. The system is managed by the SANtricity software on a management station or host, and the Persistent Monitor runs on the management station or host.</p> <p><b>Note:</b> The NetApp System stores an AutoSupport log, located in the ASUPLog directory on the array. The log provides information about status, history of transmission, and any errors encountered during the collection and delivery of the AutoSupport messages.</p>

<b>AutoSupport component</b>	<b>Description</b>
ASUP OnDemand Server (AODS)	<p>With AOD enabled, before an E-Series system can send any standard ASUP message to the ASUP backend, it must first obtain permission from the AODS.</p> <p><b>Important:</b> Under no circumstances does the AODS initiate communication with an E-Series array. The flow of communication always begins with the E-Series array, and a reply is sent back from the AODS. The reply might contain actions requested by technical support, such as to retransmit a missing ASUP message.</p>
Mail Gateway/Put Gateway	The HTTP(s) Put/Post Gateway moves the messages into the ASUP Ecosystem.
ASUP Ecosystem	The ASUP Ecosystem provides near real-time access to messages by NetApp technical support. Additionally, there is a portal within the ecosystem that NetApp customers can use to access their ASUP data. Finally, the ASUP Ecosystem provides the ability to generate reports used by various groups within NetApp.

### AutoSupport data processing

The backend processing of AutoSupport data is used in several ways:

- Routing messages based on message content
- Support case generation
- Proactive health-check and system level risk detection
- Service entitlement
- Subscription-based service fulfillment
- Persisting the support information associated with the ASUP message
- Data analysis to improve the E-Series products

### AutoSupport configuration in E-Series storage arrays

The E2800 storage array includes embedded ASUP functionality, rather than relying on an externally installed Event Monitor to configure and monitor AutoSupport (as does the E2700, E5600, and EF560 storage arrays). Because of the differences in the E2800 storage array, there are different management methods for configuring AutoSupport across the two kinds of arrays:

- For the E2800 storage array, you can use SANtricity System Manager or the RESTful API to configure AutoSupport. There are sample programs in the Java and Python client libraries which are part of the REST API SDK.
- For the E2700, E5600, and EF560 storage arrays, you can use the SANtricity Storage Manager application or SMcli.


### AutoSupport for E2800 storage arrays

With the E2800 storage array, the controller contains embedded AutoSupport functionality.



**CUSTOMER SITE**

**NETAPP**



**cURL statements for AutoSupport configuration**

This document provides cURL statements as an alternative to SMcli for the E2800. These statements connect and communicate with the embedded web services on the E2800 to accomplish these tasks, using the RESTful API.

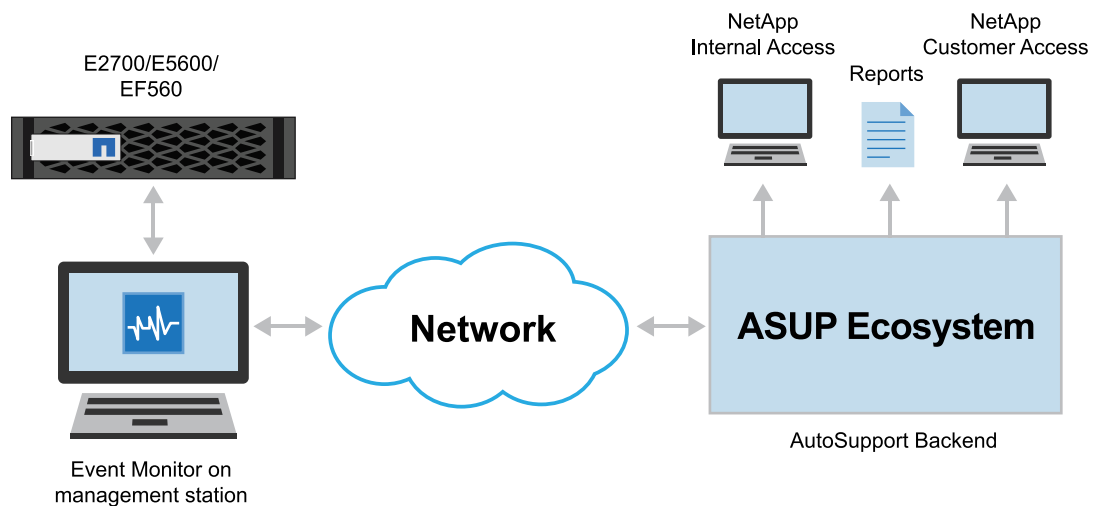
- View and update AutoSupport configuration
- View and update collection schedules
- View AutoSupport log files
- Validate AutoSupport configuration

**AutoSupport for E2700, E5600, and EF560 storage arrays**

With the E2700, E5600, or EF560 storage arrays, the Event Monitor on a storage management station on your network must be running to use AutoSupport.

**CUSTOMER SITE**

**NETAPP**



**Related references**

- [Using curl commands with the E2800 storage array](#) on page 39
- [Capture or view an AutoSupport log](#) on page 89
- [Display AutoSupport bundle collection settings](#) on page 198
- [Display AutoSupport message collection schedule](#) on page 200
- [Disable AutoSupport at the EMW level SMcli version](#) on page 197
- [Enable or disable AutoSupport at the EMW management domain level \(SMcli only\)](#) on page 210
- [Enable or disable the AutoSupport OnDemand feature at the EMW management domain level \(SMcli only\)](#) on page 211
- [Enable or disable the AutoSupport OnDemand Remote Diagnostics feature at the EMW management domain level \(SMcli only\)](#) on page 213
- [Reset AutoSupport message collection schedule](#) on page 247
- [Save storage array support data](#) on page 283
- [Specify the AutoSupport delivery method](#) on page 438
- [Test the AutoSupport configuration](#) on page 485

## Management Methods in SANtricity 11.30

SANtricity 11.30 introduces a new controller with embedded management software, web services, and Event Monitoring and AutoSupport. Previous controllers such as the E2700, E5600, and EF560 do not have this embedded functionality. Because you may have a mixed environment, with both the new E2800 storage array and older storage arrays, there are a variety of options to manage this mixed environment.

Use the table below to determine, by task, which management component provides the functionality you need.

**Management Methods**

The following high level summary describes management methods for different functionality:

<b>Task</b>	<b>E2700/E5600/EF560, all other E-Series arrays except the E2800</b>	<b>E2800</b>
Array Management GUI	SANtricity AMW - install on management station	SANtricity System Manager - On-Box
Multi-system management	SANtricity EMW (part of SANtricity Storage Manager) or Web Service proxy	SANtricity EMW (part of SANtricity Storage Manager) or Web Service proxy
Email Alerts, ASUP	Configure and dispatch from Event Monitor (part of SANtricity Storage Manager) on management station	Configure and dispatch on-box
SNMP	On management station, traps only	Embedded, adds GETs
REST APIs	Web Services proxy only	Embedded or Web Services proxy
CLI	On management station, all GUI operations +	On management station, no monitor configuration
In-band management	Full GUI and CLI, no REST	CLI and most EMW functions

Task	E2700/E5600/EF560, all other E-Series arrays except the E2800	E2800
Remote login	Telnet	SSH

### Management Use Cases

The following more detailed table is included for reference. The following list includes the acronyms used in the table:

- System Manager - SANtricity System Manager (embedded management software for an E2800 storage array)
- EMW - Enterprise Management Window (used for domain management)
- AMW - Array Management Window (storage management software for an E2700, E5600, or EF560)
- SMcli - SANtricity command line interface
- EMW script editor CLI - script editor in Enterprise Management Window
- REST - REST API endpoints, available in web proxy application for all arrays and natively for the E2800.

Task	E2700/E5600/EF560, all other E-Series arrays except the E2800	E2800
<b>Manage and discover</b>		
Discover an array in your management domain	<ul style="list-style-type: none"> <li>• EMW</li> </ul>	<ul style="list-style-type: none"> <li>• EMW</li> </ul>
Add or remove an array from your management domain	<ul style="list-style-type: none"> <li>• EMW</li> <li>• SMcli*</li> </ul>	<ul style="list-style-type: none"> <li>• EMW</li> <li>• SMcli*</li> </ul>
Launch SANtricity System Manager	<ul style="list-style-type: none"> <li>• N/A</li> </ul>	<ul style="list-style-type: none"> <li>• EMW</li> </ul>
Launch Array Management Window (AMW)	<ul style="list-style-type: none"> <li>• EMW</li> </ul>	<ul style="list-style-type: none"> <li>• N/A</li> </ul>
<b>AutoSupport and legacy support bundle collection</b>		
Enable/disable AutoSupport, AutoSupport OnDemand, and AutoSupport Remote Diagnostics features	<ul style="list-style-type: none"> <li>• EMW</li> <li>• SMcli*</li> </ul>	<ul style="list-style-type: none"> <li>• System Manager</li> <li>• REST</li> </ul>
Show AutoSupport logs for all or a select storage array	<ul style="list-style-type: none"> <li>• EMW</li> <li>• SMcli*</li> </ul>	<ul style="list-style-type: none"> <li>• System Manager</li> <li>• REST</li> </ul>
Enable or disable legacy support bundle collection for a select storage array	<ul style="list-style-type: none"> <li>• EMW</li> <li>• SMcli*</li> </ul>	<ul style="list-style-type: none"> <li>• N/A</li> </ul>

Task	E2700/E5600/EF560, all other E-Series arrays except the E2800	E2800
Specify legacy support bundle collection schedule	<ul style="list-style-type: none"> <li>• EMW</li> <li>• SMcli*</li> </ul>	<ul style="list-style-type: none"> <li>• N/A</li> </ul>
<b>Configuration and status</b>		
Display information (other than alert settings) about configured arrays	<ul style="list-style-type: none"> <li>• AMW</li> <li>• SMcli*</li> <li>• EMW script editor CLI</li> </ul>	<ul style="list-style-type: none"> <li>• System Manager</li> <li>• SMcli*</li> <li>• EMW script editor CLI</li> <li>• REST</li> </ul>
Show IP address of each array	<ul style="list-style-type: none"> <li>• AMW</li> <li>• SMcli*</li> <li>• EMW script editor CLI</li> </ul>	<ul style="list-style-type: none"> <li>• System Manager</li> <li>• SMcli*</li> <li>• EMW script editor CLI</li> <li>• REST</li> </ul>
Show WWN of each array	<ul style="list-style-type: none"> <li>• AMW</li> <li>• SMcli*</li> <li>• EMW script editor CLI</li> </ul>	<ul style="list-style-type: none"> <li>• System Manager</li> <li>• SMcli*</li> <li>• EMW script editor CLI</li> <li>• REST</li> </ul>
Show status of each array	<ul style="list-style-type: none"> <li>• AMW</li> <li>• SMcli</li> <li>• EMW script editor CLI</li> </ul>	<ul style="list-style-type: none"> <li>• System Manager</li> <li>• EMW (domain level status)</li> <li>• SMcli*</li> <li>• EMW script editor CLI</li> <li>• REST</li> </ul>
Set up remote volume mirroring groups and pairs	<ul style="list-style-type: none"> <li>• AMW</li> <li>• SMcli*</li> <li>• EMW script editor CLI</li> </ul>	<ul style="list-style-type: none"> <li>• System Manager (requires launch from EMW to correctly set up mirroring relationships)</li> <li>• SMcli*</li> <li>• EMW script editor CLI</li> <li>• REST</li> </ul>
Array-level configuration, provisioning, and tuning	<ul style="list-style-type: none"> <li>• AMW</li> <li>• SMcli*</li> <li>• EMW script editor CLI</li> </ul>	<ul style="list-style-type: none"> <li>• System Manager</li> <li>• REST</li> </ul>

Task	E2700/E5600/EF560, all other E-Series arrays except the E2800	E2800
<b>Alert and SNMP Configuration</b>		
Show global alert settings	<ul style="list-style-type: none"> <li>• EMW</li> <li>• SMcli*</li> <li>• EMW script editor CLI</li> </ul>	<ul style="list-style-type: none"> <li>• System Manager</li> <li>• REST</li> </ul>
Specify email server and other configuration for global alert settings	<ul style="list-style-type: none"> <li>• EMW</li> <li>• SMcli*</li> <li>• EMW script editor CLI</li> </ul>	<ul style="list-style-type: none"> <li>• System Manager</li> <li>• REST</li> </ul>
Remove an email from configuration for a specific array	<ul style="list-style-type: none"> <li>• EMW</li> <li>• SMcli</li> <li>• EMW script editor CLI</li> </ul>	<ul style="list-style-type: none"> <li>• System Manager</li> <li>• REST</li> </ul>
Add or remove SNMP trap information for a specific array	<ul style="list-style-type: none"> <li>• EMW</li> <li>• SMcli*</li> </ul>	<ul style="list-style-type: none"> <li>• System Manager</li> <li>• REST</li> </ul>
Send a test email based on global alert settings	<ul style="list-style-type: none"> <li>• EMW</li> <li>• SMcli*</li> </ul>	<ul style="list-style-type: none"> <li>• System Manager</li> <li>• REST</li> </ul>
<b>New features for E2800 only - see SANtricity System Manager online help for descriptions</b>		
Certificate handling - view SSL information, get a Certificate Signing Request (CSR), import a new certificate	<ul style="list-style-type: none"> <li>• N/A</li> </ul>	<ul style="list-style-type: none"> <li>• EMW</li> <li>• System Manager</li> <li>• REST</li> </ul>
More convenient syslog configuration for E2800	<ul style="list-style-type: none"> <li>• N/A</li> </ul>	<ul style="list-style-type: none"> <li>• System Manager</li> <li>• REST</li> </ul>
Save up to 30 days of historical statistical I/O data for E2800	<ul style="list-style-type: none"> <li>• N/A</li> </ul>	<ul style="list-style-type: none"> <li>• System Manager</li> <li>• REST</li> </ul>
Perform application tagging of volumes, snapshots for E2800	<ul style="list-style-type: none"> <li>• N/A</li> </ul>	<ul style="list-style-type: none"> <li>• System Manager</li> <li>• REST</li> </ul>

\* Requires installation of SANtricity Storage Manager

## Terminology differences in SANtricity 11.30

SANtricity 11.30 introduces, with the E2800 storage array, an embedded storage management software with new terminology. These terminology differences are noted in the following table.

### Software terms

The following table defines some of the most commonly used software terms. Note that different terms might apply to different hardware models.

Note that while new terminology is visible in SANtricity System Manager, some documents might still use the old terminology.

Term used for E2800 storage arrays / SANtricity System Manager	Term used for E2700 and E5600 storage arrays, and EF560 flash arrays / SANtricity Storage Manager	Definition
Allocated capacity	Configured, Used, or Physical capacity	The actual physical capacity allocated from the drives in a pool or volume group that is used for volumes and copy services operations.
Allocated capacity limit	Maximum capacity expansion	The quota or cap on how large the allocated physical capacity (repository) for a thin volume can grow.
Assign, unassign	Map, unmap	Associating a Logical Unit Number (LUN) with a host.
Default cluster	Default group	A system-defined entity that is only used if you want all the hosts connected to your storage array to have shared access to all of the volumes.  Normally, you will define individual hosts and then assign specific volumes to each of those hosts for restricted access.  Refer to the SANtricity System Manager online help for more details.
Event log	Major event log (MEL)	The Event log is a detailed record of events that occur in the storage array. You can use the Event log as a supplementary diagnostic tool to the Recovery Guru for tracing storage array events.

Term used for E2800 storage arrays / SANtricity System Manager	Term used for E2700 and E5600 storage arrays, and EF560 flash arrays / SANtricity Storage Manager	Definition
Host cluster	Host group	A host group or host cluster is a collection of hosts that share access to storage array volumes. Creating host groups lets administrators use the combined processing power of the host group computers to access the volumes. Because host groups also permit multiple hosts to access a single volume, the hosts manage this shared access.
Management client	Storage management station	For the E2800, any system that has a supported browser that can point to the storage array URL and issue management commands. If the system is performing I/O, it is a host system.  For the E2700 and E5600 storage arrays and the EF560 flash array, this is a system running the SANtricity Storage Manager application. This could also be a host system if it is performing I/O.
Objects	Elements or Components	Any managed item upon which a performance metric is being collected.
Pool	Disk Pool	Dynamic Disk Pool
Reported capacity	Volume capacity, Virtual capacity	The capacity that is reported up to the host and can accessed by the host.
Reserved capacity	Repository capacity	The physically allocated capacity that is used for any copy service operations or objects.
Reserved capacity threshold	Repository warning threshold	A generic term, used in context within the OLH.
SANtricity OS <i>n.n0</i>	Controller firmware <i>n.n0</i>	The version of bundled software code that runs on the controller.
Snapshot consistency group	Consistency group	A set of volumes that will have a snapshot image created at the same time for purposes of application consistency.
Thick volume	Standard volume	Describes any non-thin volume.
Unassigned capacity	Unconfigured capacity	Capacity that has not been configured yet into a pool or volume group.

Term used for E2800 storage arrays / SANtricity System Manager	Term used for E2700 and E5600 storage arrays, and EF560 flash arrays / SANtricity Storage Manager	Definition
Written capacity	Used capacity, Consumed capacity	The amount of capacity that has actually be written of the reserved capacity allocated for thin volumes and copy service objects or operations.

## Commands listed alphabetically

### Activate asynchronous mirroring

This command activates the Asynchronous Mirroring feature. After you activate the Asynchronous Mirroring 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

#### Related references

[Deactivate asynchronous mirroring](#) on page 176

### 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



**Related references**

- [Create iSCSI initiator](#) on page 131
- [Delete iSCSI initiator](#) on page 184
- [Diagnose controller iSCSI host cable](#) on page 193
- [Reset storage array iSCSI baseline](#) on page 252
- [Save storage array iSCSI statistics](#) on page 279
- [Set iSCSI initiator](#) on page 324
- [Set iSCSI target properties](#) on page 326
- [Show current iSCSI sessions](#) on page 390
- [Show storage array negotiation defaults](#) on page 423
- [Show storage array unconfigured iSCSI initiators](#) on page 424
- [Start iSCSI DHCP refresh](#) on page 453
- [Stop storage array iSCSI session](#) on page 479

**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
[healthCheckMelOverride=(TRUE | FALSE)]
```

**Parameter**

Parameter	Description
healthCheckMelOverride	<p>The setting to override the results of the health check of the major event log (MEL). MEL validation still occurs; it is not bypassed. If the MEL check fails, you can bypass the failure by using this parameter when running the command.</p> <p>Before any activation, the controller checks the event log to determine if an event has occurred that might prevent the new controller firmware from successfully activating. If such an event has occurred, the controller will normally not activate new firmware.</p> <p>This parameter forces the controller to activate new firmware. The default value is <code>FALSE</code>. Set this value to <code>TRUE</code> if you want to force the controller to activate new controller firmware.</p>

**Minimum firmware level**

8.10 adds the `healthCheckMelOverride` parameter.

**Related references**

- [Clear storage array firmware pending area](#) on page 101
- [Download storage array firmware/NVSRAM](#) on page 205
- [Save storage array firmware inventory](#) on page 278

## Activate Synchronous Mirroring

This command creates the mirror repository volume and activates the Synchronous Mirroring 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=(0 | 1 | 3 | 5 | 6)
repositoryDrives=(trayID1,drawerID1,slotID1 ...
trayIDn,drawerIDn,slotIDn)
[repositoryVolumeGroupUserLabel="volumeGroupName"]
[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"]
[trayLossProtect=(TRUE | FALSE)]
[drawerLossProtect=(TRUE | FALSE)]
[dataAssurance=(none | enabled)]
```

### Parameters

Parameter	Description
<code>repositoryRAIDLevel</code>	The RAID level for the mirror repository volume. Valid values are 0, 1, 3, 5, or 6.

Parameter	Description
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 24. Enclose the tray ID values, the drawer ID values, and the slot ID values in parentheses.
repositoryVolumeGroupUserLabel	The name that you want to give the new volume group in which the mirror repository volume will be located. Enclose the volume group name in double quotation marks (" ").
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.)
freeCapacityArea	<p>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:</p> <pre data-bbox="753 1157 1377 1213">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.
trayLossProtect	The setting to enforce tray loss protection when you create the mirror repository volume. To enforce tray loss protection, set this parameter to <code>TRUE</code> . The default value is <code>FALSE</code> .
drawerLossProtect	The setting to enforce drawer loss protection when you create the mirror repository volume. To enforce drawer loss protection, set this parameter to <code>TRUE</code> . The default value is <code>FALSE</code> .

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 DA-capable drives can be used for the volume group. These settings are valid:</p> <ul style="list-style-type: none"> <li>• none – The volume group does not have data assurance protection.</li> <li>• enabled – The volume group has data assurance protection. The volume group supports protected information and is formatted with protection information enabled.</li> </ul> <p><b>Note:</b> This option is only valid for enablement if the drives support DA.</p>

## 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 determines whether 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.

**Minimum firmware level**

7.10 adds RAID Level 6 capability.

7.60 adds the `drawerIDuser` input, the `driveMediaType` parameter, and the `drawerLossProtect` parameter.

7.75 adds the `dataAssurance` parameter.

8.10 removes the `driveMediaType` parameter.

**Related references**

[Deactivate synchronous mirroring](#) on page 176

**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 name of the SSD cache to which you want to add SSDs. Enclose the name in square brackets ([ ]). If the SSD cache name contains special characters or consists only of numbers, you must enclose the name 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 24 . 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

**Related references**

[Change SSD cache application type](#) on page 90

[Create SSD cache](#) on page 159

[Delete SSD cache](#) on page 188

[Remove drives from SSD cache](#) on page 235

[Rename SSD cache](#) on page 242

[Resume SSD cache](#) on page 260

[Set SSD cache for a volume](#) on page 342  
[Show SSD cache](#) on page 406  
[Show SSD cache statistics](#) on page 407  
[Start SSD cache locate](#) on page 457  
[Start SSD cache performance modeling](#) on page 458  
[Stop SSD cache locate](#) on page 475  
[Stop SSD cache performance modeling](#) on page 476  
[Suspend SSD cache](#) on page 482

## 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 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"]
addCGMemberVirtualDisk="baseVirtualDiskName"
repositoryVirtualDisk="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 double quotation marks ( " "). If the specified volume is an existing repository volume or an existing snapshot volume, the command fails.

Parameter	Description
repositoryVolume	<p>This parameter performs two functions:</p> <ul style="list-style-type: none"> <li>• In an existing consistency group that has a repository volume, this parameter identifies the repository volume.</li> <li>• When creating a new repository volume this parameter identifies either a volume group or disk pool in which to create the new repository volume.</li> </ul> <p>Enclose the volume name in double quotation marks (" ").</p>
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 parentheses around the volume group name and capacity, or the disk pool name and capacity.

### Minimum firmware level

7.83

### Related references

- [Create consistency group](#) on page 107
- [Create consistency group snapshot image](#) on page 115
- [Create consistency group snapshot volume](#) on page 116
- [Create consistency group snapshot volume mapping](#) on page 120
- [Delete consistency group](#) on page 178
- [Delete consistency group snapshot image](#) on page 179
- [Delete consistency group snapshot volume](#) on page 180
- [Remove member volume from consistency group](#) on page 237
- [Resume consistency group snapshot volume](#) on page 256
- [Set consistency group attributes](#) on page 296
- [Set consistency group snapshot volume](#) on page 298
- [Show consistency group](#) on page 384
- [Show consistency group snapshot image](#) on page 386
- [Start consistency group snapshot rollback](#) on page 443
- [Stop consistency group snapshot rollback](#) on page 466
- [Stop consistency group snapshot volume](#) on page 468
- [Stop pending snapshot images on consistency group](#) on page 471

## 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]))
```



**Parameter**

Parameter	Description
volume	The name of the primary volume that you want to add to the asynchronous mirror group. Enclose the volume name in double quotation marks (" ").
asyncMirrorGroup	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 (" ").
remotePassword	This parameter is optional. Use this parameter when the remote storage array is password protected. Enclose the password in double quotation marks (" ").
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"> <li>• Use an existing repository volume: name.</li> <li>• Create a new repository volume when you run this command.</li> </ul> <p>The name of an existing repository volume is comprised of two parts:</p> <ul style="list-style-type: none"> <li>• The term repos</li> <li>• A four digit numerical identifier that the storage management software assigns to the repository volume name</li> </ul> <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"> <li>• An integer value that represents a percentage of the base volume capacity</li> <li>• A decimal fraction value that represents a percentage of the base volume capacity</li> <li>• A specific size for the repository volume. Size is defined in units of <code>bytes</code>, <code>KB</code>, <code>MB</code>, <code>GB</code>, or <code>TB</code>.</li> </ul> <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>

**Notes**

- The Asynchronous Mirroring 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

**Related references**[Create asynchronous mirror group](#) on page 104[Delete asynchronous mirror group](#) on page 177[Remove volume from asynchronous mirror group](#) on page 239[Show asynchronous mirror groups](#) on page 382**Autoconfigure storage array**

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)]
[secureDrives=(fips | fde)]
[dataAssurance=(none | enabled)]
```

## Parameters

Parameter	Description
driveType	<p>The type of drives that you want to use for the storage array. You must use this parameter when you have more than one type of drive in your storage array.</p> <p>These drive types are valid:</p> <ul style="list-style-type: none"> <li>• fibre</li> <li>• SATA</li> <li>• SAS</li> </ul> <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. You must use this parameter when you have more than one type of drive media in your storage array.</p> <p>These drive media types are valid:</p> <ul style="list-style-type: none"> <li>• hdd – Use this option when you have hard drives.</li> <li>• ssd – Use this option when you have solid state disks.</li> <li>• unknown – Use if you are not sure what types of drive media are in the drive tray.</li> </ul> <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 4 (SSD only) 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 FALSE. To turn on cache read prefetch, set this parameter to TRUE.
readAheadMultiplier	<p>This parameter defines how many additional data blocks are read into cache. Valid values range from 0 to 65535.</p> <p><b>Note:</b> This parameter is deprecated and will be removed in a future release of storage management software. For best operation use the cacheReadPrefetch parameter.</p>

Parameter	Description
<code>securityType</code>	<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"> <li><code>none</code> – The volume group and volumes are not secure.</li> <li><code>capable</code> – The volume group and volumes are capable of having security set, but security has not been enabled.</li> <li><code>enabled</code> – The volume group and volumes have security enabled.</li> </ul>
<code>secureDrives</code>	<p>The type of secure drives to use in the volume group. These settings are valid:</p> <ul style="list-style-type: none"> <li><code>fips</code> – To use FIPS compliant drives only.</li> <li><code>fde</code> – To use FDE compliant drives.</li> </ul> <p><b>Important:</b> Use this parameter along with the <code>securityType</code> parameter. If you specify <code>none</code> for the <code>securityType</code> parameter, the value of the <code>secureDrives</code> parameter is ignored, because non-secure volume groups do not need to have secure drive types specified.</p>
<code>dataAssurance</code>	<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"> <li><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.</li> <li><code>enabled</code> – The volume group has data assurance protection. The volume group supports protected information and is formatted with protection information enabled. This is the default, if this parameter is not specified.</li> </ul> <p><b>Note:</b> This option is only valid for enablement if the drives support DA.</p>

### Drives and volume groups

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.
- You cannot mix HDD and SSD drives 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, 3, 5, and 6.
  - A volume group with RAID level 3, RAID level 5, or RAID level 6 cannot have more than 30 drives and must have a minimum of three 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.
- To enable tray/drawer loss protection, refer to the following tables for additional criteria:

Level	Criteria for Tray Loss Protection	Minimum number of trays required
Disk Pool	The disk pool contains no more than two drives in a single tray.	6
RAID 6	The volume group contains no more than two drives in a single tray.	3
RAID 3 or RAID 5	Each drive in the volume group is located in a separate tray.	3
RAID 1	Each drive in a RAID 1 pair must be located in a separate tray.	2
RAID 0	Cannot achieve Tray Loss Protection.	Not applicable

Level	Criteria for Drawer Loss Protection	Minimum number of drawers required
Disk Pool	The disk pool includes drives from all five drawers and there are an equal number of drives in each drawer. A 60-drive tray can achieve Drawer Loss Protection when the disk pool contains 15, 20, 25, 30, 35, 40, 45, 50, 55, or 60 drives.	5
RAID 6	The volume group contains no more than two drives in a single drawer.	3
RAID 3 or RAID 5	Each drive in the volume group is located in a separate drawer.	3
RAID 1	Each drive in a mirrored pair must be located in a separate drawer.	2
RAID 0	Cannot achieve Drawer Loss Protection.	Not applicable

## Hot spares

With volume groups, a valuable strategy to protect data is to assign available drives in the storage array as hot spare drives. A hot spare is a drive, containing no data, that acts as a standby in the storage array in case a drive fails in a RAID 1, RAID 3, RAID 5, or RAID 6 volume group. The hot spare adds another level of redundancy to the storage array.

Generally, hot spare drives must have capacities that are equal to or greater than the used capacity on the drives that they are protecting. Hot spare drives must be of the same media type, the same interface type, and the same capacity as the drives that they are protecting.

If a drive fails in the storage array, the hot spare is normally substituted automatically for the failed drive without requiring your intervention. If a hot spare is available when a drive fails, the controller uses redundancy data parity to reconstruct the data onto the hot spare. Data evacuation support also allows data to be copied to a hot spare before the software marks the drive "failed."

After the failed drive is physically replaced, you can use either of the following options to restore the data:

When you have replaced the failed drive, the data from the hot spare is copied back to the replacement drive. This action is called copyback.

If you designate the hot spare drive as a permanent member of a volume group, the copyback operation is not needed.

The availability of tray loss protection and drawer loss protection for a volume group depends on the location of the drives that comprise the volume group. Tray loss protection and drawer loss protection might be lost because of a failed drive and the location of the hot spare drive. To make sure that tray loss protection and drawer loss protection are not affected, you must replace a failed drive to initiate the copyback process.

The storage array automatically selects Data Assurance (DA)-capable drives for hot spare coverage of DA-enabled volumes.

Make sure you have DA-capable drives in the storage array for hot spare coverage of DA-enabled volumes. For more information about DA-capable drives, refer to Data Assurance feature.

Secure-capable (FIPS and FDE) drives can be used as a hot spare for both secure-capable and non-secure-capable drives. Non-secure-capable drives can provide coverage for other non-secure-capable drives, and for secure-capable drives if the volume group does not have the security enabled. A FIPS volume group can only use a FIPS drive as a hot spare; however, you can use a FIPS hot spare for non-secure-capable, secure-capable, and secure-enabled volume groups.

If you do not have a hot spare, you can still replace a failed drive while the storage array is operating. If the drive is part of a RAID 1, RAID 3, RAID 5, or RAID 6 volume group, the controller uses redundancy data parity to automatically reconstruct the data onto the replacement drive. This action is called reconstruction.

## 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. Valid values for the `cacheReadPrefetch` parameter are `TRUE` or `FALSE`. The default is `TRUE`.

### Security type

Use the `securityType` parameter to specify the security settings for the storage array.

Before you can set the `securityType` parameter to `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`
- `export storageArray securityKey`
- `import storageArray securityKey`
- `set storageArray securityKey`
- `enable volumeGroup [volumeGroupName] security`
- `enable diskPool [diskPoolName] security`

### Secure drives

Secure-capable drives can be either Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives. Use the `secureDrives` parameter to specify the type of secure drives to use. The values you can use are `fips` and `fde`.

### Example command

```
autoConfigure storageArray securityType=capable secureDrives=fips;
```

### Minimum firmware level

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

7.50 adds the `securityType` parameter.

7.75 adds the `dataAssurance` parameter.

8.25 adds the `secureDrives` parameter.

### Related references

[Autoconfigure storage array hot spares](#) on page 88

[Clear storage array configuration](#) on page 99

[Show storage array auto configuration](#) on page 414

## 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

### Related references

[Clear storage array configuration](#) on page 99

[Show storage array auto configuration](#) on page 414

## 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
<code>asyncMirrorGroup</code>	The name of the asynchronous mirror group for which or 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 contains special characters or consists only of numbers, you must enclose the name in double quotation marks (" ") inside square brackets.

### Minimum firmware level

7.84

### Related references

[Check asynchronous mirror group consistency](#) on page 91



[Resume asynchronous mirror group](#) on page 255

[Set asynchronous mirror group](#) on page 294

[Show asynchronous mirror group synchronization progress](#) on page 383

[Start asynchronous mirroring synchronization](#) on page 440

[Suspend asynchronous mirror group](#) on page 481

## Capture or view an AutoSupport log

The `SMcli -autoSupportLog` command allows you to view an AutoSupport log file. This file provides information about status, history of transmission activity, and any errors encountered during delivery of the AutoSupport messages. The log is available for all AutoSupport-capable and AutoSupport-enabled storage arrays.

### Description

This command allows you to view two types of logs:

- Current log- View the log captured at this point in time.
- Archive log- View the log from an archived file.

### Context

**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.

**Restriction:** This command is obsolete for the E2800 storage array. Use the RESTful API, SANtricity System Manager, or see [Using curl commands with the E2800 storage array](#) on page 39.

### Syntax

```
SMcli -autoSupportLog (all|-n "storageArrayName" | -w "wwID")
inputArchive=n outputLog=filename
```

### Parameters

The following table lists the command parameters.

Parameter	Description
all	Use this parameter if you want to create a transmission log for all of the storage arrays in the management domain with AutoSupport capability that have been AutoSupport enabled.
n	The name for the storage array for which you want to view an AutoSupport log.
w	The World Wide Identifier (WWID) of the storage array for which you want to view an AutoSupport log.
inputArchive	Allows you to specify the archived AutoSupport log, where the archive file is <i>n</i> , an integer from 0 to 5.  <b>Note:</b> Omitting this parameter means you select the current AutoSupport log (captured at this point in time).

Parameter	Description
outputLog	Allows you to specify the output AutoSupport log filename.  <b>Note:</b> This parameter is required.

### Minimum Firmware Level

8.25

### Example

```
SMcli -autosupportLog -n StorageArrayName inputArchive=0
outputLog=ASUPTransmissionLog
```

### Verification

View the AutoSupport log you requested, which is in the client directory where the storage management software was installed.

### Related references

[Using curl commands with the E2800 storage array](#) on page 39

[Display AutoSupport bundle collection settings](#) on page 198

[Display AutoSupport message collection schedule](#) on page 200

[Disable AutoSupport at the EMW level SMcli version](#) on page 197

[Enable or disable AutoSupport at the EMW management domain level \(SMcli only\)](#) on page 210

[Enable or disable the AutoSupport OnDemand feature at the EMW management domain level \(SMcli only\)](#) on page 211

[Enable or disable the AutoSupport OnDemand Remote Diagnostics feature at the EMW management domain level \(SMcli only\)](#) on page 213

[Reset AutoSupport message collection schedule](#) on page 247

[Save storage array support data](#) on page 283

[Specify the AutoSupport delivery method](#) on page 438

[Test the AutoSupport configuration](#) on page 485

## 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)
```

**Parameter**

Parameter	Description
ssdCache	The name of the SSD cache for which you want to change the application type. Enclose the SSD cache name in square brackets ([ ]). If the SSD cache name contains special characters or consists only of numbers, you must enclose the SSD cache name 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 <code>webServer</code> , <code>dataBase</code> , or <code>fileSystem</code> .

**Minimum firmware level**

7.84

**Related references**

- [Add drives to SSD cache](#) on page 77
- [Create SSD cache](#) on page 159
- [Delete SSD cache](#) on page 188
- [Remove drives from SSD cache](#) on page 235
- [Rename SSD cache](#) on page 242
- [Resume SSD cache](#) on page 260
- [Set SSD cache for a volume](#) on page 342
- [Show SSD cache](#) on page 406
- [Show SSD cache statistics](#) on page 407
- [Start SSD cache locate](#) on page 457
- [Start SSD cache performance modeling](#) on page 458
- [Stop SSD cache locate](#) on page 475
- [Stop SSD cache performance modeling](#) on page 476
- [Suspend SSD cache](#) on page 482

**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 name of the asynchronous mirror group on which to run a consistency check. Enclose the asynchronous mirror group name in square brackets ([ ]). If the asynchronous mirror group name contains special characters or consists only of numbers, you also must enclose the name in double quotation marks (" ") inside square brackets.

Parameter	Description
repositoryConsistency	Specifies that the asynchronous mirror group is checked for consistency.
localVolume	Specifies the name 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. If the local volume name contains special characters or consists only of numbers, you also must enclose the name in double quotation marks ( " ") inside 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: <pre>file="C:\Program Files\CLI\logs \repoConsistency.txt"</pre> 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

7.83

### Related references

[Test asynchronous mirror group connectivity](#) on page 484

## 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.

Parameter	Description
snapGroup	The name of the snapshot group on which to run a consistency check. Enclose the snapshot group name in square brackets ([ ]). If the snapshot group name contains special characters or consists only of numbers, you also must enclose the name in double quotation marks (" ") inside square brackets.
snapVolume	The name of the snapshot volume on which to run a consistency check. Enclose the snapshot volume name in square brackets ([ ]). If the snapshot volume name contains special characters or consists only of numbers, you also must enclose the name in double quotation marks (" ") inside square brackets.
volume	The name of the thin volume on which to run a consistency check. Enclose the thin volume name in square brackets ([ ]). If the thin volume name contains special characters or consists only of numbers, you also must enclose the name in double quotation marks (" ") inside square brackets.
volumeCopy	Specifies that the target of a volume copy relationship is checked for consistency.
target	Use only in conjunction with the volumeCopy parameter to specify the name of the volume that is the target of the volume copy relationship. If the volume name contains special characters or consists only of numbers, you also must enclose the name in double quotation marks (" ") inside square brackets.
asyncMirrorGroup	The name of the asynchronous mirror group on which to run a consistency check. Enclose the asynchronous mirror group name in square brackets ([ ]). If the asynchronous mirror group name contains special characters or consists only of numbers, you also must enclose the name in double quotation marks (" ") inside square brackets.
localVolume	Use only with the asyncMirrorGroup parameter to specify the name 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 name in double quotation marks (" ").
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: <pre>file="C:\Program Files\CLI\logs \repoConsistency.txt"</pre> 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

7.83

## 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 (" ").
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

### Related references

[Test asynchronous mirror group connectivity](#) on page 484

## 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
<code>localVolume</code>	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 must enclose the volume name in double quotation marks (" ").
<code>timeout</code>	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."

### Related references

[Diagnose synchronous mirroring](#) on page 196

[Show synchronous mirroring volume synchronization progress](#) on page 427

## 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=LBValue]
[endingLBA=LBValue]
[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 contains special characters or consists only of numbers, you must enclose the identifier 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.
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 .The lowest priority rate benefits system performance, but the parity check takes longer. The highest priority rate benefits the parity check, but system performance might be compromised.
startingLBA	The starting logical block address.
endingLBA	The ending logical block address.
verbose	The setting to capture progress details, such as percent complete, and to show the information as the volume parity is being scanned. 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.

**Related references**

[Clear volume unreadable sectors](#) on page 103

[Repair volume parity](#) on page 243

**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 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 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 contains special characters or consists only of numbers, you must enclose the name in double quotation marks (" ") inside square brackets.
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"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Enclose each of the names in double quotation marks (" ").</li> <li>• Separate each of the names with a space.</li> </ul>
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 contains special characters or consists only of numbers, you must enclose the name in double quotation marks (" ") inside square brackets.

Parameter	Description
volumes	<p>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:</p> <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Enclose each of the names in double quotation marks (" ").</li> <li>• Separate each of the names with a space.</li> </ul>

### Minimum firmware level

7.84

### Related references

[Check asynchronous mirror group consistency](#) on page 91

[Check storage array connectivity](#) on page 94

[Show asynchronous mirror group synchronization progress](#) on page 383

[Start asynchronous mirroring synchronization](#) on page 440

## Clear drive channel statistics

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

### Syntax

```
clear allDriveChannels stats
```

### Parameters

None.

### Related references

[Diagnose controller iSCSI host cable](#) on page 193

[Enable controller data transfer](#) on page 208

[Reset controller](#) on page 248

[Save controller NVSRAM](#) on page 267

[Save drive channel fault isolation diagnostic status](#) on page 267

[Set controller](#) on page 299

[Set controller DNS settings](#) on page 304

[Set controller NTP settings](#) on page 305

[Set controller service action allowed indicator](#) on page 306

[Set drive channel status](#) on page 313

[Set host channel](#) on page 321

[Show controller](#) on page 387

[Show controller NVSRAM](#) on page 390

[Show drive channel statistics](#) on page 395

[Start controller trace](#) on page 444

[Start drive channel fault isolation diagnostics](#) on page 447

[Start drive channel locate](#) on page 449

[Start iSCSI DHCP refresh](#) on page 453

[Stop drive channel fault isolation diagnostics](#) on page 469

[Stop drive channel locate](#) on page 470

## 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. You are prompted for the storage array password for this command.

### Syntax

```
clear storageArray configuration (all | volumeGroups)
[password="storageArrayPassword"]
```

### 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.
password	The password for the storage array. Enclose the password in double quotation marks (" ").

### 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.

This command also resets the storage array identifier. Any premium features enabled with a key file are invalidated, and new premium feature keys must be created and applied. Contact technical support.

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

7.10 adds these parameters:

- all
- volumeGroups

7.83 adds the recovery mode capability.

### Related references

[Clear storage array recovery mode](#) on page 101

[Load storage array DBM database](#) on page 223

[Save storage array DBM database](#) on page 273

[Save storage array DBM validator information file](#) on page 274

[Show storage array DBM database](#) on page 420

[Start storage array configuration database diagnostic](#) on page 459

[Stop storage array configuration database diagnostic](#) on page 478

[Autoconfigure storage array](#) on page 82

[Autoconfigure storage array hot spares](#) on page 88

[Show storage array auto configuration](#) on page 414

[Start storage array configuration database diagnostic](#) on page 459

[Stop storage array configuration database diagnostic](#) on page 478

## Clear storage array controller health image

**Note:** With firmware version 8.20 the `coreDumpAllowOverWrite` parameter is deprecated and will be removed in a future release of storage management software. For best operation use the `controllerHealthImageAllowOverWrite` parameter.

This command sets a flag on a controller to allow a new controller health image to overwrite an existing controller health image.

### Syntax

```
clear storageArray controllerHealthImageAllowOverWrite
```

### Parameters

None.

### Notes

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

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

### Minimum firmware level

7.83

8.20 replaces the `coreDumpAllowOverWrite` parameter with the `set storageArray controllerHealthImageAllowOverWrite` parameter.

## 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.

### Related references

[Using curl commands with the E2800 storage array](#) on page 39

[Save storage array events](#) on page 276

[Set event alert filtering](#) on page 316

[Test alerts](#) on page 484

[Show blocked events](#) on page 384

[Save storage array events](#) on page 276

## 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.

## Clear storage array recovery mode

This command forces a storage array to exit recovery mode.

## Syntax

```
clear storageArray recoveryMode
```

## Parameters

None.

## Notes

**Note:** 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.

**Caution:** If you want to return the storage array to a previous configuration, you must restore the configuration from the backup before clearing recovery mode. You must perform validation checks or check with technical support to make sure that the restore was successful. After determining that the restore was successful, recovery mode can be cleared.

## Minimum firmware level

7.83

## Related references

[Clear storage array configuration](#) on page 99  
[Load storage array DBM database](#) on page 223  
[Save storage array DBM database](#) on page 273  
[Save storage array DBM validator information file](#) on page 274  
[Show storage array DBM database](#) on page 420  
[Start storage array configuration database diagnostic](#) on page 459  
[Stop storage array configuration database diagnostic](#) on page 478  
[Clear storage array configuration](#) on page 99

## 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.

Parameter	Description
volume	The name of a volume for which you are clearing persistent reservations. Enclose the volume name in square brackets ([ ]). If the volume name contains special characters or consists only of numbers, you must enclose the name in double quotation marks (" ") inside square brackets.
volumes	The names of several volumes for which you are clearing persistent reservations. Enter the names of the member volumes using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Enclose each of the names in double quotation marks (" ").</li> <li>• Separate each of the names with a space.</li> </ul>

### Notes

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

### Related references

[Remove volume LUN mapping](#) on page 240

[Set volume mapping](#) on page 380

[Show volume reservations](#) on page 437

## 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	The name of a volume for which you are clearing unreadable sector information. Enclose the volume name in square brackets ([ ]). If the volume name contains special characters or consists only of numbers, you must enclose the name in double quotation marks (" ") inside square brackets.
volumes	The names of several volumes for which you are clearing unreadable sector information. Enter the names of the member volumes using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Enclose each of the names in double quotation marks (" ").</li> <li>• Separate each of the names with a space.</li> </ul>

**Notes**

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

**Related references**

[Check volume parity](#) on page 95

[Repair volume parity](#) on page 243

**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")
[remotePassword="password"]
interfaceType=(FC | iSCSI)
[syncInterval=integer(minutes | hours | days)]
[warningSyncThreshold=integer(minutes | hours | days)]
[warningRecoveryThreshold=integer(minutes | hours | days)]
[warningThresholdPercent=percentValue]
[autoResync=(TRUE | FALSE)]
```

**Parameters**

Parameter	Description
userLabel	The name of the new asynchronous mirror group that you want to create. Enclose the new asynchronous mirror group name in double quotation marks (" "). The name must be unique on the local and remote storage arrays.



Parameter	Description
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 (&lt; &gt;).</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>
syncInterval	<p>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.</p> <p><b>Note:</b> Do not add a space between the integer and the length of time.</p> <p>Example: 10minutes</p>
warningSyncThreshold	<p>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.</p> <p><b>Note:</b> Do not add a space between the integer and the length of time.</p> <p>Example: 30minutes</p>

Parameter	Description
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> <p><b>Note:</b> You must set the Recovery Point Threshold to be twice as large as the synchronization interval threshold.</p> <p><b>Note:</b> Do not add a space between the integer and the length of time.</p> <p>Example: 60minutes</p>
warningThresholdPercent	<p>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.</p>
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"> <li>• <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.</li> <li>• <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.</li> </ul>

### Notes

- The Asynchronous Mirroring 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 a 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

### Related references

[Add volume to asynchronous mirror group](#) on page 80

[Delete asynchronous mirror group](#) on page 177

[Establish asynchronous mirrored pair](#) on page 216

[Remove incomplete asynchronous mirrored pair from asynchronous mirror group](#) on page 236

[Remove volume from asynchronous mirror group](#) on page 239

[Reset asynchronous mirror group statistics](#) on page 246

[Resume asynchronous mirror group](#) on page 255

[Save asynchronous mirror group statistics](#) on page 265

[Set asynchronous mirror group](#) on page 294

[Show asynchronous mirror groups](#) on page 382

[Suspend asynchronous mirror group](#) on page 481

[Test asynchronous mirror group connectivity](#) on page 484

## Create consistency group

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
<code>userLabel</code>	The name of the new consistency group that you want to create. 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> .

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

### 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 snapshots. 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 table explains how to use the options for the `schedule` parameter:

Parameter	Description
<code>schedule</code>	Required for specifying schedule parameters.
<code>immediate</code>	Start the operation immediately. This item is mutually exclusive with any other scheduling parameters.
<code>enableSchedule</code>	When set to <code>true</code> , scheduling is turned on. When set to <code>false</code> , scheduling is turned off.  <b>Note:</b> The default is <code>false</code> .
<code>startDate</code>	A specific date on which to start the operation. The format for entering the date is <code>MM:DD:YY</code> . The default is the current date. An example of this option is <code>startDate=06:27:11</code> .

Parameter	Description
<code>scheduleDay</code>	<p>A day of the week on which to start the operation. Can either be all or one or more of the following values:</p> <ul style="list-style-type: none"> <li>• <code>monday</code></li> <li>• <code>tuesday</code></li> <li>• <code>wednesday</code></li> <li>• <code>thursday</code></li> <li>• <code>friday</code></li> <li>• <code>saturday</code></li> <li>• <code>sunday</code></li> </ul> <p><b>Note:</b> Enclose the value in parentheses. For example, <code>scheduleDay=(wednesday)</code>.</p> <p>More than one day can be specified by enclosing the days in a single set of parentheses and separating each day with a space. For example, <code>scheduleDay=(monday wednesday friday)</code>.</p> <p><b>Note:</b> This parameter is not compatible with a monthly schedule.</p>
<code>startTime</code>	<p>The time of a day on which to start the operation. The format for entering the time is HH:MM, where HH is the hour and MM is the minute past the hour. Uses a 24-hour clock. For example, 2:00 in the afternoon is 14:00. An example of this option is <code>startTime=14:27</code>.</p>
<code>scheduleInterval</code>	<p>An amount of time, in minutes, to have as a minimum between operations.</p> <p>Schedule interval should not be more than 1440 (24 hours) and it should be a multiple of 30.</p> <p>An example of this option is <code>scheduleInterval=180</code>.</p>
<code>endDate</code>	<p>A specific date on which to stop the operation. The format for entering the date is MM:DD:YY. If no end date is desired, you can specify <code>noEndDate</code>. An example of this option is <code>endDate=11:26:11</code>.</p>
<code>timesPerDay</code>	<p>The number of times to perform the operation in a day. An example of this option is <code>timesPerDay=4</code>.</p>
<code>timezone</code>	<p>Specifies the timezone to be used for the schedule. Can be specified in two ways:</p> <p><b>GMT±HH:MM</b></p> <p>The time zone offset from GMT. Example: <code>timezone=GMT-06:00</code>.</p> <p><b>Text String</b></p> <p>Standard time zone text string, must be enclosed in quotes. Example:<code>timezone="Central Time (US &amp; Canada)"</code></p>

Parameter	Description
scheduleDate	<p>A day of the month on which to perform the operation. The values for the days are numerical and in the range of 1-31.</p> <p><b>Note:</b> This parameter is not compatible with a weekly schedule.</p> <p>An example of the <code>scheduleDate</code> option is <code>scheduleDate=( "15" )</code>.</p>
month	<p>A specific month on which to perform the operation. The values for the months are:</p> <ul style="list-style-type: none"> <li>• jan - January</li> <li>• feb - February</li> <li>• mar - March</li> <li>• apr - April</li> <li>• may - May</li> <li>• jun - June</li> <li>• jul - July</li> <li>• aug - August</li> <li>• sep - September</li> <li>• oct - October</li> <li>• nov - November</li> <li>• dec - December</li> </ul> <p><b>Note:</b> Enclose the value in parentheses. For example, <code>month=( jan )</code>.</p> <p>More than one month can be specified by enclosing the months in a single set of parentheses and separating each month with a space. For example, <code>month=( jan jul dec )</code>.</p> <p>Use this parameter with the <code>scheduleDate</code> parameter to perform the operation on a specific day of the month.</p> <p><b>Note:</b> This parameter is not compatible with a weekly schedule.</p>

This table explains how to use the `timeZone` parameter:

Timezone Name	GMT offset
International Date Line West	GMT-12:00
Coordinated Universal Time-11	GMT-11:00
Hawaii	GMT-10:00
Alaska	GMT-09:00
Baja California	GMT-08:00
Pacific Time (US & Canada)	GMT-08:00
Arizona	GMT-07:00

<b>Timezone Name</b>	<b>GMT offset</b>
Chihuahua, La Paz, Mazatlan	GMT-07:00
Mountain Time (US & Canada)	GMT-07:00
Central America	GMT-06:00
Central Time (US & Canada)	GMT-06:00
Guadalajara, Mexico City, Monterrey	GMT-06:00
Saskatchewan	GMT-06:00
Bogota, Lima, Quito	GMT-05:00
Eastern Time (US & Canada)	GMT-05:00
Indiana (East)	GMT-05:00
Caracas	GMT-04:30
Asuncion	GMT-04:00
Atlantic Time (Canada)	GMT-04:00
Cuiaba	GMT-04:00
Georgetown, La Paz, Manaus, San Juan	GMT-04:00
Santiago	GMT-04:00
Newfoundland	GMT-03:30
Brasilia	GMT-03:00
Buenos Aires	GMT-03:00
Cayenne, Fortaleza	GMT-03:00
Greenland	GMT-03:00
Montevideo	GMT-03:00
Mid-Atlantic	GMT-02:00
Azores	GMT-01:00
Cape Verde Is.	GMT-01:00
Casablanca	GMT
Coordinated Universal Time	GMT
Dublin, Edinburgh, Lisbon, London	GMT
Monrovia, Reykjavik	GMT
Amsterdam, Berlin, Bern, Rome, Stockholm, Vienna	GMT+01:00
Belgrade, Bratislava, Budapest, Ljubljana, Prague	GMT+01:00
Brussels, Copenhagen, Madrid, Paris	GMT+01:00
Sarajevo, Skopje, Warsaw, Zagreb	GMT+01:00
West Central Africa	GMT+01:00



<b>Timezone Name</b>	<b>GMT offset</b>
Windhoek	GMT+02:00
Amman	GMT+02:00
Athens, Bucharest, Istanbul	GMT+02:00
Beirut	GMT+02:00
Cairo	GMT+02:00
Damascus	GMT+02:00
Harare, Pretoria	GMT+02:00
Helsinki, Kyiv, Riga, Sofia, Tallinn, Vilnius	GMT+02:00
Jerusalem	GMT+02:00
Minsk	GMT+02:00
Baghdad	GMT+03:00
Kuwait, Riyadh	GMT+03:00
Moscow, St. Petersburg, Volgograd	GMT+03:00
Nairobi	GMT+03:00
Tehran	GMT+03:00
Abu Dhabi, Muscat	GMT+04:00
Baku	GMT+04:00
Port Louis	GMT+04:00
Tbilisi	GMT+04:00
Yerevan	GMT+04:00
Kabul	GMT+04:30
Ekaterinburg	GMT+05:00
Islamabad, Karachi	GMT+05:00
Tashkent	GMT+05:00
Chennai, Kolkata, Mumbai, New Delhi	GMT+05:30
Sri Jayawardenepura	GMT+05:30
Kathmandu	GMT+05:45
Astana	GMT+06:00
Dhaka	GMT+06:00
Novosibirsk	GMT+06:00
Yangon (Rangoon)	GMT+06:30
Bangkok, Hanoi, Jakarta	GMT+07:00
Krasnoyarsk	GMT+07:00

Timezone Name	GMT offset
Beijing, Chongqing, Hong Kong, Urumqi	GMT+08:00
Irkutsk	GMT+08:00
Kuala Lumpur, Singapore	GMT+08:00
Perth	GMT+08:00
Taipei	GMT+08:00
Ulaanbaatar	GMT+08:00
Osaka, Sapporo, Tokyo	GMT+09:00
Seoul	GMT+09:00
Yakutsk	GMT+09:00
Adelaide	GMT+09:30
Darwin	GMT+09:30
Brisbane	GMT+10:00
Canberra, Melbourne, Sydney	GMT+10:00
Guam, Port Moresby	GMT+10:00
Hobart	GMT+10:00
Vladivostok	GMT+10:00
Magadan, Solomon Is., New Caledonia	GMT+11:00
Auckland, Wellington	GMT+12:00
Coordinated Universal Time+12	GMT+12:00
Fiji, Marshall Is.	GMT+12:00
Petropavlovsk-Kamchatsky - Old	GMT+12:00
Nuku'alofa	GMT+13:00

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.

### Related references

- [Add member to consistency group](#) on page 78
- [Create consistency group snapshot image](#) on page 115
- [Create consistency group snapshot volume](#) on page 116
- [Create consistency group snapshot volume mapping](#) on page 120
- [Delete consistency group](#) on page 178
- [Delete consistency group snapshot image](#) on page 179
- [Delete consistency group snapshot volume](#) on page 180
- [Remove member volume from consistency group](#) on page 237
- [Resume consistency group snapshot volume](#) on page 256
- [Set consistency group attributes](#) on page 296
- [Set consistency group snapshot volume](#) on page 298
- [Show consistency group](#) on page 384
- [Show consistency group snapshot image](#) on page 386
- [Start consistency group snapshot rollback](#) on page 443
- [Stop consistency group snapshot rollback](#) on page 466
- [Stop consistency group snapshot volume](#) on page 468
- [Stop pending snapshot images on consistency group](#) on page 471

## 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

**Related references**

*Add member to consistency group* on page 78  
*Create consistency group* on page 107  
*Create consistency group snapshot volume* on page 116  
*Create consistency group snapshot volume mapping* on page 120  
*Delete consistency group* on page 178  
*Delete consistency group snapshot image* on page 179  
*Delete consistency group snapshot volume* on page 180  
*Remove member volume from consistency group* on page 237  
*Resume consistency group snapshot volume* on page 256  
*Set consistency group attributes* on page 296  
*Set consistency group snapshot volume* on page 298  
*Show consistency group* on page 384  
*Show consistency group snapshot image* on page 386  
*Start consistency group snapshot rollback* on page 443  
*Stop consistency group snapshot rollback* on page 466  
*Stop consistency group snapshot volume* on page 468  
*Stop pending snapshot images on consistency group* on page 471

**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 (" ").
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"> <li>The name of the consistency group</li> <li>An identifier for the snapshot image in the consistency group.</li> </ul> <p>The identifier for the snapshot image can be one of these:</p> <ul style="list-style-type: none"> <li>newest - Use this option when you want to show the latest snapshot image created in the consistency group.</li> <li>oldest - Use this option when you want to show the earliest snapshot image created in the consistency group.</li> </ul> <p>Enclose the snapshot image name in double quotation marks (" "). You can enter more than one snapshot image name or sequence number. Enclose all of the snapshot image names inside parentheses. Separate each snapshot image name with a 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"> <li>• Use an existing repository volume: name</li> <li>• Create a new repository volume when you run this command</li> </ul> <p>The name of an existing repository volume is comprised of two parts:</p> <ul style="list-style-type: none"> <li>• The term repos</li> <li>• A four digit numerical identifier that the storage management software assigns to the repository volume name</li> </ul> <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"> <li>• An integer value that represents a percentage of the base volume capacity</li> <li>• A decimal fraction value that represents a percentage of the base volume capacity</li> <li>• A specific size for the repository volume. Size is defined in units of bytes , KB , MB , GB , or TB .</li> </ul> <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 `cg1`, `cg2`, and `cg3`. The repository volumes already exist and selected by the user in this command.

```
create cgSnapVolume userLabel="cgSnapVolume1"
cgSnapImageID="snapCG1:oldest"
members=(cg1:repos_0010 cg2:repos_0011 cg3: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 `cg1:repos_10`, member `cg1` maps to repository volume `repos_0010`.

- Creating a read/write consistency group snapshot volume on a snapshot consistency group named "snapCG1" of only members `cg1` and `cg2`:

```
create cgSnapVolume userLabel="cgSnapVolume2"
cgSnapImageID="snapCG1:14214"
members=(cg1:repos_1000 cg2:repos_1001);
```

- Creating a read-only consistency group snapshot volume on a snapshot consistency group named `snapCG1` that has three members `cg1`, `cg2`, and `cg3`:

```
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 `cg1`, `cg2`, and `cg3`:

```
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 `cg1`, `cg2`, and `cg3`:

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

## Minimum firmware level

7.83

**Related references**

- [Add member to consistency group](#) on page 78
- [Create consistency group](#) on page 107
- [Create consistency group snapshot image](#) on page 115
- [Create consistency group snapshot volume mapping](#) on page 120
- [Delete consistency group](#) on page 178
- [Delete consistency group snapshot image](#) on page 179
- [Delete consistency group snapshot volume](#) on page 180
- [Remove member volume from consistency group](#) on page 237
- [Resume consistency group snapshot volume](#) on page 256
- [Set consistency group attributes](#) on page 296
- [Set consistency group snapshot volume](#) on page 298
- [Show consistency group](#) on page 384
- [Show consistency group snapshot image](#) on page 386
- [Start consistency group snapshot rollback](#) on page 443
- [Stop consistency group snapshot rollback](#) on page 466
- [Stop consistency group snapshot volume](#) on page 468
- [Stop pending snapshot images on consistency group](#) on page 471

**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="cgSnapVolumeName"
(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

### Related references

[Add member to consistency group](#) on page 78  
[Create consistency group](#) on page 107  
[Create consistency group snapshot image](#) on page 115  
[Create consistency group snapshot volume](#) on page 116  
[Delete consistency group](#) on page 178  
[Delete consistency group snapshot image](#) on page 179  
[Delete consistency group snapshot volume](#) on page 180  
[Remove member volume from consistency group](#) on page 237  
[Resume consistency group snapshot volume](#) on page 256  
[Set consistency group attributes](#) on page 296  
[Set consistency group snapshot volume](#) on page 298  
[Show consistency group](#) on page 384  
[Show consistency group snapshot image](#) on page 386  
[Start consistency group snapshot rollback](#) on page 443  
[Stop consistency group snapshot rollback](#) on page 466  
[Stop consistency group snapshot volume](#) on page 468  
[Stop pending snapshot images on consistency group](#) on page 471

## 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)
userLabel="diskPoolName"
[driveType=SAS]
[driveCount=driveCountValue]
[trayLossProtect=(TRUE | FALSE)]
[drawerLossProtect=(TRUE | FALSE)]
[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 )]
[secureDrives=(fips | fde )]
[driveMediaType=(hdd | ssd | allMedia | unknown)]
[dataAssurance=(none|enabled)]
```

**Parameters**

Parameter	Description
drives	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 24. Enclose the tray ID values, the drawer ID values, and the slot ID values in parentheses.
driveType	The type of drive that you want to use in the disk pool. You cannot mix drive types. This drive type is valid: <ul style="list-style-type: none"> <li>SAS</li> </ul> <p><b>Note:</b> If you do not specify a drive type, the command defaults to SAS. Because only one drive type is currently supported, this parameter is retained for future additions. There is no need to specify it.</p>
userLabel	The name that you want to give the new disk pool. Enclose the disk pool name in double quotation marks (" ").
driveCount	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.
trayLossProtect	The setting to enforce Tray Loss Protection when you create the disk pool. To enforce Tray Loss Protection, set this parameter to <code>TRUE</code> . The default value is <code>FALSE</code> .
drawerLossProtect	The setting to enforce Drawer Loss Protection when you create the disk pool. To enforce Drawer Loss Protection, set this parameter to <code>TRUE</code> . The default value is <code>FALSE</code> .
warningThreshold	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. Valid values are from 0 to 100. The default value is 50. Setting this parameter to 0 (zero) disables warning alerts. If you set this to <code>default</code> , the warning alert threshold value is determined by the controller firmware.

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 (zero) 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.</p> <p>These settings are valid:</p> <ul style="list-style-type: none"> <li>• none – The volume candidates are not secure.</li> <li>• capable – The volume candidates are capable of having security set, but security has not been enabled.</li> <li>• enabled – The volume candidates have security enabled.</li> </ul> <p>The default value is none.</p>

Parameter	Description
secureDrives	<p>The type of secure drives to use in the volume group. These settings are valid:</p> <ul style="list-style-type: none"> <li>• <code>fips</code> – To use FIPS compliant drives only.</li> <li>• <code>fde</code> – To use FDE compliant drives.</li> </ul> <p><b>Important:</b> Use this parameter along with the <code>securityType</code> parameter. If you specify <code>none</code> for the <code>securityType</code> parameter, the value of the <code>secureDrives</code> parameter is ignored, because non-secure disk pools do not need to have secure drive types specified.</p> <p><b>Note:</b> This parameter is ignored unless you are also using the <code>driveCount</code> parameter. If you are specifying the drives to use for the disk pool instead of providing a count, specify the appropriate type of drives in the selection list based on the security type you desire.</p>
driveMediaType	<p>The type of drive media that you want to use for the disk pool. You must use this parameter when you have more than one type of drive media in your storage array.</p> <p>These drive media are valid:</p> <ul style="list-style-type: none"> <li>• <code>hdd</code> – Use this option when you have hard drives.</li> <li>• <code>ssd</code> – Use this option when you have solid-state disks.</li> <li>• <code>unknown</code> – Use this option if you are not sure what types of drive media are in the drive tray.</li> <li>• <code>allMedia</code> – Use this option when you want to use all types of drive media that are in the drive tray.</li> </ul> <p>The default value is <code>hdd</code>.</p> <p><b>Note:</b> The controller firmware does not mix <code>hdd</code> and <code>ssd</code> drive media in the same disk pool, regardless of using the setting you select.</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.</p> <p>These settings are valid:</p> <ul style="list-style-type: none"> <li>• <code>none</code> – The disk pool does not have data assurance protection.</li> <li>• <code>enabled</code> – The disk pool has data assurance protection. The disk pool supports protected information and is formatted with protection information enabled.</li> </ul> <p>The default value is <code>none</code>.</p> <p><b>Note:</b> This option is only valid for enablement if the drives support DA.</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.

## Tray Loss Protection and Drawer Loss Protection

Level	Criteria for Tray Loss Protection	Minimum number of trays required
Disk Pool	The disk pool contains no more than two drives in a single tray.	6
RAID 6	The volume group contains no more than two drives in a single tray.	3
RAID 3 or RAID 5	Each drive in the volume group is located in a separate tray.	3
RAID 1	Each drive in a RAID 1 pair must be located in a separate tray.	2
RAID 0	Cannot achieve Tray Loss Protection.	Not applicable

Level	Criteria for Drawer Loss Protection	Minimum number of drawers required
Disk Pool	The disk pool includes drives from all five drawers and there are an equal number of drives in each drawer. A 60-drive tray can achieve Drawer Loss Protection when the disk pool contains 15, 20, 25, 30, 35, 40, 45, 50, 55, or 60 drives.	5
RAID 6	The volume group contains no more than two drives in a single drawer.	3

Level	Criteria for Drawer Loss Protection	Minimum number of drawers required
RAID 3 or RAID 5	Each drive in the volume group is located in a separate drawer.	3
RAID 1	Each drive in a mirrored pair must be located in a separate drawer.	2
RAID 0	Cannot achieve Drawer Loss Protection.	Not applicable

**Note:** If you have a storage array configuration in which a disk pool 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 can set the `drawerLossProtect` parameter to `TRUE` or leave it unspecified, but you cannot set it to `FALSE`.

### 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 as follows:

- **Warning** – This is the first level of alert. This level 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. This level 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 always must 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 supported background operations have the following relative priority levels:

1. Reconstruction
2. Format
3. IAF

## 4. DCE

### Security type

Use the `securityType` parameter to specify the security settings for the storage array.

Before you can set the `securityType` parameter to `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`
- `export storageArray securityKey`
- `import storageArray securityKey`
- `set storageArray securityKey`
- `enable volumeGroup [volumeGroupName] security`
- `enable diskPool [diskPoolName] security`

### Secure drives

Secure-capable drives can be either Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives. Use the `secureDrives` parameter to specify the type of secure drives to use. The values you can use are `fips` and `fde`.

### Example command

```
create diskPool driveType=SAS userLabel="FIPS_Pool" driveCount=11
securityType=capable secureDrives=fips;
```

### Minimum firmware level

7.83

8.20 adds these parameters:

- `trayLossProtect`
- `drawerLossProtect`

8.25 adds the `secureDrives` parameter.

### Related references

[Delete disk pool](#) on page 181

## 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.

**Related references**

- [Create host group](#) on page 128
- [Create host port](#) on page 129
- [Delete host](#) on page 182
- [Delete host group](#) on page 182
- [Delete host port](#) on page 183
- [Set host](#) on page 320
- [Set host channel](#) on page 321
- [Set host group](#) on page 322
- [Set host port label](#) on page 323
- [Show host ports](#) on page 398
- [Show storage array host topology](#) on page 421

**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

## Related references

- [Create host](#) on page 127
- [Create host port](#) on page 129
- [Delete host](#) on page 182
- [Delete host group](#) on page 182
- [Delete host port](#) on page 183
- [Set host](#) on page 320
- [Set host channel](#) on page 321
- [Set host group](#) on page 322
- [Set host port label](#) on page 323
- [Show host ports](#) on page 398
- [Show storage array host topology](#) on page 421

## 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 (" ").
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"> <li>• FC – Fibre Channel</li> <li>• SAS – Serial-Attached SCSI</li> <li>• IB – Infiniband</li> <li>• iSCSI</li> </ul> <p>An FC or a SAS selection requires an 8-byte WWID. An IB selection requires a 16-byte group identifier (gid). An iSCSI selection requires an IQN.</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.

**Related references**

[Create host](#) on page 127

[Create host group](#) on page 128

[Delete host](#) on page 182

[Delete host group](#) on page 182

[Delete host port](#) on page 183

[Set host](#) on page 320

[Set host channel](#) on page 321

[Set host group](#) on page 322

[Set host port label](#) on page 323

[Show host ports](#) on page 398

[Show storage array host topology](#) on page 421

## 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 (" ").
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

### Related references

[Activate iSCSI initiator](#) on page 72

[Delete iSCSI initiator](#) on page 184

[Diagnose controller iSCSI host cable](#) on page 193

[Reset storage array iSCSI baseline](#) on page 252

[Save storage array iSCSI statistics](#) on page 279

[Set iSCSI initiator](#) on page 324

[Set iSCSI target properties](#) on page 326

[Show current iSCSI sessions](#) on page 390

[Show storage array negotiation defaults](#) on page 423

[Show storage array unconfigured iSCSI initiators](#) on page 424

[Start iSCSI DHCP refresh](#) on page 453

[Stop storage array iSCSI session](#) on page 479

## 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 \(manual drive select\)](#) on page 140.

### Syntax

```
create volume driveCount=numberOfDrives
[volumeGroupUserLabel="volumeGroupName"]
raidLevel=(0 | 1 | 3 | 5 | 6)
[userLabel="volumeName"]
driveMediaType=(HDD | SSD | unknown | allMedia)
[driveType= 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)]
[secureDrives=(fips | fde )]
[dataAssurance=(none | enabled)]
```

### Parameters

Parameter	Description
<code>driveCount</code>	The number of unassigned drives that you want to use in the volume group.
<code>volumeGroupUserLabel</code>	The name that you want to give the new volume group. Enclose the new volume group name in double quotation marks (" ").  <b>Note:</b> If you do not specify a user label for the volume group, the controller firmware assigns it a number.
<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 to the new volume. Enclose the new volume name in double quotation marks (" ").  <b>Note:</b> If you do not specify a user label for the volume, the controller firmware assigns it a number.

Parameter	Description
driveMediaType	<p>The type of drive media that you want to use for the volume group. These drive media are valid:</p> <ul style="list-style-type: none"> <li>• HDD – Use this option when you have hard drives in the drive tray.</li> <li>• SSD – Use this option when you have solid state drives in the drive tray.</li> <li>• unknown – Use this option if you are not sure what types of drive media are in the drive tray.</li> <li>• allMedia – Use this option when you want to use all types of drive media that are in the drive tray.</li> </ul>
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>These drive types are valid:</p> <ul style="list-style-type: none"> <li>• fibre</li> <li>• SATA</li> <li>• SAS</li> </ul> <p>If you do not specify a drive type, the command defaults to SAS.</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>

Parameter	Description
<code>drawerLossProtect</code>	The setting to enforce drawer loss protection when you create the mirror repository volume group. To enforce drawer loss protection, set this parameter to <code>TRUE</code> . The default value is <code>FALSE</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: <ul style="list-style-type: none"> <li>• <code>none</code> – The volume group and volumes are not secure.</li> <li>• <code>capable</code> – The volume group and volumes are capable of having security set, but security has not been enabled.</li> <li>• <code>enabled</code> – The volume group and volumes have security enabled.</li> </ul>
<code>secureDrives</code>	The type of secure drives to use. These settings are valid: <ul style="list-style-type: none"> <li>• <code>fips</code> – To use FIPS compliant drives only.</li> <li>• <code>fde</code> – To use FDE compliant drives.</li> </ul> <p><b>Important:</b> Use this parameter along with the <code>securityType</code> parameter. If you specify <code>none</code> for the <code>securityType</code> parameter, the value of the <code>secureDrives</code> parameter is ignored, because non-secure volume groups do not need to have secure drive types specified.</p>
<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"> <li>• <code>none</code> – The volume group does not have data assurance protection. This is the default.</li> <li>• <code>enabled</code> – The volume group has data assurance protection. The volume group supports protected information and is formatted with protection information enabled.</li> </ul> <p><b>Note:</b> This option is only valid for enablement if the drives support DA.</p>

## 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.

### 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.

### Usage Hint

**Note:** 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. The segment size and cache read prefetch settings for various usage hints are shown in the following table:

Usage hint	Segment size setting	Dynamic cache read prefetch setting
File system	128 KB	Enabled
Database	128 KB	Enabled
Multimedia	256 KB	Enabled

### 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. Valid values for the `cacheReadPrefetch` parameter are `TRUE` or `FALSE`. The default is `TRUE`.

### Security type

Use the `securityType` parameter to specify the security settings for the storage array.

Before you can set the `securityType` parameter to `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`

- `export storageArray securityKey`
- `import storageArray securityKey`
- `set storageArray securityKey`
- `enable volumeGroup [volumeGroupName] security`
- `enable diskPool [diskPoolName] security`

### Secure drives

Secure-capable drives can be either Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives. Use the `secureDrives` parameter to specify the type of secure drives to use. The values you can use are `fips` and `fde`.

### Tray loss protection and drawer loss protection

To enable tray/drawer loss protection, refer to the following tables for additional criteria:

Level	Criteria for Tray Loss Protection	Minimum number of trays required
Disk Pool	The disk pool contains no more than two drives in a single tray.	6
RAID 6	The volume group contains no more than two drives in a single tray.	3
RAID 3 or RAID 5	Each drive in the volume group is located in a separate tray.	3
RAID 1	Each drive in a RAID 1 pair must be located in a separate tray.	2
RAID 0	Cannot achieve Tray Loss Protection.	Not applicable

Level	Criteria for Drawer Loss Protection	Minimum number of drawers required
Disk Pool	The disk pool includes drives from all five drawers and there are an equal number of drives in each drawer. A 60-drive tray can achieve Drawer Loss Protection when the disk pool contains 15, 20, 25, 30, 35, 40, 45, 50, 55, or 60 drives.	5
RAID 6	The volume group contains no more than two drives in a single drawer.	3
RAID 3 or RAID 5	Each drive in the volume group is located in a separate drawer.	3
RAID 1	Each drive in a mirrored pair must be located in a separate drawer.	2
RAID 0	Cannot achieve Drawer Loss Protection.	Not applicable

### Example command

```
create volume driveCount=2 volumeGroupUserLabel="FIPS_VG" raidLevel=1
userLabel="FIPS_V" driveMediaType=HDD securityType=capable
secureDrives=fips
```



### Minimum firmware level

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.

8.25 adds the `secureDrives` parameter.

### Related references

[Create RAID volume \(free extent based select\)](#) on page 137

[Create RAID volume \(manual drive select\)](#) on page 140

## 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
<code>volumeGroup</code>	The name of a specific volume group in your storage array. Enclose the volume group name in double quotation marks (" ").
<code>userLabel</code>	The name that you want to give the new volume. Enclose the new volume name in double quotation marks (" ").
<code>freeCapacityArea</code>	<p>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:</p> <pre>freeCapacityArea=2</pre> <p><b>Note:</b> Run the <code>show volumeGroup</code> command to determine whether the free capacity area exists.</p> <p><b>Note:</b> If this parameter is not specified, the volume is created in the lowest-numbered free capacity area with sufficient capacity for the volume.</p>

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.
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 <code>TRUE</code> . To turn off cache read prefetch, set this parameter to <code>FALSE</code> .
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.
usageHint	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> .
dssPreAllocate	The setting to make sure that reserve capacity is allocated for future segment size increases. The default value is <code>TRUE</code> .
securityType	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.
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"> <li><code>none</code> – The volume group does not have data assurance protection.</li> <li><code>enabled</code> – The volume group has data assurance protection. The volume group supports protected information and is formatted with protection information enabled.</li> </ul> <p><b>Note:</b> This option is only valid for enablement if the drives are DA-capable.</p>

## 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.

### Usage Hint

**Note:** 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. The segment size and cache read prefetch settings for various usage hints are shown in the following table:

Usage hint	Segment size setting	Dynamic cache read prefetch setting
File system	128 KB	Enabled
Database	128 KB	Enabled
Multimedia	256 KB	Enabled

### 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. Valid values for the `cacheReadPrefetch` parameter are `TRUE` or `FALSE`. The default is `TRUE`.

### Security type

Use the `securityType` parameter to specify the security settings for the storage array.

Before you can set the `securityType` parameter to `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`

- `export storageArray securityKey`
- `import storageArray securityKey`
- `set storageArray securityKey`
- `enable volumeGroup [volumeGroupName] security`
- `enable diskPool [diskPoolName] security`

### Minimum firmware level

7.10 adds the `dssPreAllocate` parameter.

7.50 adds the `securityType` parameter.

7.75 adds the `dataAssurance` parameter.

### Related references

[Create RAID volume \(automatic drive select\)](#) on page 132

[Create RAID volume \(manual drive select\)](#) on page 140

## Create RAID volume (manual drive select)

This command creates a new volume group and volume, and lets you specify the drives for the 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 24 . Enclose the tray ID values, the drawer ID values, and the slot ID values in parentheses.

Parameter	Description
volumeGroupUserLabel	The name that you want to give the new volume group. Enclose the volume group name in double quotation marks (" ").  <b>Note:</b> If you do not specify a user label for the volume group, the controller firmware assigns it a number.
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 the new volume. Enclose the new volume name in double quotation marks (" ").  <b>Note:</b> If you do not specify a user label for the volume, the controller firmware assigns it a number.
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 off cache read prefetch, set this parameter to FALSE. To turn on cache read prefetch, set this parameter to TRUE.
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 4 (SSD only) 8, 16, 32, 64, 128, 256, or 512.
usageHint	The settings for both the 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.
trayLossProtect	The setting to enforce tray loss protection when you create the repository. 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 mirrored repository volume. To enforce drawer loss protection, set this parameter to TRUE. The default value is FALSE.
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"> <li>• none – The volume group and volumes are not secure.</li> <li>• capable – The volume group and volumes are capable of having security set, but security has not been enabled.</li> <li>• enabled – The volume group and volumes have security enabled.</li> </ul>

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"> <li>• <code>none</code> – The volume group does not have data assurance protection.</li> <li>• <code>enabled</code> – The volume group has data assurance protection. The volume group supports protected information and is formatted with protection information enabled.</li> </ul> <p><b>Note:</b> This option is only valid for enablement if the drives are DA-capable.</p>

### 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 level 1:

- There are an even number of drives in the group, by RAID 1 definition
- The first half of the group, in listed order (and in stripe order), are primary drives
- The second half of the group, in listed order (and in stripe order), are the corresponding mirror drives

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.

### Tray loss protection and drawer loss protection

For tray loss protection to work, your configuration must adhere to the following guidelines:

Level	Criteria for Tray Loss Protection	Minimum number of trays required
Disk Pool	The disk pool contains no more than two drives in a single tray.	6
RAID 6	The volume group contains no more than two drives in a single tray.	3

Level	Criteria for Tray Loss Protection	Minimum number of trays required
RAID 3 or RAID 5	Each drive in the volume group is located in a separate tray.	3
RAID 1	Each drive in a RAID 1 pair must be located in a separate tray.	2
RAID 0	Cannot achieve Tray Loss Protection.	Not applicable

For drawer loss protection to work (in high density enclosure environments), your configuration must adhere to the following guidelines:

Level	Criteria for Drawer Loss Protection	Minimum number of drawers required
Disk Pool	The disk pool includes drives from all five drawers and there are an equal number of drives in each drawer. A 60-drive tray can achieve Drawer Loss Protection when the disk pool contains 15, 20, 25, 30, 35, 40, 45, 50, 55, or 60 drives.	5
RAID 6	The volume group contains no more than two drives in a single drawer.	3
RAID 3 or RAID 5	Each drive in the volume group is located in a separate drawer.	3
RAID 1	Each drive in a mirrored pair must be located in a separate drawer.	2
RAID 0	Cannot achieve Drawer Loss Protection.	Not applicable

### 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.

### Usage Hint

**Note:** 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. The segment size and cache read prefetch settings for various usage hints are shown in the following table:

Usage hint	Segment size setting	Dynamic cache read prefetch setting
File system	128 KB	Enabled
Database	128 KB	Enabled
Multimedia	256 KB	Enabled

### 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. Valid values for the `cacheReadPrefetch` parameter are `TRUE` or `FALSE`. The default is `TRUE`.

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

Use the `securityType` parameter to specify the security settings for the storage array.

Before you can set the `securityType` parameter to `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`
- `export storageArray securityKey`
- `import storageArray securityKey`
- `set storageArray securityKey`
- `enable volumeGroup [volumeGroupName] security`
- `enable diskPool [diskPoolName] security`

### Minimum firmware level

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

7.50 adds the `securityType` parameter.

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

7.75 adds the `dataAssurance` parameter.

### Related references

[Create RAID volume \(automatic drive select\)](#) on page 132

[Create RAID volume \(free extent based select\)](#) on page 137



## 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 (" ").
<code>snapImageID</code>	<p>The name of a snapshot image that you want to add to the new read-only snapshot volume. The name of a snapshot image is comprised of two parts:</p> <ul style="list-style-type: none"> <li>The name of the snapshot group</li> <li>An identifier for the snapshot image in the snapshot group</li> </ul> <p>The identifier for the snapshot image can be one of these:</p> <ul style="list-style-type: none"> <li>An integer value that is the sequence number of the snapshot in the snapshot group.</li> <li><code>NEWEST</code> - Use this option when you want to show the latest snapshot image created in the snapshot group.</li> <li><code>OLDEST</code> - Use this option when you want to show the earliest snapshot image created in the snapshot group.</li> </ul> <p>Enclose the snapshot image name in double quotation marks (" ").</p>
<code>readOnly</code>	<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 <code>TRUE</code>.</p> <p><b>Note:</b> With the <code>readOnly</code> parameter, snapshot volumes are not created.</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 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

### Related references

- [Create snapshot volume](#) on page 156
- [Delete snapshot volume](#) on page 187
- [Rename snapshot volume](#) on page 242
- [Resume snapshot volume](#) on page 259
- [Revive snapshot volume](#) on page 263
- [Set read-only snapshot volume to read/write volume](#) on page 328
- [Set snapshot volume media scan](#) on page 339
- [Set snapshot volume repository volume capacity](#) on page 340
- [Set SSD cache for a volume](#) on page 342
- [Show snapshot volumes](#) on page 403
- [Stop snapshot volume](#) on page 474

## 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"> <li>• Use an existing repository volume: name</li> <li>• Create a new repository volume when you run this command</li> </ul> <p>The name of an existing repository volume is comprised of two parts:</p> <ul style="list-style-type: none"> <li>• The term repos</li> <li>• A four digit numerical identifier that the storage management software assigns to the repository volume name</li> </ul> <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"> <li>• An integer value that represents a percentage of the base volume capacity</li> <li>• A decimal fraction value that represents a percentage of the base volume capacity</li> <li>• A specific size for the repository volume. Size is defined in units of bytes , KB , MB , GB , or TB .</li> </ul> <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>
rollBackPriority	<p>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>.</p>

Parameter	Description
<code>repositoryFullLimit</code>	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.
<code>autoDeleteLimit</code>	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.
<code>enableSchedule</code>	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> .
<code>schedule</code>	<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"> <li>• <code>immediate</code></li> <li>• <code>startDate</code></li> <li>• <code>scheduleDay</code></li> <li>• <code>startTime</code></li> <li>• <code>scheduleInterval</code></li> <li>• <code>endDate</code></li> <li>• <code>timesPerDay</code></li> <li>• <code>timeZone</code></li> <li>• <code>scheduleDate</code></li> <li>• <code>month</code></li> </ul> <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 table explains how to use the options for the `schedule` parameter:

Parameter	Description
<code>schedule</code>	Required for specifying schedule parameters.
<code>immediate</code>	Start the operation immediately. This item is mutually exclusive with any other scheduling parameters.
<code>enableSchedule</code>	When set to <code>true</code> , scheduling is turned on. When set to <code>false</code> , scheduling is turned off.  <b>Note:</b> The default is <code>false</code> .
<code>startDate</code>	A specific date on which to start the operation. The format for entering the date is MM:DD:YY. The default is the current date. An example of this option is <code>startDate=06:27:11</code> .
<code>scheduleDay</code>	A day of the week on which to start the operation. Can either be all or one or more of the following values: <ul style="list-style-type: none"> <li>• <code>monday</code></li> <li>• <code>tuesday</code></li> <li>• <code>wednesday</code></li> <li>• <code>thursday</code></li> <li>• <code>friday</code></li> <li>• <code>saturday</code></li> <li>• <code>sunday</code></li> </ul> <b>Note:</b> Enclose the value in parentheses. For example, <code>scheduleDay=(wednesday)</code> .  More than one day can be specified by enclosing the days in a single set of parentheses and separating each day with a space. For example, <code>scheduleDay=(monday wednesday friday)</code> .  <b>Note:</b> This parameter is not compatible with a monthly schedule.
<code>startTime</code>	The time of a day on which to start the operation. The format for entering the time is HH:MM, where HH is the hour and MM is the minute past the hour. Uses a 24-hour clock. For example, 2:00 in the afternoon is 14:00. An example of this option is <code>startTime=14:27</code> .
<code>scheduleInterval</code>	An amount of time, in minutes, to have as a minimum between operations.  Schedule interval should not be more than 1440 (24 hours) and it should be a multiple of 30.  An example of this option is <code>scheduleInterval=180</code> .
<code>endDate</code>	A specific date on which to stop the operation. The format for entering the date is MM:DD:YY. If no end date is desired, you can specify <code>noEndDate</code> . An example of this option is <code>endDate=11:26:11</code> .
<code>timesPerDay</code>	The number of times to perform the operation in a day. An example of this option is <code>timesPerDay=4</code> .

Parameter	Description
timezone	<p>Specifies the timezone to be used for the schedule. Can be specified in two ways:</p> <p><b>GMT±HH:MM</b></p> <p>The time zone offset from GMT. Example:  <code>timezone=GMT-06:00.</code></p> <p><b>Text String</b></p> <p>Standard time zone text string, must be enclosed in quotes.  Example:<code>timezone="Central Time (US &amp; Canada)"</code></p>
scheduleDate	<p>A day of the month on which to perform the operation. The values for the days are numerical and in the range of 1-31.</p> <p><b>Note:</b> This parameter is not compatible with a weekly schedule.</p> <p>An example of the <code>scheduleDate</code> option is <code>scheduleDate="15"</code>.</p>
month	<p>A specific month on which to perform the operation. The values for the months are:</p> <ul style="list-style-type: none"> <li>• <code>jan</code> - January</li> <li>• <code>feb</code> - February</li> <li>• <code>mar</code> - March</li> <li>• <code>apr</code> - April</li> <li>• <code>may</code> - May</li> <li>• <code>jun</code> - June</li> <li>• <code>jul</code> - July</li> <li>• <code>aug</code> - August</li> <li>• <code>sep</code> - September</li> <li>• <code>oct</code> - October</li> <li>• <code>nov</code> - November</li> <li>• <code>dec</code> - December</li> </ul> <p><b>Note:</b> Enclose the value in parentheses. For example,  <code>month=(jan)</code>.</p> <p>More than one month can be specified by enclosing the months in a single set of parentheses and separating each month with a space. For example, <code>month=(jan jul dec)</code>.</p> <p>Use this parameter with the <code>scheduleDate</code> parameter to perform the operation on a specific day of the month.</p> <p><b>Note:</b> This parameter is not compatible with a weekly schedule.</p>

This table explains how to use the `timeZone` parameter:

<b>Timezone Name</b>	<b>GMT offset</b>
International Date Line West	GMT-12:00
Coordinated Universal Time-11	GMT-11:00
Hawaii	GMT-10:00
Alaska	GMT-09:00
Baja California	GMT-08:00
Pacific Time (US & Canada)	GMT-08:00
Arizona	GMT-07:00
Chihuahua, La Paz, Mazatlan	GMT-07:00
Mountain Time (US & Canada)	GMT-07:00
Central America	GMT-06:00
Central Time (US & Canada)	GMT-06:00
Guadalajara, Mexico City, Monterrey	GMT-06:00
Saskatchewan	GMT-06:00
Bogota, Lima, Quito	GMT-05:00
Eastern Time (US & Canada)	GMT-05:00
Indiana (East)	GMT-05:00
Caracas	GMT-04:30
Asuncion	GMT-04:00
Atlantic Time (Canada)	GMT-04:00
Cuiaba	GMT-04:00
Georgetown, La Paz, Manaus, San Juan	GMT-04:00
Santiago	GMT-04:00
Newfoundland	GMT-03:30
Brasilia	GMT-03:00
Buenos Aires	GMT-03:00
Cayenne, Fortaleza	GMT-03:00
Greenland	GMT-03:00
Montevideo	GMT-03:00
Mid-Atlantic	GMT-02:00
Azores	GMT-01:00
Cape Verde Is.	GMT-01:00
Casablanca	GMT
Coordinated Universal Time	GMT
Dublin, Edinburgh, Lisbon, London	GMT



<b>Timezone Name</b>	<b>GMT offset</b>
Monrovia, Reykjavik	GMT
Amsterdam, Berlin, Bern, Rome, Stockholm, Vienna	GMT+01:00
Belgrade, Bratislava, Budapest, Ljubljana, Prague	GMT+01:00
Brussels, Copenhagen, Madrid, Paris	GMT+01:00
Sarajevo, Skopje, Warsaw, Zagreb	GMT+01:00
West Central Africa	GMT+01:00
Windhoek	GMT+02:00
Amman	GMT+02:00
Athens, Bucharest, Istanbul	GMT+02:00
Beirut	GMT+02:00
Cairo	GMT+02:00
Damascus	GMT+02:00
Harare, Pretoria	GMT+02:00
Helsinki, Kyiv, Riga, Sofia, Tallinn, Vilnius	GMT+02:00
Jerusalem	GMT+02:00
Minsk	GMT+02:00
Baghdad	GMT+03:00
Kuwait, Riyadh	GMT+03:00
Moscow, St. Petersburg, Volgograd	GMT+03:00
Nairobi	GMT+03:00
Tehran	GMT+03:00
Abu Dhabi, Muscat	GMT+04:00
Baku	GMT+04:00
Port Louis	GMT+04:00
Tbilisi	GMT+04:00
Yerevan	GMT+04:00
Kabul	GMT+04:30
Ekaterinburg	GMT+05:00
Islamabad, Karachi	GMT+05:00
Tashkent	GMT+05:00
Chennai, Kolkata, Mumbai, New Delhi	GMT+05:30
Sri Jayawardenepura	GMT+05:30
Kathmandu	GMT+05:45

<b>Timezone Name</b>	<b>GMT offset</b>
Astana	GMT+06:00
Dhaka	GMT+06:00
Novosibirsk	GMT+06:00
Yangon (Rangoon)	GMT+06:30
Bangkok, Hanoi, Jakarta	GMT+07:00
Krasnoyarsk	GMT+07:00
Beijing, Chongqing, Hong Kong, Urumqi	GMT+08:00
Irkutsk	GMT+08:00
Kuala Lumpur, Singapore	GMT+08:00
Perth	GMT+08:00
Taipei	GMT+08:00
Ulaanbaatar	GMT+08:00
Osaka, Sapporo, Tokyo	GMT+09:00
Seoul	GMT+09:00
Yakutsk	GMT+09:00
Adelaide	GMT+09:30
Darwin	GMT+09:30
Brisbane	GMT+10:00
Canberra, Melbourne, Sydney	GMT+10:00
Guam, Port Moresby	GMT+10:00
Hobart	GMT+10:00
Vladivostok	GMT+10:00
Magadan, Solomon Is., New Caledonia	GMT+11:00
Auckland, Wellington	GMT+12:00
Coordinated Universal Time+12	GMT+12:00
Fiji, Marshall Is.	GMT+12:00
Petropavlovsk-Kamchatsky - Old	GMT+12:00
Nuku'alofa	GMT+13:00

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.

### Related references

[Delete snapshot group](#) on page 185

[Revive snapshot group](#) on page 263

[Set snapshot group attributes](#) on page 331

[Set snapshot group media scan](#) on page 333

[Set snapshot group repository volume capacity](#) on page 334

[Set snapshot group schedule](#) on page 336

[Set SSD cache for a volume](#) on page 342

[Show snapshot group](#) on page 399

[Stop snapshot group pending snapshot images](#) on page 472

## 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	The name of the snapshot group that will contain the snapshot image. Enclose the name of the snapshot group in double quotation marks (" ").
snapGroups	The names of several snapshot groups that will contain the snapshot image. Enter the names of the snapshot groups using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in in parentheses.</li> <li>• Enclose each of the names in double quotation marks (" ").</li> <li>• Separate each of the names with a space.</li> </ul>

## 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

## Related references

[Delete snapshot image](#) on page 186

[Resume snapshot image rollback](#) on page 258

[Show snapshot image](#) on page 401

[Start snapshot image rollback](#) on page 456

[Stop snapshot image rollback](#) on page 473

## 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"> <li>• The name of the snapshot group</li> <li>• An identifier for the snapshot image in the snapshot group</li> </ul> <p>The identifier for the snapshot image can be one of these:</p> <ul style="list-style-type: none"> <li>• An integer value that is the sequence number of the snapshot in the snapshot group.</li> <li>• <b>NEWEST</b> - Use this option when you want to show the latest snapshot image created in the snapshot group.</li> <li>• <b>OLDEST</b> - Use this option when you want to show the earliest snapshot image created in the snapshot group.</li> </ul> <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"> <li>• Use an existing repository volume name</li> <li>• Create a new repository volume when you run this command</li> </ul> <p>The name of an existing repository volume is comprised of two parts:</p> <ul style="list-style-type: none"> <li>• The term repos</li> <li>• A four digit numerical identifier that the storage management software assigns to the repository volume name</li> </ul> <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"> <li>• An integer value that represents a percentage of the base volume capacity</li> <li>• A decimal fraction value that represents a percentage of the base volume capacity</li> <li>• A specific size for the repository volume. Size is defined in units of bytes , KB , MB , GB , or TB .</li> </ul> <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

### Related references

- [Create read-only snapshot volume](#) on page 145
- [Delete snapshot volume](#) on page 187
- [Rename snapshot volume](#) on page 242
- [Resume snapshot volume](#) on page 259
- [Revive snapshot volume](#) on page 263
- [Set read-only snapshot volume to read/write volume](#) on page 328
- [Set snapshot volume media scan](#) on page 339
- [Set snapshot volume repository volume capacity](#) on page 340
- [Set SSD cache for a volume](#) on page 342
- [Show snapshot volumes](#) on page 403
- [Stop snapshot volume](#) on page 474

## 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)]
[securityType=(none|capable|enabled)]
```

### Parameters

Parameter	Description
<code>userLabel</code>	The name that you want to name the new SSD cache. Enclose the SSD cache name in double quotation marks (" ").

Parameter	Description
<code>drives</code>	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 24 . Enclose the tray ID values, the drawer ID values, and the slot ID values in parentheses.
<code>updateExistingVolumes</code>	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 <code>TRUE</code> . If you want to enable the SSD cache for individual volumes later, set this parameter to <code>FALSE</code> . The default value is <code>TRUE</code> .
<code>securityType</code>	This optional parameter specifies the security level when creating an SSD cache. These settings are valid: <ul style="list-style-type: none"> <li>• <code>none</code> - The SSD cache is not secure.</li> <li>• <code>capable</code> - The SSD cache is capable of having security set, but security has not been enabled.</li> <li>• <code>enabled</code> - The SSD cache has security enabled.</li> </ul>

### Notes

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

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 images.

If all of the SSDs in the SSD cache are data assurance (DA)-capable and the DA 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

8.20\_M3 - FDE drives are now supported. Added the `securityType` parameter.

### Related references

[Add drives to SSD cache](#) on page 77

[Change SSD cache application type](#) on page 90

[Delete SSD cache](#) on page 188

[Remove drives from SSD cache](#) on page 235

[Rename SSD cache](#) on page 242

[Resume SSD cache](#) on page 260

[Set SSD cache for a volume](#) on page 342

[Show SSD cache](#) on page 406

[Show SSD cache statistics](#) on page 407



[Start SSD cache locate](#) on page 457

[Start SSD cache performance modeling](#) on page 458

[Stop SSD cache locate](#) on page 475

[Stop SSD cache performance modeling](#) on page 476

[Suspend SSD cache](#) on page 482

## 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:  <code>file="C:\Program Files\CLI\sup\seckey.slk"</code>  <b>IMPORTANT</b> – If you are creating the security key for storage array running the storage management software GUI, you must add a file extension of <code>.slk</code> 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 (" ").
commitSecurityKey	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:
  - 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`.

**Note:** Starting with firmware 8.10, the `writeOrder` parameter and the `writeMode` parameter cannot be used on storage arrays that have either an E2700 controller or an E5500 controller. The parameters are still valid for storage arrays that have an E2600 controller or an E5400 controller.

## 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)]
[role=(primary | secondary)]
[force=(TRUE | FALSE)]
```

## Parameters

Parameter	Description
primary	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 (" ").
secondary	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 (" ").
remoteStorageArrayName	The name of the remote storage array. Enclose the remote storage array name in double quotation marks (" ").
remoteStorageArrayWwn	The World Wide Identifier (WWID) of the remote storage array. Enclose the WWID in double quotation marks (" ").
remotePassword	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 (" ").
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> .
autoResync	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"> <li><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.</li> <li><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.</li> </ul>
writeOrder	The write order for data transmission between the primary volume and the secondary volume. Valid values are <code>preserved</code> or <code>notPreserved</code> .
writeMode	How the primary volume writes to the secondary volume. Valid values are <code>synchronous</code> or <code>asynchronous</code> .

Parameter	Description
role	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> .
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> .

### 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 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.

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

8.10 the `writeOrder` parameter and the `writeMode` parameter cannot be used on storage arrays that have either an E2700 controller or an E5500 controller.

### Related references

[Re-create synchronous mirroring repository volume](#) on page 230

[Remove synchronous mirroring](#) on page 238

[Resume synchronous mirroring](#) on page 261

[Set synchronous mirroring](#) on page 360

[Show synchronous mirroring volume candidates](#) on page 426

## Create volume copy

This command creates a volume copy and starts the volume copy operation. This command is valid for 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 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 image, also called *offline* volume copy
- Volume copy with snapshot image, also called *online* volume copy

If you use volume copy without 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 image. You can select volume copy with snapshot image through the optional parameters in the command syntax.

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.

Parameter	Description
<code>copyType</code>	<p>Use this parameter to create a volume copy with a snapshot image. Creating a volume copy with a 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 image, set this parameter to <code>online</code>. To create a volume copy without a snapshot image, set this parameter to <code>offline</code>.</p> <p>If you do not use this parameter, the volume copy is created without a snapshot image.</p>
<code>repositoryPercentOfBase</code>	<p>This parameter determines the size of the repository volume for the snapshot image when you are creating a volume copy with a 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 <code>20</code>, <code>40</code>, <code>60</code>, <code>120</code>, and <code>default</code>. The default value is <code>20</code>. 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>
<code>repositoryGroupPreference</code>	<p>This parameter determines to which volume group the snapshot image repository volume is written. You have these choices:</p> <ul style="list-style-type: none"> <li><code>sameAsSource</code> – The snapshot image volume is written to the same volume group as the source volume if space is available.</li> <li><code>otherThanSource</code> – Deprecated parameter. Do not use</li> <li><code>default</code> – The snapshot image repository volume is written to any volume group that has space.</li> </ul> <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.

### Related references

[Set volume copy](#) on page 377

[Show volume copy](#) on page 431

## 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=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 24. 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 name that you want to give the new volume group. Enclose the volume group name in double quotation marks (" ").
driveMediaType	The type of drive media that you want to use for the volume group. You must use this parameter when you have more than one type of drive media in your storage array. These drive media are valid: <ul style="list-style-type: none"> <li>HDD – Use this option when you have hard drives in the drive tray.</li> <li>SSD – Use this option when you have solid state drives in the drive tray.</li> <li>unknown – Use this option if you are not sure what types of drive media are in the drive tray.</li> <li>allMedia – Use this option when you want to use all types of drive media that are in the drive tray.</li> </ul>

Parameter	Description
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>These drive types are valid:</p> <ul style="list-style-type: none"> <li>• fibre</li> <li>• SATA</li> <li>• SAS</li> </ul> <p>If you do not specify a drive type, the command defaults to SAS.</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"> <li>• none – The volume group and volumes are not secure.</li> <li>• capable – The volume group and volumes are capable of having security set, but security has not been enabled.</li> <li>• enabled – The volume group and volumes have security enabled.</li> </ul>
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"> <li>• none – The volume group does not have data assurance protection.</li> <li>• enabled – The volume group has data assurance protection. The volume group supports protected information and is formatted with protection information enabled.</li> </ul> <p><b>Note:</b> This option is only valid for enablement if the drives support DA.</p>

### Drives and volume groups

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.
- You cannot mix HDD and SSD drives 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, 3, 5, and 6 .
  - 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.
- To enable tray/drawer loss protection, refer to the following tables for additional criteria:

Level	Criteria for Tray Loss Protection	Minimum number of trays required
Disk Pool	The disk pool contains no more than two drives in a single tray.	6
RAID 6	The volume group contains no more than two drives in a single tray.	3
RAID 3 or RAID 5	Each drive in the volume group is located in a separate tray.	3
RAID 1	Each drive in a RAID 1 pair must be located in a separate tray.	2
RAID 0	Cannot achieve Tray Loss Protection.	Not applicable

Level	Criteria for Drawer Loss Protection	Minimum number of drawers required
Disk Pool	The disk pool includes drives from all five drawers and there are an equal number of drives in each drawer. A 60-drive tray can achieve Drawer Loss Protection when the disk pool contains 15, 20, 25, 30, 35, 40, 45, 50, 55, or 60 drives.	5
RAID 6	The volume group contains no more than two drives in a single drawer.	3
RAID 3 or RAID 5	Each drive in the volume group is located in a separate drawer.	3
RAID 1	Each drive in a mirrored pair must be located in a separate drawer.	2
RAID 0	Cannot achieve Drawer Loss Protection.	Not applicable

## Hot spares

With volume groups, a valuable strategy to protect data is to assign available drives in the storage array as hot spare drives. A hot spare is a drive, containing no data, that acts as a standby in the storage array in case a drive fails in a RAID 1, RAID 3, RAID 5, or RAID 6 volume group. The hot spare adds another level of redundancy to the storage array.

Generally, hot spare drives must have capacities that are equal to or greater than the used capacity on the drives that they are protecting. Hot spare drives must be of the same media type, the same interface type, and the same capacity as the drives that they are protecting.

If a drive fails in the storage array, the hot spare is normally substituted automatically for the failed drive without requiring your intervention. If a hot spare is available when a drive fails, the controller uses redundancy data parity to reconstruct the data onto the hot spare. Data evacuation support also allows data to be copied to a hot spare before the software marks the drive "failed."

After the failed drive is physically replaced, you can use either of the following options to restore the data:

When you have replaced the failed drive, the data from the hot spare is copied back to the replacement drive. This action is called copyback.

If you designate the hot spare drive as a permanent member of a volume group, the copyback operation is not needed.

The availability of tray loss protection and drawer loss protection for a volume group depends on the location of the drives that comprise the volume group. Tray loss protection and drawer loss protection might be lost because of a failed drive and the location of the hot spare drive. To make sure that tray loss protection and drawer loss protection are not affected, you must replace a failed drive to initiate the copyback process.

The storage array automatically selects Data Assurance (DA)-capable drives for hot spare coverage of DA-enabled volumes.

Make sure you have DA-capable drives in the storage array for hot spare coverage of DA-enabled volumes. For more information about DA-capable drives, refer to Data Assurance feature.

Secure-capable (FIPS and FDE) drives can be used as a hot spare for both secure-capable and non-secure-capable drives. Non-secure-capable drives can provide coverage for other non-secure-capable drives, and for secure-capable drives if the volume group does not have the security enabled. A FIPS volume group can only use a FIPS drive as a hot spare; however, you can use a FIPS hot spare for non-secure-capable, secure-capable, and secure-enabled volume groups.

If you do not have a hot spare, you can still replace a failed drive while the storage array is operating. If the drive is part of a RAID 1, RAID 3, RAID 5, or RAID 6 volume group, the controller uses redundancy data parity to automatically reconstruct the data onto the replacement drive. This action is called reconstruction.

## 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.

### Usage Hint

**Note:** 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. The segment size and cache read prefetch settings for various usage hints are shown in the following table:

Usage hint	Segment size setting	Dynamic cache read prefetch setting
File system	128 KB	Enabled
Database	128 KB	Enabled
Multimedia	256 KB	Enabled

### 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. Valid values for the `cacheReadPrefetch` parameter are `TRUE` or `FALSE`. The default is `TRUE`.

### Security type

Use the `securityType` parameter to specify the security settings for the storage array.

Before you can set the `securityType` parameter to `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`
- `export storageArray securityKey`
- `import storageArray securityKey`
- `set storageArray securityKey`
- `enable volumeGroup [volumeGroupName] security`
- `enable diskPool [diskPoolName] security`

### Secure drives

Secure-capable drives can be either Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives. Use the `secureDrives` parameter to specify the type of secure drives to use. The values you can use are `fips` and `fde`.

## Data assurance management

The Data Assurance (DA) feature increases data integrity across the entire storage system. DA enables the storage array to check for errors that might occur when data is moved between the hosts and the drives. When this feature is enabled, the storage array appends error-checking codes (also known as cyclic redundancy checks or CRCs) to each block of data in the volume. After a data block is moved, the storage array uses these CRC codes to determine if any errors occurred during transmission. Potentially corrupted data is neither written to disk nor returned to the host.

If you want to use the DA feature, start with a volume group or disk pool that includes only drives that support DA. Then, create DA-capable volumes. Finally, map these DA-capable volumes to the host using an I/O interface that is capable of DA. I/O interfaces that are capable of DA include Fibre Channel, SAS, and iSER InfiniBand (iSCSI Extensions for RDMA/IB). DA is not supported by iSCSI over TCP/IP, or by the SRP InfiniBand.

**Note:** When all the drives are DA-capable, you can set the `dataAssurance` parameter to `enabled` and then use DA with certain operations. For example, you can create a volume group that includes DA-capable drives, and then create a volume within that volume group that is DA-enabled. Other operations that use a DA-enabled volume have options to support the DA feature.

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.

## 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.

## Related references

[Delete volume group](#) on page 191

[Enable volume group security](#) on page 216

[Set volume group](#) on page 378

[Set volume group forced state](#) on page 380

[Show volume group](#) on page 433

## Create volume in disk pool

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"> <li>capacity=500MB</li> <li>capacity=2GB</li> </ul>
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.
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 none. If you want to map now, set this parameter to default. The volume is mapped to all hosts that have access to the storage pool.</p> <p>The default value is none.</p>

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"> <li>• none – The volume group does not have any explicit data assurance defined.</li> <li>• enabled – The volume group has data assurance protection. The volume group supports protected information and is formatted with protection information enabled.</li> </ul> <p>The default value is enabled if the disk pool is data assurance capable.</p> <p><b>Note:</b> This option is only available for enablement if the drives support DA.</p>
existingRepositoryLabel	<p>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.</p>
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"> <li>• <code>capacity=500MB</code></li> <li>• <code>capacity=2GB</code></li> </ul> <p>The default value is 50 percent of the virtual capacity.</p>
repositoryMaxCapacity	<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"> <li>• <code>capacity=500MB</code></li> <li>• <code>capacity=2GB</code></li> </ul>

Parameter	Description
<code>warningThresholdPercent</code>	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.
<code>repositoryExpansionPolicy</code>	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.
<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> .

### 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` parameter specifies the virtual capacity of the volume, and the `repositoryCapacity` parameter specifies the capacity of the volume created as the repository volume. Use the `existingRepositoryLabel` parameter 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

The Data Assurance (DA) feature increases data integrity across the entire storage system. DA enables the storage array to check for errors that might occur when data is moved between the hosts and the drives. When this feature is enabled, the storage array appends error-checking codes (also known as cyclic redundancy checks or CRCs) to each block of data in the volume. After a data block is moved, the storage array uses these CRC codes to determine if any errors occurred during transmission. Potentially corrupted data is neither written to disk nor returned to the host.

If you want to use the DA feature, start with a volume group or disk pool that includes only drives that support DA. Then, create DA-capable volumes. Finally, map these DA-capable volumes to the host using an I/O interface that is capable of DA. I/O interfaces that are capable of DA include Fibre Channel, SAS, and iSER InfiniBand (iSCSI Extensions for RDMA/IB). DA is not supported by iSCSI over TCP/IP, or by the SRP InfiniBand.

**Note:** When all drives are DA-capable, you can set the `dataAssurance` parameter to `enabled` and then use DA with certain operations. For example, you can create a volume group that includes DA-capable drives, and then create a volume within that volume group that is DA-enabled. Other operations that use a DA-enabled volume have options to support the DA feature.

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 data assurance drives.

#### Minimum firmware level

7.83

#### Related references

[Delete volume from disk pool](#) on page 190

[Increase capacity of volume in disk pool or volume group](#) on page 219

## Deactivate asynchronous mirroring

This command deactivates the Asynchronous Mirroring 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

#### Related references

[Activate asynchronous mirroring](#) on page 72

## Deactivate synchronous mirroring

This command deactivates the synchronous mirroring 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



**Related references**

[Activate Synchronous Mirroring](#) on page 74

**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 (allAsyncMirrorGroups | asyncMirrorGroup["asyncMirrorGroupName"]
|
asyncMirrorGroups ["asyncMirrorGroupName_1" ...
"asyncMirrorGroupName_n"])
```

**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 double quotation marks ( " ") inside square brackets ( [ ] ).
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"> <li>• Enclose all of the names in square brackets ( [ ] ).</li> <li>• Enclose each of the names in double quotation marks ( " " ).</li> <li>• Separate each of the names with a space.</li> </ul>

**Minimum firmware level**

7.84

**Related references**

[Add volume to asynchronous mirror group](#) on page 80

[Cancel asynchronous mirror group role reversal](#) on page 88

[Check asynchronous mirror group consistency](#) on page 91

[Check storage array connectivity](#) on page 94

[Clear asynchronous mirroring fault](#) on page 96

[Create asynchronous mirror group](#) on page 104

[Deactivate asynchronous mirroring](#) on page 176

[Establish asynchronous mirrored pair](#) on page 216

[Remove incomplete asynchronous mirrored pair from asynchronous mirror group](#) on page 236

[Remove volume from asynchronous mirror group](#) on page 239

- [Reset asynchronous mirror group statistics](#) on page 246
- [Reset iSCSI IP address](#) on page 250
- [Resume asynchronous mirror group](#) on page 255
- [Save asynchronous mirror group statistics](#) on page 265
- [Set asynchronous mirror group](#) on page 294
- [Show asynchronous mirror groups](#) on page 382
- [Show asynchronous mirror group synchronization progress](#) on page 383
- [Start asynchronous mirroring synchronization](#) on page 440
- [Suspend asynchronous mirror group](#) on page 481
- [Test asynchronous mirror group connectivity](#) on page 484

## 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 double quotation marks ( " ") inside square brackets ( [ ] ).
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

### Related references

- [Add member to consistency group](#) on page 78
- [Create consistency group](#) on page 107
- [Create consistency group snapshot image](#) on page 115
- [Create consistency group snapshot volume](#) on page 116
- [Create consistency group snapshot volume mapping](#) on page 120
- [Delete consistency group snapshot image](#) on page 179
- [Delete consistency group snapshot volume](#) on page 180
- [Remove member volume from consistency group](#) on page 237
- [Resume consistency group snapshot volume](#) on page 256
- [Set consistency group attributes](#) on page 296

[Set consistency group snapshot volume](#) on page 298

[Show consistency group](#) on page 384

[Show consistency group snapshot image](#) on page 386

[Start consistency group snapshot rollback](#) on page 443

[Stop consistency group snapshot rollback](#) on page 466

[Stop consistency group snapshot volume](#) on page 468

[Stop pending snapshot images on consistency group](#) on page 471

## 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.
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

**Related references**

[Add member to consistency group](#) on page 78  
[Create consistency group](#) on page 107  
[Create consistency group snapshot image](#) on page 115  
[Create consistency group snapshot volume](#) on page 116  
[Create consistency group snapshot volume mapping](#) on page 120  
[Delete consistency group](#) on page 178  
[Delete consistency group snapshot volume](#) on page 180  
[Remove member volume from consistency group](#) on page 237  
[Resume consistency group snapshot volume](#) on page 256  
[Set consistency group attributes](#) on page 296  
[Set consistency group snapshot volume](#) on page 298  
[Show consistency group](#) on page 384  
[Show consistency group snapshot image](#) on page 386  
[Start consistency group snapshot rollback](#) on page 443  
[Stop consistency group snapshot rollback](#) on page 466  
[Stop consistency group snapshot volume](#) on page 468  
[Stop pending snapshot images on consistency group](#) on page 471

**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

**Related references**

[Add member to consistency group](#) on page 78

[Create consistency group](#) on page 107  
[Create consistency group snapshot image](#) on page 115  
[Create consistency group snapshot volume](#) on page 116  
[Create consistency group snapshot volume mapping](#) on page 120  
[Delete consistency group](#) on page 178  
[Delete consistency group snapshot image](#) on page 179  
[Remove member volume from consistency group](#) on page 237  
[Resume consistency group snapshot volume](#) on page 256  
[Set consistency group attributes](#) on page 296  
[Set consistency group snapshot volume](#) on page 298  
[Show consistency group](#) on page 384  
[Show consistency group snapshot image](#) on page 386  
[Start consistency group snapshot rollback](#) on page 443  
[Stop consistency group snapshot rollback](#) on page 466  
[Stop consistency group snapshot volume](#) on page 468  
[Stop pending snapshot images on consistency group](#) on page 471

## 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)]
```

### Parameters

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 contains special characters or consists only of numbers, you must enclose the name in double quotation marks (" ") inside square brackets.
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

**Related references**[Create disk pool](#) on page 121**Delete host**

This command deletes one or more hosts.

**Syntax**

```
delete host [hostName]
```

**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 must enclose the host name in double quotation marks (" ") inside square brackets.

**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

**Related references**[Create host](#) on page 127[Create host group](#) on page 128[Create host port](#) on page 129[Delete host group](#) on page 182[Delete host port](#) on page 183[Set host](#) on page 320[Set host channel](#) on page 321[Set host group](#) on page 322[Set host port label](#) on page 323[Show host ports](#) on page 398[Show storage array host topology](#) on page 421**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 contains special characters or consists only of numbers, you must enclose the host group name in double quotation marks (" ") inside square brackets.

## 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

## Related references

- [Create host](#) on page 127
- [Create host group](#) on page 128
- [Create host port](#) on page 129
- [Delete host](#) on page 182
- [Delete host port](#) on page 183
- [Set host](#) on page 320
- [Set host channel](#) on page 321
- [Set host group](#) on page 322
- [Set host port label](#) on page 323
- [Show host ports](#) on page 398
- [Show storage array host topology](#) on page 421

## 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 ([ ]). If the host port name contains special characters or consists only of numbers, you also must enclose the name in double quotation marks (" ") inside 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

**Related references**

- [Create host](#) on page 127
- [Create host group](#) on page 128
- [Create host port](#) on page 129
- [Delete host](#) on page 182
- [Delete host group](#) on page 182
- [Set host](#) on page 320
- [Set host channel](#) on page 321
- [Set host group](#) on page 322
- [Set host port label](#) on page 323
- [Show host ports](#) on page 398
- [Show storage array host topology](#) on page 421

**Delete iSCSI initiator**

This command deletes a specific iSCSI initiator object.

**Syntax**

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

**Parameters**

Parameter	Description
iscsiInitiator	<p>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.</p> <p>Enclose an iSCSI ID in double quotation marks ( " ") inside angle brackets (&lt;&gt;).</p> <p>Enclose the name in double quotation marks ( " ") inside square brackets ([ ]).</p>

**Minimum firmware level**

7.10

**Related references**

- [Activate iSCSI initiator](#) on page 72
- [Create iSCSI initiator](#) on page 131
- [Diagnose controller iSCSI host cable](#) on page 193
- [Reset storage array iSCSI baseline](#) on page 252
- [Save storage array iSCSI statistics](#) on page 279
- [Set iSCSI initiator](#) on page 324



[Set iSCSI target properties](#) on page 326

[Show current iSCSI sessions](#) on page 390

[Show storage array negotiation defaults](#) on page 423

[Show storage array unconfigured iSCSI initiators](#) on page 424

[Start iSCSI DHCP refresh](#) on page 453

[Stop storage array iSCSI session](#) on page 479

## 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

### Related references

[Create snapshot group](#) on page 146

[Revive snapshot group](#) on page 263

[Set snapshot group attributes](#) on page 331

[Set snapshot group media scan](#) on page 333

[Set snapshot group repository volume capacity](#) on page 334

[Set snapshot group schedule](#) on page 336

[Set SSD cache for a volume](#) on page 342

[Show snapshot group](#) on page 399

[Stop snapshot group pending snapshot images](#) on page 472

## 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
snapGroup	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 (" ").
snapGroups	The names of several snapshot groups that has the snapshot image that you want to delete. Enter the names of the snapshot groups using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in parentheses.</li> <li>• Enclose each of the names in double quotation marks (" ").</li> <li>• Separate each of the names with a space.</li> </ul>
deleteCount	The number of snapshot images that you want to delete from the snapshot 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. 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.
retainCount	The number of snapshot images that you want to keep in the snapshot group. Use integer values. This parameter keeps the most recent snapshot images in the snapshot group and deletes older snapshot images. 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.

Parameter	Description
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"> <li>• TRUE - Use this value to delete the snapshot image even though the snapshot image has an associated snapshot volume.</li> <li>• FALSE - Use this value to keep the snapshot image when the snapshot image has an associated snapshot volume.</li> </ul> <p>The default value is TRUE .</p>
snapImageID	The snapImageID parameter accepts only the OLDEST option. This parameter deletes the earliest snapshot image created.

### 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.

If you do not use any other parameters with the snapGroup parameter or snapGroups parameter, by default the oldest snapshot image is deleted.

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

### Related references

[Create snapshot image](#) on page 155

[Resume snapshot image rollback](#) on page 258

[Show snapshot image](#) on page 401

[Start snapshot image rollback](#) on page 456

[Stop snapshot image rollback](#) on page 473

## 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 name 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 <code>FALSE</code> . To delete the repository members, set this parameter to <code>TRUE</code> . The default value is <code>TRUE</code> . If you do not use this parameter the repository members are automatically deleted.

**Minimum firmware level**

7.83

**Related references**

- [Create read-only snapshot volume](#) on page 145
- [Create snapshot volume](#) on page 156
- [Rename snapshot volume](#) on page 242
- [Resume snapshot volume](#) on page 259
- [Revive snapshot volume](#) on page 263
- [Set read-only snapshot volume to read/write volume](#) on page 328
- [Set snapshot volume media scan](#) on page 339
- [Set snapshot volume repository volume capacity](#) on page 340
- [Set SSD cache for a volume](#) on page 342
- [Show snapshot volumes](#) on page 403
- [Stop snapshot volume](#) on page 474

**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 name of the SSD cache that you want to delete. Enclose the SSD cache name in square brackets ( [ ] ). If the SSD cache name contains special characters or consists only of numbers, you must enclose the SSD cache name in double quotation marks ( " ") inside square brackets.

**Minimum firmware level**

7.84

**Related references**

- [Add drives to SSD cache](#) on page 77

[Change SSD cache application type](#) on page 90  
[Create SSD cache](#) on page 159  
[Remove drives from SSD cache](#) on page 235  
[Rename SSD cache](#) on page 242  
[Resume SSD cache](#) on page 260  
[Set SSD cache for a volume](#) on page 342  
[Show SSD cache](#) on page 406  
[Show SSD cache statistics](#) on page 407  
[Start SSD cache locate](#) on page 457  
[Start SSD cache performance modeling](#) on page 458  
[Stop SSD cache locate](#) on page 475  
[Stop SSD cache performance modeling](#) on page 476  
[Suspend SSD cache](#) on page 482

## Delete volume

This command deletes one or more standard 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	The name of a volume that you want to delete. Enclose the volume name in square brackets ([ ]). If the volume name has special characters or consists only of numbers, you must enclose the volume name in double quotation marks (" ") inside square brackets.
volumes	<p>The names of several volumes that you want to delete. Enter the names of the volumes using these rules:</p> <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Separate each of the names with a space.</li> </ul> <p>If the volume names have special characters or consist only of numbers, enter the names using these rules:</p> <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Enclose each of the names in double quotation marks (" ").</li> <li>• Separate each of the names with a space.</li> </ul>

Parameter	Description
<code>removeVolumeGroup</code>	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 <code>TRUE</code> . To keep standalone volume groups intact, set this parameter to <code>FALSE</code> .
<code>force</code>	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 <code>TRUE</code> . To wait until the controllers have finished performing other operations, do not use this parameter or set this parameter to <code>FALSE</code> .

### 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.

### Related references

[Create RAID volume \(automatic drive select\)](#) on page 132

[Create RAID volume \(free extent based select\)](#) on page 137

[Create RAID volume \(manual drive select\)](#) on page 140

[Create volume in disk pool](#) on page 172

[Delete volume from disk pool](#) on page 190

## 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"]
[force=(TRUE | FALSE)]
[schedule]
[retainRepositoryMembers=(TRUE | FALSE)])
```

## Parameters

Parameter	Description
<code>allVolumes</code>	This parameter deletes all of the volumes on a disk pool.
<code>volume</code>	The name of a specific volume that you want to delete. Enclose the volume name in square brackets ([ ]). If the volume name contains special characters or consists only of numbers, you must enclose the volume name in double quotation marks (" ") inside square brackets.
<code>volumes</code>	The names of several volumes that you want to delete. Enter the names of the volumes using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Separate each of the names with a space.</li> </ul>
<code>force</code>	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 <code>TRUE</code> . To wait until the controllers have finished performing other operations, do not use this parameter or set this parameter to <code>FALSE</code> .
<code>schedule</code>	This parameter deletes any schedule for related to a specific disk pool volume. Only the schedule is deleted, the disk pool volume remains.
<code>retainRepositoryMembers</code>	When you delete a thin volume, the associated repository volume is deleted by default. However, when the <code>retainRepositoryMembers</code> is set to <code>TRUE</code> , the repository volume is retained. For normal volumes, this parameter has no effect.

## 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

## Related references

[Create volume in disk pool](#) on page 172

[Delete volume](#) on page 189

[Increase capacity of volume in disk pool or volume group](#) on page 219

## 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 name of the volume group that you want to delete. Enclose the volume group name in square brackets ([ ]).

## Minimum firmware level

6.10

## Related references

[Create volume group](#) on page 167

## 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 .
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"> <li>• 1 – Read test</li> <li>• 2 – Write test</li> <li>• 3 – Data loop-back test</li> <li>• discreteLines – Discrete lines diagnostic test</li> </ul>



Parameter	Description
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. 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 controller iSCSI host cable

This command runs diagnostic tests on the copper cables between iSCSI host interface cards and a controller. You can run diagnostics on a selected port or all ports. The ports must be able to support the cable diagnostics. If the ports do not support cable diagnostics, an error is returned.

### Syntax

```
diagnose controller [(a|b)]
iscsiHostPort (all | [portLabel])
testID=cableDiagnostics
```

## Parameters

Parameter	Description
controller	The controller on which you want to run the cable diagnostic test. 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.
iscsiHostPort	The iSCSI host port on which you want to run the diagnostic tests. You can run the diagnostics on all iSCSI host ports, or you can run the diagnostics on a specific iSCSI host port. See the following for more information: <a href="#">"Identifying an iSCSI host port label"</a>
testID	The identifier for the diagnostic test that you want to run. For this diagnostic test, the only choice is <code>cableDiagnostics</code> .

## Identifying an iSCSI host port label

You must specify a label for the host port. Follow these steps to specify the host port label:

1. If you do not know the port label for the iSCSI host port, run the `show controller` command.
2. In the Host interface section of the results, locate the host port you want to select.  
**Note:** The port label is the complete value returned for the `Port` field.
3. Enclose the entire value of the port label in both quotes and square brackets: `["portLabel"]`. For example, if the port label is `Ch 2`, specify the iSCSI host port as follows:

```
iscsiHostPort["ch 2"]
```

**Note:** If you are using a Windows command line and the label contains a pipe (`|`), the character should be escaped (using `^`); otherwise, it will be interpreted as a command. For example, if the port label is `e0b|0b`, specify the iSCSI host port as follows:

```
iscsiHostPort["e0b^|0b"]
```

**Note:** For backward compatibility, the `iscsiPortNumber`, enclosed by braces [ ] rather than quotes and braces [" "] can still be used for E2700, E5600, or EF560 controllers (and other previous generations of E-Series or EF-Series controllers). For those controllers, valid values for `iscsiPortNumber` are as follows:

- For controllers with integrated host ports, the numbering is 3, 4, 5, or 6.
- For controllers with host ports on a host interface card only, the numbering is 1, 2, 3, or 4.

An example of the prior syntax is as follows:

```
iscsiHostPort[3]
```

## Notes

When you run the cable diagnostic test, the firmware returns the following information:

- Host Port: The port on which the diagnostic test was run.

- HIC: The host interface card associated with this port.
- The date and time the test was run.
- Status:
  - OK: All of the pairs of cables are good and do not have any faults.
  - Open: One or more of the four pairs of cables are open.
  - Short: One or more of the four pairs of cables are shorted.
  - Incomplete: One or more of the four pairs returned incomplete or invalid test results.
- Length – The length of the cables are listed in meters and the following information about the cables is returned:
  - When the cable status is OK, the approximate lengths of the cable pairs are returned. The lengths of the cable pairs are shown as a range (L1-L2), which are the shortest and the longest lengths of the cable pairs.
  - If the cable status is Open or Short, the approximate distance to the failure in the cable pairs. If there is one failure, the length is reported for that cable pair. If there is more than one failure, the information returned is both the shortest and longest lengths to the failures. The lengths are listed as a range (L1-L2) where L1<L2.
  - If the cable status is Incomplete, the information returned are the lengths for the shortest and longest cable pairs that the firmware can successfully test. The lengths are listed for the valid cable pairs as a range (L1-L2) where L1<L2.
- Register values for the cable diagnostic registers. The values are in a hexadecimal format:
  - Two bytes show the combined cable status (four bits per port).
  - Four two-byte numbers show the length of each channel.

### Minimum firmware level

7.77

8.10 revises the numbering system for iSCSI host ports.

### Related references

[Clear drive channel statistics](#) on page 98  
[Enable controller data transfer](#) on page 208  
[Reset controller](#) on page 248  
[Save controller NVSRAM](#) on page 267  
[Save drive channel fault isolation diagnostic status](#) on page 267  
[Set controller](#) on page 299  
[Set controller DNS settings](#) on page 304  
[Set controller NTP settings](#) on page 305  
[Set controller service action allowed indicator](#) on page 306  
[Set drive channel status](#) on page 313  
[Set host channel](#) on page 321  
[Show controller](#) on page 387  
[Show controller NVSRAM](#) on page 390  
[Show drive channel statistics](#) on page 395  
[Start controller trace](#) on page 444

[Start drive channel fault isolation diagnostics](#) on page 447  
[Start drive channel locate](#) on page 449  
[Start iSCSI DHCP refresh](#) on page 453  
[Stop drive channel fault isolation diagnostics](#) on page 469  
[Stop drive channel locate](#) on page 470  
[Activate iSCSI initiator](#) on page 72  
[Create iSCSI initiator](#) on page 131  
[Delete iSCSI initiator](#) on page 184  
[Reset storage array iSCSI baseline](#) on page 252  
[Save storage array iSCSI statistics](#) on page 279  
[Set iSCSI initiator](#) on page 324  
[Set iSCSI target properties](#) on page 326  
[Show current iSCSI sessions](#) on page 390  
[Show storage array negotiation defaults](#) on page 423  
[Show storage array unconfigured iSCSI initiators](#) on page 424  
[Start iSCSI DHCP refresh](#) on page 453  
[Stop storage array iSCSI session](#) on page 479

## 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 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
<code>primary</code>	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 or consists only of numbers, you must enclose the primary volume name in double quotation marks (" ") inside square brackets.
<code>primaries</code>	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 space. The names of several volumes that you want you want to use as primary volumes. Enter the names of the primary volumes using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Enclose each of the names in double quotation marks (" ").</li> <li>• Separate each of the names with a space.</li> </ul>

**Minimum firmware level**

6.10

**Related references**[Check synchronous mirroring status](#) on page 95[Set synchronous mirroring](#) on page 360**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.

**Restriction:** This command is obsolete for the E2800 storage array. Use the RESTful API, SANtricity System Manager, or see [Using curl commands with the E2800 storage array](#) on page 39.

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

**Related references**[Using curl commands with the E2800 storage array](#) on page 39[Capture or view an AutoSupport log](#) on page 89[Display AutoSupport bundle collection settings](#) on page 198[Display AutoSupport message collection schedule](#) on page 200[Enable or disable AutoSupport at the EMW management domain level \(SMcli only\)](#) on page 210[Enable or disable the AutoSupport OnDemand feature at the EMW management domain level \(SMcli only\)](#) on page 211[Enable or disable the AutoSupport OnDemand Remote Diagnostics feature at the EMW management domain level \(SMcli only\)](#) on page 213[Reset AutoSupport message collection schedule](#) on page 247[Save storage array support data](#) on page 283[Specify the AutoSupport delivery method](#) on page 438[Test the AutoSupport configuration](#) on page 485**Disable storage array feature**

This command disables a storage array feature. Run the `show storageArray` command to show a list of the feature identifiers for all enabled features in the storage array.

## Syntax

```
disable storageArray (featurePack | feature=featureAttributeList)
```

The *featureAttributeList* can be one or more of these attribute values. If you enter more than one attribute value, separate the values with a space.

- `driveSecurity`

## Parameters

None.

## Minimum firmware level

8.20 adds the `driveSecurity` attribute and removes all other attributes.

## Related references

[Enable storage array feature](#) on page 214

[Show storage array](#) on page 410

## Display AutoSupport bundle collection settings

The `SMcli -autoSupportConfig show` command displays the AutoSupport bundle collection settings for the management domain. These settings apply to all of the storage arrays in your management domain for which you have enabled AutoSupport.

The command output from this command shows the following configuration information:

- The activation status for this management domain of the AutoSupport feature, the AutoSupport OnDemand feature, and the AutoSupport Remote Diagnostics feature
- Delivery parameters that depend on the configured delivery method:
  - Delivery method is SMTP: Destination email address, mail relay server, and sender email address are shown
  - Delivery method is HTTP or HTTPS:
    - Direct connection - destination IP address is shown
    - Proxy server connection - Host address, port number, and authentication details are shown
    - Automatic proxy configuration script (PAC) connection - script location is shown
- A table that lists, for each storage array, the name of the array, whether or not it is an ASUP-capable array, the chassis serial numbers, and the IP addresses for the controllers in the array

## Context

**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.

**Restriction:** This command is obsolete for the E2800 storage array. Use the RESTful API, SANtricity System Manager, or see [Using curl commands with the E2800 storage array](#) on page 39.

## Syntax

```
SMcli -autoSupportConfig show
```

## Parameters

None.

## Minimum Firmware Level

8.25

## Example output for the activation statuses

```
The AutoSupport feature is activated on this storage management station.
The AutoSupport OnDemand feature is de-activated on this storage
management station.
The AutoSupport Remote Diagnostics feature is de-activated on this
storage management station.
```

## Example output for the SMTP delivery method

```
Delivery method: SMTP
Destination email address: http://asupserver.corp.netapp.com/put/
AsupPut/
Mail relay server: server.eng.netapp.com
Sender email address: user@netapp.com
```

## Example output for the HTTP or HTTPS delivery method with a direct connection

```
Delivery method: HTTP
Destination IP address: http://asupserver.corp.netapp.com/put/AsupPut/
Connection: Direct
```

## Example output for the HTTP delivery method with a proxy connection

```
Connection: Proxy server
Host Address: 10.227.76.123
Port number: 8080
Authentication required: Yes
User name: admin
Password: *****
```

## Example output for the HTTP delivery method with an automatic proxy configuration script (PAC)

```
Connection: Automatic proxy configuration script (PAC)
Script location: http://esgweb.eng.netapp.com/~user/proxy.pac
```

## Example output for array table (all delivery methods)

Name	AutoSupport Capable	Chassis Serial Number	IP Address
SA_1	Yes (enabled)	SX94500434	10.113.173.123,
			10.113.173.456
SA_2	Yes (disabled)	SX94607107	10.113.174.789,

10.113.174.345		
SA_3 No	Not available	10.113.59.58,
10.113.59.59		

### Related references

[Using curl commands with the E2800 storage array](#) on page 39

[Capture or view an AutoSupport log](#) on page 89

[Display AutoSupport message collection schedule](#) on page 200

[Disable AutoSupport at the EMW level SMcli version](#) on page 197

[Enable or disable AutoSupport at the EMW management domain level \(SMcli only\)](#) on page 210

[Enable or disable the AutoSupport OnDemand feature at the EMW management domain level \(SMcli only\)](#) on page 211

[Enable or disable the AutoSupport OnDemand Remote Diagnostics feature at the EMW management domain level \(SMcli only\)](#) on page 213

[Reset AutoSupport message collection schedule](#) on page 247

[Save storage array support data](#) on page 283

[Specify the AutoSupport delivery method](#) on page 438

[Test the AutoSupport configuration](#) on page 485

## 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 AutoSupport message collection schedule

The `SMcli -autoSupportSchedule show` command shows the schedule for collection and processing of both daily and weekly AutoSupport messages.

### Context

**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.

**Restriction:** This command is obsolete for the E2800 storage array. Use the RESTful API, SANtricity System Manager, or see [Using curl commands with the E2800 storage array](#) on page 39.

Note the following with regard to the scheduling of Automatic Support Bundle collection:



- If ASUP is enabled, the management software sends both daily AutoSupport messages and weekly AutoSupport messages.
- The management software randomly selects a time of day for both daily and weekly messages and a day of the week for weekly messages.
- The management software makes all attempts to ensure that no two storage arrays within a management domain send scheduled AutoSupport messages at the same time.
- Using the `SMcli -autoSupportSchedule reset` command, you can specify a range (hour-level granularity) of preferred times to send both daily and weekly messages for the storage arrays in your management domain. For the weekly messages schedule, you use the command to select preferred days of the week.

## Syntax

```
SMcli -autoSupportSchedule show
```

## Parameters

None.

## Minimum Firmware Level

8.25

## Example

The output shows both the preferred time ranges, followed by the daily and weekly schedule.

**Note:** If you do not select preferred time ranges for both schedules, and preferred days of the week, then this command shows the daily schedule and weekly schedule that the storage management software picked at random.

```
SMcli -autoSupportSchedule show
The AutoSupport daily schedule preferred time range is from 12:00 AM to
01:00 AM.
The AutoSupport weekly schedule preferred time range is from 10:00 PM to
11:00 PM on Thursday, Saturday.
```

Name	Daily Schedule	Weekly Schedule
Accounting	12:55 AM	10:08 PM on Thursdays
Finance	12:02 AM	10:30 PM on Saturdays

```
SMcli completed successfully.
```

## Related references

[Using curl commands with the E2800 storage array](#) on page 39

[Capture or view an AutoSupport log](#) on page 89

[Display AutoSupport bundle collection settings](#) on page 198

[Disable AutoSupport at the EMW level SMcli version](#) on page 197

[Enable or disable AutoSupport at the EMW management domain level \(SMcli only\)](#) on page 210

[Enable or disable the AutoSupport OnDemand feature at the EMW management domain level \(SMcli only\)](#) on page 211

[Enable or disable the AutoSupport OnDemand Remote Diagnostics feature at the EMW management domain level \(SMcli only\)](#) on page 213

[Reset AutoSupport message collection schedule](#) on page 247

[Save storage array support data](#) on page 283

[Specify the AutoSupport delivery method](#) on page 438

[Test the AutoSupport configuration](#) on page 485

## 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.

**Note:** The `content` parameter is deprecated. Use the `firmware` parameter instead.

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]|drives [trayID1, drawerID1,
slotID1.....trayIDN, drawerIDN, slotIDN] [online|offline] firmware
file="Filename"
```

### Parameters

Parameter	Description
drive or drives	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 24 . Enclose the tray ID values, the drawer ID values, and the slot ID values in brackets ([ ]).

Parameter	Description
online or offline	Select <code>online</code> to perform an online drive firmware download, otherwise select <code>offline</code> . The default value is <code>offline</code> .
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: <pre>file="C:\Program Files\CLI\dnlld\drvfrm.dlp"</pre> Valid file names have a <code>.dlp</code> extension.
firmware	Indicates to download controller firmware.

## Notes

Before trying to download drive firmware, take these precautions:

- Unless you are using the `online` parameter, stop all I/O activity to the storage array before you download the firmware image. The `download drive firmware` command without the `online` parameter 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.
- For online drive firmware update requests issued using this command, the controller will remove RAID 0 volume group drives from the list of drives to which new firmware is downloaded. Individual drive return status will be set to `Never Attempted With Reason`.

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

7.60 adds the `drawerID` user input.

8.25 adds the `online` parameter.

**Related references**

[Show drive download progress](#) on page 396

**Download environmental card firmware**

This command downloads environmental services module (ESM) firmware.

**Important:** For the E2800 storage array only - to ensure the downloads are not reverted, disable the IOM service on the E2800 controller before the download. When you are finished updating the firmware, re-enable the IOM service.

**Syntax**

```
download (allTrays | tray [trayID])
firmware file="filename"
```

**Parameters**

Parameter	Description
allTrays	This parameter downloads new firmware to all of the compatible trays in the storage array. The firmware package you select determines which trays are compatible. Incompatible trays are skipped. No error message is displayed about incompatible trays.
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 ([ ]).  <b>Note:</b> There can be no leading zero in the tray ID. For example, if the displayed tray ID is "02," this command requires it to be specified as [2], not [02].
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\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\sadrvm.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

**Related references**

[Activate storage array firmware](#) on page 73

[Download storage array firmware/NVSRAM](#) on page 205

[Download storage array NVSRAM](#) on page 207

**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)
healthCheckMelOverride=(TRUE | FALSE)]
```

**Parameters**

<b>Parameter</b>	<b>Description</b>
NVSRAM	This parameter downloads 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: <pre>file="C:\Program Files\CLI\dnld\safrm.dlp"</pre> Valid file names have a <code>.dlp</code> extension.
<i>NVSRAM-filename</i>	The file path and the file name that contains the NVSRAM values. Enclose the NVSRAM file name in double quotation marks (" "). For example: <pre>file="C:\Program Files\CLI\dnld\safrm.dlp"</pre> Valid file names have a <code>.dlp</code> extension. Include a comma before the file name when you download both firmware and NVSRAM.
downgrade	<b>Attention: Possible damage to the storage array configuration</b> – 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 technical support before using this parameter.  The setting to load firmware that is a previous version. The default value is <code>FALSE</code> . Set the <code>downgrade</code> parameter to <code>TRUE</code> if you want to download an earlier version of firmware.
activateNow	The setting to activate the firmware image and the NVSRAM image. The default value is <code>TRUE</code> . If you set the <code>activateNow</code> parameter to <code>FALSE</code> , you must run the <code>activate storageArray firmware</code> command to activate the firmware and the NVSRAM at a later time.
healthCheckMelOverride	The setting to override the results of the health check of the major event log (MEL). MEL validation still occurs, it is not bypassed. If the MEL check fails, you can bypass the failure by using this parameter when running the command.  Before any download, the controller checks the event log to determine if an event has occurred that might prevent the new controller firmware from successfully downloading. If such an event has occurred, the controller will normally not download new firmware.  This parameter forces the controller to download new firmware. The default value is <code>FALSE</code> . Set this value to <code>TRUE</code> if you want to force the controller to download new controller firmware.

**Minimum firmware level**

5.00

8.10 adds the `healthCheckMelOverride` parameter.

**Related references**

- [Activate storage array firmware](#) on page 73
- [Download storage array drive firmware](#) on page 204
- [Download storage array NVSRAM](#) on page 207

**Download storage array NVSRAM**

This command downloads the NVSRAM values for the storage array controller.

**Syntax**

```
download storageArray NVSRAM file="filename"
[healthCheckMelOverride=(TRUE | FALSE)]
```

**Parameters**

Parameter	Description
file	<p>The file path and the file name that contains the NVSRAM values. Enclose the NVSRAM file name in double quotation marks (" "). For example:</p> <pre>file="C:\Program Files\CLI\dnld\afrm.dlp"</pre> <p>Valid file names have a .dlp extension.</p>
healthCheckMelOverride	<p>The setting to override the results of the health check of the major event log (MEL). MEL validation still occurs, it is not bypassed. If the MEL check fails, you can bypass the failure by using this parameter when running the command.</p> <p>Before any download, the controller checks the event log to determine if an event has occurred that might prevent the new NVSRAM from successfully downloading. If such an event has occurred, the controller will normally not download new NVSRAM.</p> <p>This parameter forces the controller to download new NVSRAM. The default value is FALSE . Set this value to TRUE if you want to force the controller to download new NVSRAM.</p>

**Minimum firmware level**

6.10

8.10 adds the `healthCheckMelOverride` parameter.

**Related references**

- [Download storage array drive firmware](#) on page 204
- [Download storage array firmware/NVSRAM](#) on page 205

**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
allTrays	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 ([ ]).
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: <pre>file="C:\Program Files\CLI\dnld\trayset.dlp"</pre> Valid file names have a .dlp extension.

## 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 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

## Related references

[Reset controller](#) on page 248



## 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

### Related references

[Set disk pool](#) on page 307

[Set disk pool \(modify disk pool\)](#) on page 310

[Show disk pool](#) on page 392

## Enable or disable ODX

This command turns on or turns off Offloaded Data Transfer (ODX) for a storage array. The storage array comes with ODX turned on. Use this command to turn off ODX if you do not want to run ODX, or if ODX is causing issues with the storage array. The only storage array that is affected is the storage array on which the command is run.

### Syntax

```
set storageArray odxEnabled=(TRUE | FALSE)
```

### Parameter

Parameter	Description
odxEnabled	The setting to turn on or turn off ODX. To turn on ODX set this parameter to TRUE . To turn off ODX set this parameter to FALSE . The default setting is that ODX is turned on.

**Notes**

ODX provides a way to transfer data without using buffered read and buffered write operations and does not require direct host involvement with the data transfer operation. When ODX is not enabled, data is read from the source storage to the host, and then written to the target storage from the host. With ODX enabled, the data transfer operations are directly managed by the storage infrastructure. The data is moved directly from the source storage to the target storage without going through the host.

**Minimum firmware level**

8.20

**Related references**

[Enable or disable VAAI](#) on page 214

[Show storage array ODX setting](#) on page 423

## Enable or disable AutoSupport at the EMW management domain level (SMcli only)

The `SMcli enable autoSupportFeature` command turns on the AutoSupport (ASUP) feature for all managed storage arrays and makes it possible to transmit messages to the NetApp 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.

**Context**

**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.


**Note:** The Persistent Monitor must be enabled for the AutoSupport feature to work.

**Note:** No customer data is ever transmitted to the NetApp technical support site.

**Restriction:** This command is obsolete for the E2800 storage array. Use the RESTful API, SANtricity System Manager, or see [Using curl commands with the E2800 storage array](#) on page 39.

This command impacts all ASUP-capable storage arrays in the management domain. After enabling this feature for the management domain, you have control over the feature enablement for each storage array.

After enabling this feature, you can next enable the AutoSupport OnDemand feature (if desired), and then enable the AutoSupport Remote Diagnostics feature (if desired). You must enable the three features in this order.

**Syntax**

```
SMcli enable|disable autoSupportFeature
```

**Parameters**

None.

**Minimum Firmware Level**

7.86

**Verification**

Use the `SMcli -autoSupportConfig show` command to see whether you have enabled the feature. The initial line of the displayed output shows the enable status:

```
The AutoSupport feature is activated on this storage management station.
```

**Related references**

[Using curl commands with the E2800 storage array](#) on page 39

[Capture or view an AutoSupport log](#) on page 89

[Display AutoSupport bundle collection settings](#) on page 198

[Display AutoSupport message collection schedule](#) on page 200

[Disable AutoSupport at the EMW level SMcli version](#) on page 197

[Enable or disable the AutoSupport OnDemand feature at the EMW management domain level \(SMcli only\)](#) on page 211

[Enable or disable the AutoSupport OnDemand Remote Diagnostics feature at the EMW management domain level \(SMcli only\)](#) on page 213

[Reset AutoSupport message collection schedule](#) on page 247

[Save storage array support data](#) on page 283

[Specify the AutoSupport delivery method](#) on page 438

[Test the AutoSupport configuration](#) on page 485

**Enable or disable the AutoSupport OnDemand feature at the EMW management domain level (SMcli only)**

The `SMcli enable/disable autoSupportOnDemand` command turns on or turns off the AutoSupport OnDemand feature. This feature allows technical support to coordinate AutoSupport data transmission, and allows them to request the re-transmission of missing support data.

**Context**

**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.

**Note:** This command impacts all ASUP-capable storage arrays in the management domain.

**Restriction:** This command is obsolete for the E2800 storage array. Use the RESTful API, SANtricity System Manager, or see [Using curl commands with the E2800 storage array](#) on page 39.

Before enabling this feature, you must first enable the AutoSupport feature on the management domain. After enabling this feature, you can next enable the AutoSupport Remote Diagnostics feature (if desired). The AutoSupport Remote Diagnostics features is enabled, by default, when the AutoSupport OnDemand feature is enabled.



## Syntax

```
SMcli enable|disable autoSupportOnDemand
```

## Parameters

None.

## Minimum Firmware Level

8.25

## Verification

Use the `SMcli -autoSupportConfig show` command to see if you have enabled the feature. The initial two lines of the displayed output show the enable status of the AutoSupport feature, followed by the AutoSupport OnDemand feature:

```
The AutoSupport feature is activated on this storage management
station.
The AutoSupport OnDemand feature is activated on this storage
management station.
```

## Related references

[Using curl commands with the E2800 storage array](#) on page 39

[Capture or view an AutoSupport log](#) on page 89

[Display AutoSupport bundle collection settings](#) on page 198

[Display AutoSupport message collection schedule](#) on page 200

[Disable AutoSupport at the EMW level SMcli version](#) on page 197

[Enable or disable AutoSupport at the EMW management domain level \(SMcli only\)](#) on page 210

[Enable or disable the AutoSupport OnDemand Remote Diagnostics feature at the EMW management domain level \(SMcli only\)](#) on page 213

[Reset AutoSupport message collection schedule](#) on page 247

[Save storage array support data](#) on page 283

[Specify the AutoSupport delivery method](#) on page 438

[Test the AutoSupport configuration](#) on page 485

## Enable or disable the AutoSupport OnDemand Remote Diagnostics feature at the EMW management domain level (SMcli only)

The `SMcli enable/disable autoSupportRemoteDiag` command turns on or turns off the AutoSupport On Demand Remote Diagnostics feature. This feature enables technical support to request support data to diagnose problems remotely.


### Context

**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.

**Note:** This command impacts all ASUP-capable storage arrays in the management domain.

**Restriction:** This command is obsolete for the E2800 storage array. Use the RESTful API, SANtricity System Manager, or see [Using curl commands with the E2800 storage array](#) on page 39.

Before enabling this feature, first enable the AutoSupport feature and then the AutoSupport OnDemand feature on the management domain. You must enable the three features in this order.



### Syntax

```
SMcli enable|disable autoSupportRemoteDiag
```

### Parameters

None.

### Minimum Firmware Level

8.25

### Verification

Use the `SMcli -autoSupportConfig show` command to see whether you have enabled the feature. The initial three lines of the displayed output show the enable status of the AutoSupport feature, followed by the AutoSupport OnDemand feature, followed by the status of the AutoSupport Remote Diagnostics feature:

```
The AutoSupport feature is activated on this storage management
station.
The AutoSupport OnDemand feature is activated on this storage management
station.
The AutoSupport Remote Diagnostics feature is activated on this storage
management station.
```

### Related references

[Using curl commands with the E2800 storage array](#) on page 39

[Capture or view an AutoSupport log](#) on page 89

- [Display AutoSupport bundle collection settings](#) on page 198
- [Display AutoSupport message collection schedule](#) on page 200
- [Disable AutoSupport at the EMW level SMcli version](#) on page 197
- [Enable or disable AutoSupport at the EMW management domain level \(SMcli only\)](#) on page 210
- [Enable or disable the AutoSupport OnDemand feature at the EMW management domain level \(SMcli only\)](#) on page 211
- [Reset AutoSupport message collection schedule](#) on page 247
- [Save storage array support data](#) on page 283
- [Specify the AutoSupport delivery method](#) on page 438
- [Test the AutoSupport configuration](#) on page 485

## Enable or disable VAAI

This command turns on or turns off VMware vStorage API Array Architecture (VAAI) for a storage array. The storage array comes with VAAI turned on. Use this command to turn off VAAI if you do not want to run VAAI, or if VAAI is causing issues with the storage array. The only storage array that is affected is the storage array on which the command is run.

### Syntax

```
set storageArray vaaiEnabled=(TRUE | FALSE)
```

### Parameter

Parameter	Description
vaaiEnabled	The setting to turn on or turn off VAAI. To turn on VAAI set this parameter to <code>TRUE</code> . To turn off VAAI set this parameter to <code>FALSE</code> . The default setting is that VAAI is turned on.

### Notes

**Note:** VAAI enables some storage tasks to be offloaded from the server hardware to a storage array. Offloading the storage tasks to a storage array increases host performance by reducing the tasks the host must perform.

**Note:** View the VAAI setting using the `show storageArray odxSetting name`.

### Minimum firmware level

8.20

### Related references

- [Enable or disable ODX](#) on page 209
- [Show storage array ODX setting](#) on page 423

## Enable storage array feature

This command enables a 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 features, such as Storage Partitioning and Synchronous Mirroring. These features are combined for the convenience of the users. When users install a feature pack, all of the features in the feature pack are installed at one time.

Each 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 features installed on the storage array, which features can be evaluated for a trial period, which features are enabled, and which 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 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 space.

- `driveSecurity`

### Minimum firmware level

8.25 removes all attributes that are no longer valid.

### Related references

[Disable storage array feature](#) on page 197

[Show storage array](#) on page 410

[Disable storage array feature](#) on page 197

## 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 name of the volume group that you want to place in the Security Enabled state. Enclose the volume group name 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 feature must be enabled.
- The storage array security key has to be set.
- The volume group is Optimal, and it does not have 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

### Related references

[Show volume group](#) on page 433

## 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 that comprise 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

## Related references

- [Add volume to asynchronous mirror group](#) on page 80
- [Cancel asynchronous mirror group role reversal](#) on page 88
- [Check asynchronous mirror group consistency](#) on page 91
- [Check storage array connectivity](#) on page 94
- [Clear asynchronous mirroring fault](#) on page 96
- [Create asynchronous mirror group](#) on page 104
- [Deactivate asynchronous mirroring](#) on page 176
- [Delete asynchronous mirror group](#) on page 177
- [Remove incomplete asynchronous mirrored pair from asynchronous mirror group](#) on page 236
- [Remove volume from asynchronous mirror group](#) on page 239
- [Reset asynchronous mirror group statistics](#) on page 246
- [Reset iSCSI IP address](#) on page 250
- [Resume asynchronous mirror group](#) on page 255
- [Save asynchronous mirror group statistics](#) on page 265
- [Set asynchronous mirror group](#) on page 294
- [Show asynchronous mirror groups](#) on page 382
- [Show asynchronous mirror group synchronization progress](#) on page 383
- [Start asynchronous mirroring synchronization](#) on page 440
- [Suspend asynchronous mirror group](#) on page 481
- [Test asynchronous mirror group connectivity](#) on page 484

## 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. Enclose the pass phrase in double quotation marks (" ").
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" <b>IMPORTANT</b> – You must add a file extension of .slk 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" <b>IMPORTANT</b> – 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 or volume group

This command increases the capacity of either a standard volume or a repository volume in a disk pool or volume group. 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 volume in a disk pool or volume group for which you want to increase capacity. Enclose the volume name in double quotation marks (" ").
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 24 . Enclose the tray ID value, drawer ID value, and the slot ID value in parentheses.  <b>Important:</b> The addDrives parameter can only be used to increase the capacity of a volume group. The parameter cannot be used to increase the capacity of a disk pool.

## 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

## Related references

[Create volume in disk pool](#) on page 172

[Delete volume from disk pool](#) on page 190

[Set volume attributes for a volume in a disk pool](#) on page 366

## Initialize thin volume

This command initializes or re-initializes a thin volume. The action depends on the parameters used:

- 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 on which you are starting the initialization. Enclose the volume name in square brackets ([ ]). If the volume name contains special characters or numbers, you must enclose the volume name in double quotation marks (" ") inside square brackets.
existingRepositoryLabel	<p>This parameter replaces the repository volume by the candidate volume specified.</p> <ul style="list-style-type: none"> <li>• 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.</li> <li>• If the newly specified repository volume is on a different disk pool, the thin volume will change ownership to that pool.</li> <li>• The old repository volume will be deleted by default.</li> </ul> <p>If the existing repository volume name contains special characters or consists only of numbers, you must enclose the volume name in double quotation marks (" ").</p>
diskPool	<p>The name of the disk pool in which you want to create a new repository volume with the specified capacity. If the disk pool name contains special characters or consists only of numbers, you must enclose the disk pool name in double quotation marks (" ").</p> <p>You must use this parameter with the <code>capacity</code> parameter to create a new repository volume with the specified capacity.</p>

Parameter	Description
capacity	<p>The size that you want to set for the repository volume that you are creating. Size is defined in units of <code>bytes</code>, <code>KB</code>, <code>MB</code>, <code>GB</code>, or <code>TB</code>.</p> <p>The minimum physical capacity is 4 GB.</p> <p>The maximum physical capacity is 257 TB.</p> <p>You must use this parameter with the <code>diskPool</code> parameter to create a new repository volume with the specified capacity.</p>
retainRepositoryMembers	<p>If this parameter is set to <code>TRUE</code>, the old repository is retained. By default, the old repository is deleted. This parameter is ignored if the existing repository is reused.</p>

### 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.

**Restriction:** Immediate Availability Format (IAF) does not apply to volumes greater than 64 TB.

The following table lists the capacity limits for a thin volume.

Type of Capacity	Size
Minimum virtual capacity	32 MB
Maximum virtual capacity	256 TB
Minimum physical capacity	4 GB
Maximum physical capacity	257 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 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

8.30 increases the maximum capacity of a thin volume to 256 TB.

### Related references

[Set thin volume attributes](#) on page 362

[Show thin volume](#) on page 428

## 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 includes data for RAID configuration, volume groups and disk pools.

Before using this command with the `file` option, you must first obtain a validator string (a security code) from technical support. To obtain a validator, use the `save storageArray dbmValidator` command to generate an XML file that contains validator information. Technical support 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= all])
```

### Parameters

Parameter	Description
<code>file</code>	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: <code>file="C:\Array Backups\DBMbackup_03302010.dbm"</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.
<code>validator</code>	The alphanumeric security code that is required to restore a storage array to an existing configuration. You must get the validator from technical support. To get the validator, generate the required validation information XML file by running the <code>save storageArray dbmValidator</code> command. After the validation information XML file is available, contact technical support to get the validator. Enclose the validator in double quotation marks (" ").
<code>sourceLocation</code>	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> .
<code>controller</code>	This parameter specifies the controller from which data will be exclusively retrieved. 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.
<code>contentType</code>	This parameter specifies the content type of the data that will be retrieved. This parameter is set to <code>all</code> by default, so that all of the data including disk pool configuration data is retrieved.

## 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

## Related references

[Save storage array DBM database](#) on page 273

[Save storage array DBM validator information file](#) on page 274

[Show storage array DBM database](#) on page 420

[Save storage array DBM database](#) on page 273

[Save storage array DBM validator information file](#) on page 274

[Show storage array DBM database](#) on page 420

## 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 snapshot image volume copy pairs.

This command works with volume copy pairs that you created 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 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 must enclose the target volume name in double quotation marks (" ") inside square brackets.



Parameter	Description
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 must enclose the source volume name in double quotation marks (" ") inside square brackets.
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

7.83 removes the copyType=(online | offline) parameter.

## Recover disabled drive ports

This command is used to tell the controller that corrective action has been taken to recover from a mis-wire condition. The controller will re-enable any SAS ports that were disabled when a mis-wire was detected.

### Syntax

```
recover disabled drivePorts;
```

### Examples

```
recover disabled drivePorts;
```

### Related references

[Recover SAS port mis-wire](#) on page 229

## 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 (recovery-profile.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 24. 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 <code>show storageArray profile</code> command.)
newVolumeGroup	The name that you want to give a new volume group. Enclose the new volume group name in double quotation marks (" ").
userLabel	<p>The name of the volume that you want to recover. The user label has two parts:</p> <ul style="list-style-type: none"> <li>The volume name. Enclose the volume name in double quotation marks (" ").</li> <li>The world wide name of the volume, <code>volumeWWN</code>, in the form of a 16 byte identifier, for example, <code>60080E500017B4320000000049887D77</code>. Enclose the identifier in double quotation marks (" ").</li> </ul> <p>You must enter both the volume name and the world wide name of the volume. You must enter the volume name first. For example:</p> <pre>userLabel="engdata" volumeWWN=60080E500017B4320000000049887D77</pre>
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.

Parameter	Description
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 <code>TRUE</code> . To turn off allocation, set this parameter to <code>FALSE</code> .
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 <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.
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> .
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"> <li>• <code>none</code> – The volume group does not have data assurance protection.</li> <li>• <code>enabled</code> – The volume group has data assurance protection. The volume group supports protected information and is formatted with protection information enabled.</li> </ul> <p><b>Note:</b> This option is only valid for enablement if the drives support DA.</p>

## 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 pre-allocates the space in a volume for future segment size changes. The pre-allocated space is the maximum allowable segment size. The `dssPreAllocate` parameter is necessary for properly recovering volume configurations that are not retrievable from the controller database. 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.

### Related references

[Check volume parity](#) on page 95

[Repair volume parity](#) on page 243

## Recover SAS port mis-wire

This command tells a controller that corrective action has been taken to recover from a mis-wire condition. The controller can then re-enable any SAS ports that were disabled because a mis-wire was detected. Run this command after you have corrected any SAS mis-wire conditions.

SAS mis-wire conditions can be one of these:

- Invalid SAS topologies
- Controller host port to expansion tray connection
- Initiator to drive topology connection
- ESM mis-wire

### Syntax

```
recover sasPort miswire
```

### Parameters

None.

### Notes

When a mis-wire condition is detected the controller firmware reports the following:

- A critical MEL event indicating the presence of a mis-wire condition
- A “Needs Attention” condition
- A Recovery Action indicating the mis-wire

In addition, the controller reports “Redundancy Loss” information for the drives, ESMs, and trays for which it has only one access path due to the mis-wire (for example the bottom two trays in a storage array). The Redundancy Loss information is not unique to SAS topologies.

Following are the event notifications related to SAS port mis-wire:

- SAS mis-wire detected
- HBA connected to drive channel
- Drive channels cross wired

For mis-wire conditions where the controller disables one or more SAS ports to protect the integrity of the SAS domain, take these steps to recover from the condition:

1. Determine which cable is in the wrong place and remove it. If the last action taken was to install a new cable, this is the most likely candidate. If not, look for cables attached to devices that are not part of the storage array or cables attached between channels. The recovery event should report the channel on which the mis-wire was detected, and possibly the tray.
2. After you remove a cable or move the cable to the correct location, run the `recover sasPort miswire` command. The controller will attempt to re-enable any SAS ports that were disabled when the mis-wire was detected.

3. If the cable that you removed or moved to correct for the mis-wire provided the controller its only access to the disabled port, the controller will not be able to re-enable the SAS port that was disabled. To complete the recovery, you must power cycle the trays in the storage array.

### Minimum firmware level

8.10

### Related references

[Recover disabled drive ports](#) on page 225

## Re-create synchronous mirroring repository volume

**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 returns 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,drawerID1,slotID1
... trayIDN,drawerIDN,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
<code>repositoryRAIDLevel</code>	The RAID level for the mirror repository volume. Valid values are 1, 3, 5, or 6.

Parameter	Description
repositoryDrives	<p>The drives that you want to use 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 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 24. Enclose the tray ID value, the drawer ID value, and the slot ID value in square brackets ([ ]). Enter the names of the repository volumes using these rules:</p> <ul style="list-style-type: none"> <li>• Enclose all of the tray ID values, the drawer ID values, and the slot ID value sin square brackets ([ ]).</li> <li>• Separate the tray ID value, the drawer ID value, and the slot ID value with commas.</li> <li>• Separate each of the drive positions with a space.</li> </ul>
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 data-bbox="721 1121 1377 1171">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. 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"> <li>• fibre</li> <li>• SATA</li> <li>• SAS</li> </ul> <p>If you do not specify a drive type, the command defaults to <code>fibre</code>.</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 <code>TRUE</code> . The default value is <code>FALSE</code> .

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"> <li>• none – The volume group does not have data assurance protection.</li> <li>• enabled – The volume group has data assurance protection. The volume group supports protected information and is formatted with protection information enabled.</li> </ul> <p><b>Note:</b> This option is only valid for enablement if the drives support DA.</p>

### 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.

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.

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.

### Related references

[Check synchronous mirroring status](#) on page 95

[Diagnose synchronous mirroring](#) on page 196

[Set synchronous mirroring](#) on page 360



[Show synchronous mirroring volume candidates](#) on page 426

[Show synchronous mirroring volume synchronization progress](#) on page 427

## Reduce disk pool capacity

This command reduces the capacity of the disk pool by logically removing the drives from the pool. This command uses the Dynamic Capacity Reduction (DCR) feature, which enables you to remove selected drives from a disk pool and re-use the removed drives as needed.

The maximum number of drives that you can remove from a disk pool in a single DCR operation is 12. You cannot reduce the disk pool to below the minimum disk pool size.

### Syntax

```
set diskPool [diskPoolName]
removeDrives=(trayID1,drawerID1,slotID1
... trayIDn,drawerIDn,slotIDn)
```

### Parameter

Parameter	Description
diskPool	The name of the disk pool for which you want to reduce capacity. Enclose the disk pool name in square brackets ([ ]).
removeDrives	The drives that you want to remove from the disk pool. 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 remove. For low-capacity drive trays, specify the tray ID value and the slot ID value for each drive 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 24 . Enclose the tray ID values, the drawer ID values, and the slot ID values in parentheses ([ ]).

### 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 might want to remove selected drives from a disk pool if you need to create a separate group or remove unneeded capacity from the storage array. The removed drives become unassigned drives and the data that existed on them is redistributed across the remaining drives in the disk pool. The capacity of the disk pool is decreased by the capacity of the drives that you remove. Removing drives from a disk pool is always carried out as a background task and the volumes remain fully accessible during this operation. The progress of the removal operation is reported as part of the long running operations status.

This command fails if there is not enough free capacity on the drives that remain in the disk pool to store all of the data that was on the drives you remove from the disk pool. Reserve capacity in the disk pool might be used to store data that was on the drives you remove from the disk pool. However, if the reserve capacity falls below the critical threshold, a critical event is written to the major event log.

Decreasing the disk pool size might result in a smaller necessary reserve capacity. In some cases, in order to decrease the likelihood that this command will fail, reduce your reserve capacity before running this command in order to increase the amount of space available for disk pool data.

The controller firmware determines how many drives that you can remove and you can then select the specific drives to be removed. The controller firmware bases the count on the amount of free space needed to redistribute the data across the remaining drives without consuming any capacity reserved for reconstructions. If the free capacity of the disk pool is already less than the Reserved Reconstruction Drive Amount, the controller does not allow a DCR operation to begin.

A DCR operation might result in the configured capacity exceeding one or both of the Pool Utilization Thresholds. If so, normal threshold alerts are issued.

### Minimum firmware level

8.10

### Related references

[Set disk pool](#) on page 307

[Set disk pool \(modify disk pool\)](#) on page 310

## Register SNMP community

This command creates a new Simple Network Management Protocol (SNMP) community, community name string and registers the new community as a known entity to the SNMP agent. Any attempt to register a new community with the same name as that of an existing community is rejected and the firmware returns an error message.

### Syntax

```
create snmpCommunity communityName="snmpCommunityName"
```

### Parameters

Parameter	Description
communityName	The name of the SNMP community that you want to create. Enclose the SNMP community name in double quotation marks (" ").

### Minimum firmware level

8.30

## Register SNMP trap destination

This command creates a new Simple Network Management Protocol (SNMP) trap destination. A trap destination is the SNMP manager that you want to have receive the trap messages.

### Syntax

```
create snmpTrapDestination trapReceiverIP=ipAddress
communityName="communityName"
[sendAuthenticationFailureTraps=(TRUE | FALSE)]
```

### Parameters

Parameter	Description
trapReceiverIP	The IP address of the SNMP manager to which you want to send trap messages.

Parameter	Description
communityName	The name of the SNMP community for which you want to send trap messages.
sendAuthenticationFailureTraps	This parameter turns on or turns off sending authentication failure messages to an SNMP manager. To send authentication failure messages, set the parameter to <code>TRUE</code> . To prevent sending authentication failure messages, set the parameter to <code>FALSE</code> . The default is <code>TRUE</code> .

### Minimum firmware level

8.30

## 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)
```

### Parameters

Parameter	Description
ssdCache	The name of the SSD cache from which you want to remove SSDs. Enclose the SSD cache name in square brackets ( [ ] ). If the SSD cache name has special characters or consists only of numbers, you must enclose the SSD cache name in double quotation marks ( " ") inside square brackets.
removeDrives	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 24 . 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

### Related references

[Add drives to SSD cache](#) on page 77

[Change SSD cache application type](#) on page 90

[Create SSD cache](#) on page 159

[Delete SSD cache](#) on page 188

[Rename SSD cache](#) on page 242  
[Resume SSD cache](#) on page 260  
[Set SSD cache for a volume](#) on page 342  
[Show SSD cache](#) on page 406  
[Show SSD cache statistics](#) on page 407  
[Start SSD cache locate](#) on page 457  
[Start SSD cache performance modeling](#) on page 458  
[Stop SSD cache locate](#) on page 475  
[Stop SSD cache performance modeling](#) on page 476  
[Suspend SSD cache](#) on page 482

## 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 ( " ") 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

### Related references

[Add volume to asynchronous mirror group](#) on page 80  
[Cancel asynchronous mirror group role reversal](#) on page 88  
[Check asynchronous mirror group consistency](#) on page 91  
[Check storage array connectivity](#) on page 94  
[Clear asynchronous mirroring fault](#) on page 96  
[Create asynchronous mirror group](#) on page 104

[Deactivate asynchronous mirroring](#) on page 176  
[Delete asynchronous mirror group](#) on page 177  
[Establish asynchronous mirrored pair](#) on page 216  
[Remove volume from asynchronous mirror group](#) on page 239  
[Reset asynchronous mirror group statistics](#) on page 246  
[Reset iSCSI IP address](#) on page 250  
[Resume asynchronous mirror group](#) on page 255  
[Save asynchronous mirror group statistics](#) on page 265  
[Set asynchronous mirror group](#) on page 294  
[Show asynchronous mirror groups](#) on page 382  
[Show asynchronous mirror group synchronization progress](#) on page 383  
[Start asynchronous mirroring synchronization](#) on page 440  
[Suspend asynchronous mirror group](#) on page 481  
[Test asynchronous mirror group connectivity](#) on page 484

## Remove member volume from consistency group

This command removes a member volume from 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.

### Minimum firmware level

7.83

### Related references

[Add member to consistency group](#) on page 78  
[Create consistency group](#) on page 107  
[Create consistency group snapshot image](#) on page 115  
[Create consistency group snapshot volume](#) on page 116  
[Create consistency group snapshot volume mapping](#) on page 120  
[Delete consistency group](#) on page 178  
[Delete consistency group snapshot image](#) on page 179  
[Delete consistency group snapshot volume](#) on page 180

- [Resume consistency group snapshot volume](#) on page 256
- [Set consistency group attributes](#) on page 296
- [Set consistency group snapshot volume](#) on page 298
- [Show consistency group](#) on page 384
- [Show consistency group snapshot image](#) on page 386
- [Start consistency group snapshot rollback](#) on page 443
- [Stop consistency group snapshot rollback](#) on page 466
- [Stop consistency group snapshot volume](#) on page 468
- [Stop pending snapshot images on consistency group](#) on page 471

## Remove synchronous mirroring

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])
```

### Parameters

Parameter	Description
<code>localVolume</code>	The name of the primary volume (the volume on the local storage array) that you want to remove. Enclose the volume name in square brackets ([ ]). If the volume name has special characters or consists only of numbers, you must enclose the volume name in double quotation marks (" ") inside square brackets.
<code>localVolumes</code>	The names of several primary volumes (the volumes on the local storage array) that you want to remove. Enter the names of the primary volumes using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Separate each of the names with a space.</li> </ul> If the primary volume names have special characters or consist only of numbers, enter the names using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Enclose each of the names in double quotation marks (" ").</li> <li>• Separate each of the names with a space.</li> </ul>

### Minimum firmware level

6.10

### Related references

- [Check synchronous mirroring status](#) on page 95

[Create synchronous mirroring](#) on page 162

[Re-create synchronous mirroring repository volume](#) on page 230

[Resume synchronous mirroring](#) on page 261

[Set synchronous mirroring](#) on page 360

## 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 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 or consists only of numbers, you must enclose the target volume name in double quotation marks (" ") inside square brackets.
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 or consists only of numbers, you must enclose the source volume name in double quotation marks (" ") inside square brackets.

### Minimum firmware level

5.40

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 ([ ]).

Parameter	Description
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

### Related references

- [Add volume to asynchronous mirror group](#) on page 80
- [Cancel asynchronous mirror group role reversal](#) on page 88
- [Check asynchronous mirror group consistency](#) on page 91
- [Check storage array connectivity](#) on page 94
- [Clear asynchronous mirroring fault](#) on page 96
- [Create asynchronous mirror group](#) on page 104
- [Deactivate asynchronous mirroring](#) on page 176
- [Delete asynchronous mirror group](#) on page 177
- [Establish asynchronous mirrored pair](#) on page 216
- [Remove incomplete asynchronous mirrored pair from asynchronous mirror group](#) on page 236
- [Reset asynchronous mirror group statistics](#) on page 246
- [Reset iSCSI IP address](#) on page 250
- [Resume asynchronous mirror group](#) on page 255
- [Save asynchronous mirror group statistics](#) on page 265
- [Set asynchronous mirror group](#) on page 294
- [Show asynchronous mirror groups](#) on page 382
- [Show asynchronous mirror group synchronization progress](#) on page 383
- [Start asynchronous mirroring synchronization](#) on page 440
- [Suspend asynchronous mirror group](#) on page 481
- [Test asynchronous mirror group connectivity](#) on page 484

## Remove volume LUN mapping

This command removes the logical unit number (LUN) mapping from one or more volumes.

### Syntax for removing LUN mapping from one volume

```
remove (volume ["volumeName"] | accessVolume) lunMapping
(host="hostName" | hostGroup=("hostGroupName" | defaultGroup))
```

### Syntax for removing LUN mapping from more than one volume LUN mapping

```
remove (allVolumes | volumes [volumeName1 ... volumeNameN])
lunMapping
```



## Parameters

Parameter	Description
volume	The name of a volume from which you are removing the LUN mapping. Enclose the volume name double quotation marks ( " ") inside square brackets ( [ ] ).
accessVolume	This parameter removes the access volume.
allVolumes	This parameter removes the LUN mapping from all of the volumes.
volumes	<p>The names of several volumes from which you are removing the LUN mapping. Enter the names of the volumes using these rules:</p> <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ( [ ] ).</li> <li>• Separate each of the names with a space.</li> </ul> <p>If the volume names have special characters or consist only of numbers, enter the names using these rules:</p> <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ( [ ] ).</li> <li>• Enclose each of the names in double quotation marks ( " " ).</li> <li>• Separate each of the names with a space.</li> </ul>
host	<p>The name of the host to which the volume is mapped. Enclose the host name in double quotation marks ( " " ).</p> <p>This parameter works only when you are removing LUN mapping from one volume.</p>
hostGroup	<p>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.</p> <p>This parameter works only when you are removing LUN mapping from one volume.</p>

## 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

**Related references**

[Set volume mapping](#) on page 380

[Show volume reservations](#) on page 437

**Rename snapshot volume**

This command renames an existing snapshot volume.

**Syntax**

```
set snapVolume [ "snapVolumeName" ] userLabel="snapImageVolumeName "
```

**Parameters**

Parameter	Description
snapVolume	The name of the snapshot volume that you want to rename. Enclose the snapshot volume name in double quotation marks ( " ") inside of square brackets ( [ ] ).
userLabel	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

**Related references**

[Create read-only snapshot volume](#) on page 145

[Create snapshot volume](#) on page 156

[Delete snapshot volume](#) on page 187

[Resume snapshot volume](#) on page 259

[Revive snapshot volume](#) on page 263

[Set read-only snapshot volume to read/write volume](#) on page 328

[Set snapshot volume media scan](#) on page 339

[Set snapshot volume repository volume capacity](#) on page 340

[Set SSD cache for a volume](#) on page 342

[Show snapshot volumes](#) on page 403

[Stop snapshot volume](#) on page 474

**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 name of the SSD cache that you want to rename. Enclose the SSD cache name in square brackets ([ ]). If the SSD cache name has special characters or consists only of numbers, you also must enclose the SSD cache name 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

**Related references**

- [Add drives to SSD cache](#) on page 77
- [Change SSD cache application type](#) on page 90
- [Create SSD cache](#) on page 159
- [Delete SSD cache](#) on page 188
- [Remove drives from SSD cache](#) on page 235
- [Resume SSD cache](#) on page 260
- [Set SSD cache for a volume](#) on page 342
- [Show SSD cache](#) on page 406
- [Show SSD cache statistics](#) on page 407
- [Start SSD cache locate](#) on page 457
- [Start SSD cache performance modeling](#) on page 458
- [Stop SSD cache locate](#) on page 475
- [Stop SSD cache performance modeling](#) on page 476
- [Suspend SSD cache](#) on page 482

**Repair volume parity**

This command repairs the parity errors on a volume.

**Caution:** If volume parity and volume data do not match, volume data might be defective rather than volume parity. Repairing volume parity destroys the ability to repair the defective volume data, if necessary.

**Syntax**

```
repair volume [volumeName] parity
parityErrorFile="filename"
[verbose=(TRUE | FALSE)]
```

**Parameters**

Parameter	Description
volume	The name of the volume for which you are repairing parity. Enclose the volume name in square brackets ([ ]). If the volume name has special characters or numbers, you 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"  <b>Note:</b> A parityErrorFile is the result of a previous check volume parity operation.
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

**Related references**[Check volume parity](#) on page 95**Replace drive**

This command replaces a drive in a volume group. Replacing a drive in a volume group 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.

Optionally, you can copy the data from the drive that you want to replace and then fail the drive so that it will not be reused. You can then replace the failed drive with a good drive.

**Syntax**

```
replace drive([trayID,drawerID,slotID] | <"wwID">)
replacementDrive=trayID,drawerID,slotID
[copyDrive]
[failDrive]
```

**Parameters**

Parameter	Description
drive	The location of the drive that you want to replace. 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 24 . Enclose the tray ID value, the drawer ID value, and the slot ID value in square brackets ([ ]).

Parameter	Description
<code>drive</code>	The World Wide Identifier (WWID) of the drive that you want to replace. Enclose the WWID in double quotation marks ( " ") inside angle brackets (< >).
<code>replacementDrive</code>	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 24 .
<code>copyDrive</code>	This parameter copies the data from the drive that you are replacing to the replacement drive.
<code>failDrive</code>	This parameter causes the drive that you are replacing to be failed after the data has been copied to the replacement drive.  This parameter can be used only when the <code>copyDrive</code> parameter is used.

## 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.

8.20 adds these parameters:

- `copyDrive`
- `failDrive`

## Related references

[Set drive hot spare](#) on page 314

[Set drive service action allowed indicator](#) on page 314

[Set drive state](#) on page 315

[Set FIPS drive security identifier](#) on page 318

[Set foreign drive to native](#) on page 319

[Show drive](#) on page 393

[Start drive channel fault isolation diagnostics](#) on page 447

[Start drive initialize](#) on page 449

[Start drive reconstruction](#) on page 451

[Stop drive channel fault isolation diagnostics](#) on page 469

[Reset drive](#) on page 249

[Set drive state](#) on page 315

[Stop drive replace](#) on page 471

## 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 or numbers, you 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"> <li><code>all</code>— Data for all three sample types are reset.</li> <li><code>mostRecent</code> — Statistics are reset for the most recent 50 resynchronization samples.</li> <li><code>longestSyncTime</code> — Statistics are reset for the most recent 20 longest resynchronization samples.</li> <li><code>errors</code>— Statistics are reset for the most recent 20 failed resynchronization samples.</li> </ul>

### 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

### Related references

[Save asynchronous mirror group statistics](#) on page 265

## Reset AutoSupport message collection schedule

The `reset autoSupportSchedule reset` command resets the daily and weekly times and days of the week that AutoSupport messages are sent. You can accept default random values generated by the management software, or you can specify values with this command.

### Context

**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.

**Restriction:** This command is obsolete for the E2800 storage array. Use the RESTful API, SANtricity System Manager, or see [Using curl commands with the E2800 storage array](#) on page 39.

- If AutoSupport is enabled, the management software sends both daily AutoSupport messages and weekly AutoSupport messages.
- The management software randomly selects a time of day for both daily and weekly messages and a day of the week for weekly messages. Management software makes all attempts to ensure that no two storage arrays within a management domain sends scheduled AutoSupport messages at the same time.
- You can specify a range (hour-level granularity) of when to send daily and when to send weekly messages for the storage arrays in your management domain.
- For the weekly schedule, you select preferred days of the week for AutoSupport collection and transmission.

### Syntax

```
SMcli -autoSupportSchedule reset | (dailyTime=<startTime>-<endTime>
  dayOfWeek=[Sunday Monday Tuesday Wednesday Thursday Friday Saturday]
  weeklyTime=<startTime>-<endTime>)
```

### Parameters

Parameter	Description
reset	Resets and generates new schedules using a 24 hour day and 7 day weeks (Sunday to Saturday) time interval.
dailyTime	<startTime> - <endTime> Specifies the time of day that you want to start and end the collection of AutoSupport data for all storage arrays. The startTime and endTime must be in the form of HH:MM[am pm].

Parameter	Description
dayOfWeek	[Sunday Monday Tuesday Wednesday Thursday Friday Saturday] Specifies the preferred days of the week (Sunday through Saturday) that you want to collect AutoSupport bundle collection data for all storage arrays. The <i>dayOfWeek</i> parameter must be surrounded by brackets and separated with a space.
weeklyTime	<startTime> - <endTime> Specifies the time of day that you want to start and end the collection of the AutoSupport bundle collection data for each day of the week that you have selected. The <i>startTime</i> and <i>endTime</i> must be in the form of HH:MM[am pm].

### Minimum Firmware Level

8.25

### Example

```
SMcli -autoSupportSchedule dailyTime=10:00am-11:00am
dayOfWeek=[Monday Thursday Friday] weeklyTime=2:00am-3:00am
```

```
SMcli -autoSupportSchedule dailyTime=10:00am-11:00am dayOfWeek=[Monday
Thursday Friday] weeklyTime=2:00am-3:00am
```

### Verification

Use the `SMcli -autoSupportSchedule show` command to see the resulting change to the schedule in the management domain.

### Related references

[Using curl commands with the E2800 storage array](#) on page 39

[Capture or view an AutoSupport log](#) on page 89

[Display AutoSupport bundle collection settings](#) on page 198

[Display AutoSupport message collection schedule](#) on page 200

[Disable AutoSupport at the EMW level SMcli version](#) on page 197

[Enable or disable AutoSupport at the EMW management domain level \(SMcli only\)](#) on page 210

[Enable or disable the AutoSupport OnDemand feature at the EMW management domain level \(SMcli only\)](#) on page 211

[Enable or disable the AutoSupport OnDemand Remote Diagnostics feature at the EMW management domain level \(SMcli only\)](#) on page 213

[Save storage array support data](#) on page 283

[Specify the AutoSupport delivery method](#) on page 438

[Test the AutoSupport configuration](#) on page 485

## 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.



**Caution:** If you use in-band management, you cannot control which controller any command is sent to, and this command can have unexpected results.

## Syntax

```
reset controller [(a|b)]
```

## Parameters

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, the storage management software 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

## Related references

[Set controller](#) on page 299

[Show controller](#) on page 387

## Reset drive

This command power cycles a drive in a volume group or a disk pool to aid in the recovery of a drive that is exhibiting inconsistent or non-optimal behavior.

By power cycling a drive, you can clear some errors that cause inconsistent or non-optimal behavior. This avoids replacing a drive when it is experiencing only a transient, non-fatal error, and the drive can remain operational. Resetting a drive in this way reduces disruptions and avoids replacing a drive.

If the problem cannot be corrected by power cycling the drive, the data is copied from the drive and the drive is powered down for replacement.

## Syntax

```
reset drive([trayID,drawerID,slotID] | <"wwID">)
```

## Parameters

Parameter	Description
drive	The location of the drive that you want to replace. 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 24 . Enclose the tray ID value, the drawer ID value, and the slot ID value in square brackets ([ ]).
drive	The World Wide Identifier (WWID) of the drive that you want to replace. Enclose the WWID in double quotation marks (" ") inside angle 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.

After power cycling the drive, the controller must verify the drive to be functional before it can be placed back into use. If the drive cannot be verified it is marked as a failed drive. Power cycling a failed drive requires global hot spare drives and full copy back to a replacement drive after it is brought into service.

Failed drives cannot be reset by a power cycle more than once in a 24 hour period, and in some cases cannot be reset more than once. Thresholds and counters for drive power cycles are kept in persistent storage, and are included in state capture data. Informational events are logged when a drive is power cycled.

## Minimum firmware level

8.20

## Related references

[Replace drive](#) on page 244

[Set drive state](#) on page 315

[Stop drive replace](#) on page 471

## Reset iSCSI IP address

This command resets the IP address of 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 has changed and needs to be updated.

When establishing an asynchronous mirroring relationship with an iSCSI connection, both the local and the remote storage arrays store a record of the IP address of the remote storage array in 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 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 technical support.

### Syntax

```
reset storageArray diagnosticData
```

### Parameters

None.

### Minimum firmware level

6.16

**Related references**

[Save storage array diagnostic data](#) on page 275

[Start storage array configuration database diagnostic](#) on page 459

[Stop storage array configuration database diagnostic](#) on page 478

**Reset storage array Infiniband statistics baseline**

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

**Related references**

[Save storage array InfiniBand statistics](#) on page 278

**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

**Related references**

[Activate iSCSI initiator](#) on page 72  
[Create iSCSI initiator](#) on page 131  
[Delete iSCSI initiator](#) on page 184  
[Diagnose controller iSCSI host cable](#) on page 193  
[Save storage array iSCSI statistics](#) on page 279  
[Set iSCSI initiator](#) on page 324  
[Set iSCSI target properties](#) on page 326  
[Show current iSCSI sessions](#) on page 390  
[Show storage array negotiation defaults](#) on page 423  
[Show storage array unconfigured iSCSI initiators](#) on page 424  
[Start iSCSI DHCP refresh](#) on page 453  
[Stop storage array iSCSI session](#) on page 479  
[Save storage array iSCSI statistics](#) on page 279  
[Show storage array unconfigured iSCSI initiators](#) on page 424  
[Stop storage array iSCSI session](#) on page 479

**Reset storage array iSER baseline**

This command resets the iSER baseline to 0 for the storage array.

**Syntax**

```
reset storageArray iserStatsBaseline
```

**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**

8.20

**Related references**

[Save storage array iSER statistics](#) on page 280  
[Set iSER target](#) on page 327

## 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 ILevel

5.00

### Related references

[Save storage array RLS counts](#) on page 281

## 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.

### Related references

[Save storage array SAS PHY counts](#) on page 282

## 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

**Related references**

[Save storage array SOC counts](#) on page 282

**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"]
[deleteRecoveryPointIfNecessary]
```

**Parameters**

Parameter	Description
asyncMirrorGroup	The name of the asynchronous mirror group for which you want to start synchronization. Enclose the asynchronous mirror group name in double quotation marks ( " ") inside square brackets ([ ]).
deleteRecoveryPointIfNecessary	The parameter to delete the recovery point if the recoverable synchronization data has exceeded time threshold for recovery. Recovery point age is measured from the time the data was frozen on the primary storage array.

**Minimum firmware level**

7.84

8.10 adds the `deleteRecoveryPointIfNecessary` parameter.**Related references**

[Suspend asynchronous mirror group](#) on page 481

**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"
```

**Parameters**

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 ([ ]).



Parameter	Description
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"> <li>• The name of the consistency group</li> <li>• An identifier for the snapshot image in the consistency group.</li> </ul> <p>The identifier for the snapshot image can be one of these:</p> <ul style="list-style-type: none"> <li>• An integer value that is the sequence number of the snapshot in the consistency group.</li> <li>• <b>NEWEST</b> - Use this option when you want to show the latest snapshot image created in the consistency group.</li> <li>• <b>OLDEST</b> - Use this option when you want to show the earliest snapshot image created in the consistency group.</li> </ul> <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

### Related references

- [Add member to consistency group](#) on page 78
- [Create consistency group](#) on page 107
- [Create consistency group snapshot image](#) on page 115
- [Create consistency group snapshot volume](#) on page 116
- [Create consistency group snapshot volume mapping](#) on page 120
- [Delete consistency group](#) on page 178
- [Delete consistency group snapshot image](#) on page 179
- [Delete consistency group snapshot volume](#) on page 180
- [Remove member volume from consistency group](#) on page 237
- [Set consistency group attributes](#) on page 296
- [Set consistency group snapshot volume](#) on page 298
- [Show consistency group](#) on page 384
- [Show consistency group snapshot image](#) on page 386
- [Start consistency group snapshot rollback](#) on page 443
- [Stop consistency group snapshot rollback](#) on page 466

[Stop consistency group snapshot volume](#) on page 468

[Stop pending snapshot images on consistency group](#) on page 471

## 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"> <li>The name of the snapshot group</li> <li>An identifier for the snapshot image in the snapshot group</li> </ul> <p>The identifier for the snapshot image can be one of these:</p> <ul style="list-style-type: none"> <li>An integer value that is the sequence number of the snapshot in the snapshot group.</li> <li>NEWEST - Use this option when you want to show the latest snapshot image created in the snapshot group.</li> <li>OLDEST - Use this option when you want to show the earliest snapshot image created in the snapshot group.</li> </ul> <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

**Related references**

- [Create snapshot image](#) on page 155
- [Delete snapshot image](#) on page 186
- [Show snapshot image](#) on page 401
- [Start snapshot image rollback](#) on page 456
- [Stop snapshot image rollback](#) on page 473

**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 snapshot volume for which you want to resume operations. Enclose the snapshot volume name in double quotation marks ( " ") inside square brackets ( [ ] ).
snapImage	<p>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:</p> <ul style="list-style-type: none"> <li>• The name of the snapshot group</li> <li>• An identifier for the snapshot image in the snapshot group</li> </ul> <p>The identifier for the snapshot image can be one of these:</p> <ul style="list-style-type: none"> <li>• An integer value that is the sequence number of the snapshot in the snapshot group.</li> <li>• <b>NEWEST</b> - Use this option when you want to show the latest snapshot image created in the snapshot group.</li> <li>• <b>OLDEST</b> - Use this option when you want to show the earliest snapshot image created in the snapshot group.</li> </ul> <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 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 ["snapVol1"] 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 ["snapVol1"] snapImage="snapGroup1:newest";
```

### Minimum firmware level

7.83

### Related references

- [Create read-only snapshot volume](#) on page 145
- [Create snapshot volume](#) on page 156
- [Delete snapshot volume](#) on page 187
- [Rename snapshot volume](#) on page 242
- [Revive snapshot volume](#) on page 263
- [Set read-only snapshot volume to read/write volume](#) on page 328
- [Set snapshot volume media scan](#) on page 339
- [Set snapshot volume repository volume capacity](#) on page 340
- [Set SSD cache for a volume](#) on page 342
- [Show snapshot volumes](#) on page 403
- [Stop snapshot volume](#) on page 474

## 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 name of the SSD cache for which you want to resume caching operations. Enclose the SSD cache name in square brackets ([ ]). If the SSD cache name has special characters or consists only of numbers, you must enclose the SSD cache name in double quotation marks (" ") inside square brackets.

### Minimum firmware level

7.84

### Related references

- [Add drives to SSD cache](#) on page 77
- [Change SSD cache application type](#) on page 90
- [Create SSD cache](#) on page 159
- [Delete SSD cache](#) on page 188
- [Remove drives from SSD cache](#) on page 235
- [Rename SSD cache](#) on page 242
- [Set SSD cache for a volume](#) on page 342
- [Show SSD cache](#) on page 406
- [Show SSD cache statistics](#) on page 407

[Start SSD cache locate](#) on page 457

[Start SSD cache performance modeling](#) on page 458

[Stop SSD cache locate](#) on page 475

[Stop SSD cache performance modeling](#) on page 476

[Suspend SSD cache](#) on page 482

## 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
<code>primary</code>	The name of the primary volume for which you want to resume a Synchronous Mirroring operation. Enclose the primary volume name in square brackets ([ ]). If the primary volume name has special characters or numbers, you must enclose the primary volume name in double quotation marks (" ") inside square brackets.
<code>primaries</code>	The names of several primary volumes for which you want to resume a Synchronous Mirroring operation. Enter the names of the primary volumes using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Separate each of the names with a space.</li> </ul> If the primary volume names have special characters or consist only of numbers, enter the names using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Enclose each of the names in double quotation marks (" ").</li> <li>• Separate each of the names with a space.</li> </ul>
<code>writeConsistency</code>	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 <code>TRUE</code> . For the volumes to be separate, set this parameter to <code>FALSE</code> .

### 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

**Related references**[Show synchronous mirroring volume synchronization progress](#) on page 427[Start synchronous mirroring synchronization](#) on page 462[Suspend synchronous mirroring](#) on page 482**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 technical support.

**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 24 . 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.

**Related references**[Replace drive](#) on page 244[Set drive state](#) on page 315[Set FIPS drive security identifier](#) on page 318[Show drive](#) on page 393

[Start drive initialize](#) on page 449

[Start drive reconstruction](#) on page 451

[Start FDE secure drive erase](#) on page 455

## 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 square brackets ([ ]). If the snapshot group name has special characters or consists only of numbers, you must enclose the snapshot group 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

### Related references

[Create snapshot group](#) on page 146

[Delete snapshot group](#) on page 185

[Set snapshot group attributes](#) on page 331

[Set snapshot group media scan](#) on page 333

[Set snapshot group repository volume capacity](#) on page 334

[Set snapshot group schedule](#) on page 336

[Set SSD cache for a volume](#) on page 342

[Show snapshot group](#) on page 399

[Stop snapshot group pending snapshot images](#) on page 472

## 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 volume name in square brackets ([ ]). If the snapshot volume name has special characters or consists only of numbers, you must 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

## Related references

- [Create read-only snapshot volume](#) on page 145
- [Create snapshot volume](#) on page 156
- [Delete snapshot volume](#) on page 187
- [Rename snapshot volume](#) on page 242
- [Resume snapshot volume](#) on page 259
- [Set read-only snapshot volume to read/write volume](#) on page 328
- [Set snapshot volume media scan](#) on page 339
- [Set snapshot volume repository volume capacity](#) on page 340
- [Set SSD cache for a volume](#) on page 342
- [Show snapshot volumes](#) on page 403
- [Stop snapshot volume](#) on page 474

## 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 technical support.

## Syntax

```
revive volumeGroup [volumeGroupName]
```

## Parameter

Parameter	Description
volumeGroup	The name of the volume group to be set to the Optimal state. Enclose the volume group name 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 saving synchronization statistics. Enclose the asynchronous mirror group name in square brackets ([ ]). If the asynchronous mirror group name has special characters or numbers, you 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\asynchdata.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.
volume	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.  If statistics are collected for more than one volume in an asynchronous mirror group, all the data will be written to the same file.  Enclose the volume name in double quotation marks (" ").

Parameter	Description
sampleType	<p>This parameter is optional. The default value for <code>sampleType</code> is <code>all</code>.</p> <ul style="list-style-type: none"> <li><code>all</code>— Data for all three sample types are collected and written to the same file.</li> <li><code>mostRecent</code> — Statistics are recorded for the most recent 50 resynchronization samples.</li> <li><code>longestSyncTime</code> — Statistics are collected for the most recent 20 longest resynchronization samples.</li> <li><code>errors</code>— Statistics are recorded for the most recent 20 failed resynchronization samples. These samples include a failure code.</li> </ul>
recordLimit	<p>This parameter is optional. The default value for <code>record limit</code> is <code>no limit</code>. 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

### Related references

[Reset asynchronous mirror group statistics](#) on page 246

## Save controller NVSRAM

This command saves a selected set of controller NVSRAM regions to a file.

### 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: <pre>file="C:\Program Files\CLI\logs\nvsramb.txt"</pre> The default name of the file that contains the NVSRAM values is <code>nvsram-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

### Related references

[Show controller NVSRAM](#) on page 390

## Save drive channel fault isolation diagnostic status

**Note:** With firmware version 8.10, the `save driveChannel faultDiagnostics` command is deprecated.

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](#) on page 447 for more information.

### Syntax

```
save driveChannel faultDiagnostics file="filename"
```

**Parameter**

Parameter	Description
file	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:  file="C:\Program Files\CLI\sup\fltdiag.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.

**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 legacy controller trays.

**Related references**

[Start drive channel fault isolation diagnostics](#) on page 447

[Stop drive channel fault isolation diagnostics](#) on page 469

[Start drive channel fault isolation diagnostics](#) on page 447

[Stop drive channel fault isolation diagnostics](#) on page 469

**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	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\lgsendat.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

### Related references

- [Revive drive](#) on page 262
- [Save drive channel fault isolation diagnostic status](#) on page 267
- [Set drive hot spare](#) on page 314
- [Set drive service action allowed indicator](#) on page 314
- [Set drive state](#) on page 315
- [Set FIPS drive security identifier](#) on page 318
- [Set foreign drive to native](#) on page 319
- [Show drive](#) on page 393

## Save input output controller (IOC) dump

This command saves the IOC dumps from the controllers in a storage array to a file on a host.

### Syntax

```
save IOCLog [file="filename"]
```

### Parameter

Parameter	Description
file	<p>The file path and the file name to which you save the IOC dump. Enclose the file name in double quotation marks (" "). For example:</p> <pre>file="C:\Array Backups\IOCCoredump1.7z"</pre> <p>This command saves the data to a compressed file and adds the file extension <code>.7z</code> to the saved file. The default name of the file that contains the IOC-related data for controllers in the storage array uses the WWN of the storage array.</p>

### Notes

This command retrieves from both controllers the IOC dump log data and the IOC dump metadata. Using a 7zip file format, the retrieved data is archived and compressed into a single file with a file name of your choice. The 7zip archive file contains the following:

- filename+"IOCLog"+[A|B].gz – the IOC logs retrieved from controller A or controller B if available
- filename+"IOCLogInfo"+[A|B].txt – the IOC logs metadata info retrieved from controller A or controller B . If the IOC log data cannot be retrieved from a controller, the metadata .txt file will contain the error condition and reason.

Following are the error conditions:

- The controller platform and HIC do not support IOC dump.
- The controllers have not collected IOC dump data.

The compressed logs are not in a human-readable format. You must return the logs to technical support to be evaluated.

### Minimum firmware level

8.20

**Related references**

[Start input output controller \(IOC\) dump](#) on page 452

**Save auto-load balancing statistics**

The Automatic Load Balancing feature provides automated I/O workload balancing and ensures that incoming I/O traffic from the hosts is dynamically managed and balanced across both controllers with SANtricity 11.30 and SANtricity OS (controller software) 8.30 or later.

This command saves the automatic load balancing statistics of the storage array to a file. These statistics show the ownership change history on each volume in the array.

**Note:** Send this file to technical support for interpretation.

**Syntax**

```
save storageArray autoLoadBalanceStatistics file="filename"
(numberOfDays=numberOfDays);
```

**Parameters**

Parameter	Description
storageArray	Specifies that this command works on a storage array.
file	The file path and the file name to which you want to save the auto-load balancing statistics. Enclose the file name in double quotation marks (" ").
numberOfDays	The number of days worth of statistics to save. This parameter is optional and the default value is 0, which indicates all available data.

**Example**

```
SMcli -n Array1 -c "save storageArray autoLoadBalanceStatistics
file="filename" numberOfDays=30;"
```

**Minimum firmware level**

8.30

**Related references**

[Set storage array to enable or disable Automatic Load Balancing](#) on page 349

**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	<p>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:</p> <pre>file="C:\Program Files\CLI\logs\saconf.cfg"</pre> <p>The default name of the file that contains the configuration settings is <code>storage-array-configuration.cfg</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>
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 <code>TRUE</code> .)
globalSettings	The setting to save the global settings to the file. To save the global settings, set this parameter to <code>TRUE</code> . To prevent saving the global settings, set this parameter to <code>FALSE</code> . The default value is <code>TRUE</code> .
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 <code>TRUE</code> . To prevent saving the volume configuration settings and global settings, set this parameter to <code>FALSE</code> . The default value is <code>TRUE</code> .
hostTopology	The setting to save the host topology to the file. To save the host topology, set this parameter to <code>TRUE</code> . To prevent saving the host topology, set this parameter to <code>FALSE</code> . The default value is <code>FALSE</code> .
lunMappings	The setting to save the LUN mapping to the file. To save the LUN mapping, set this parameter to <code>TRUE</code> . To prevent saving the LUN mapping, set this parameter to <code>FALSE</code> . The default value is <code>FALSE</code> .

## 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. You must either use the `allConfig` parameter or one or more of the other four parameters.

## Minimum firmware level

6.10

## Related references

- [Autoconfigure storage array](#) on page 82
- [Autoconfigure storage array hot spares](#) on page 88
- [Clear storage array configuration](#) on page 99
- [Show storage array auto configuration](#) on page 414

[Start storage array configuration database diagnostic](#) on page 459

[Stop storage array configuration database diagnostic](#) on page 478

## Save storage array controller health image

**Note:** With firmware version 8.20 the `coreDump` parameter is removed. Instead, use the `controllerHealthImage` parameter.

This command saves a controller health image to a specified file on a host.

### Syntax

```
save storageArray controllerHealthImage file="filename"
```

### Parameter

Parameter	Description
file	The file path and the file name to which you save the controller health image. Enclose the file name in double quotation marks ( " "). For example:  <pre>file="C:\Array Backups\DBMbackup_03302010.core"</pre>

### Notes

A controller health image file might be one to two gigabytes in size, and take several minutes to save. In addition to saving the controller health image itself, this command generates an XML descriptor file based on the controller health image metadata. This descriptor file is saved in ZIP file format to the same path as the controller health image. The following example show the XML format for the descriptor file.

```
<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
- <DPLcontrollerHealthImageInfo>
<dplcontrollerHealthImageTag>sometag#</dplcontrollerHealthImageTag>
<captureTime>12/22/10 3:58:53 PM IST</captureTime>
- <StorageArrayData>
<ArraySAID>600a0b80006e006a000000004c65efc1</ArraySAID>
<ManagementApiVersion>devmgr.v1083api01.Manager</ManagementApiVersion>
<fwVersion>07.83.01.00</fwVersion>
<platformType>7091</platformType>
</StorageArrayData>
<fullcontrollerHealthImageCtrl>controllerserialNumber1
  </fullcontrollerHealthImageCtrl>
<fullcontrollerHealthImageSize>fullCaptureSize</
fullcontrollerHealthImageSize>
<altcontrollerHealthImageCtrl>controllerserialNumber2</
altcontrollerHealthImage
  Ctrl>
<altcontrollerHealthImageSize>altCaptureSize</
altcontrollerHealthImageSize>
<triggerReason>Exception</triggerReason>-
<DPLcontrollerHealthImageDetail>
<dplcontrollerHealthImageReason>DPLcontrollerHealthImageReason
  </dplcontrollerHealthImageReason>
- <fwLocation >
<filename>filename</filename>
<lineNumber>line#</lineNumber>
</fwLocation >
<panicReason>panicString</panicReason>
</DPLcontrollerHealthImageDetail>
</DPLcontrollerHealthImageInfo>
```



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

### Minimum firmware level

7.83

8.20 replaces the `coreDump` parameter with the `controllerHealthImage` parameter.

### Related references

[Set storage array controller health image allow overwrite](#) on page 348

[Show storage array controller health image](#) on page 419

[Start storage array controller health image](#) on page 460

## 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= 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>
sourceLocation	<p>This parameter specifies the location from which to obtain backup database information.</p> <ul style="list-style-type: none"> <li>disk indicates that data comes directly from the database on the drive</li> <li>onboard indicates that data comes from the RPA memory location</li> </ul> <p>The default location is onboard .</p>
controller	<p>This parameter specifies the controller from which data will be exclusively retrieved, if <code>sourceLocation</code> is set to <code>onboard</code> . If the <code>controller</code> parameter is not specified, data might be retrieved from either controller.</p> <p>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 .</p>

Parameter	Description
contentType	This parameter specifies the content type of the data that will be retrieved. This parameter is set to <code>all</code> by default, so that all of the data including disk pool configuration data is retrieved.

### 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`

### Related references

[Load storage array DBM database](#) on page 223

[Save storage array DBM validator information file](#) on page 274

[Show storage array DBM database](#) on page 420

[Load storage array DBM database](#) on page 223

[Save storage array DBM validator information file](#) on page 274

[Show storage array DBM database](#) on page 420

## Save storage array DBM validator information file

This command saves the database management (DBM) validation information for a storage array in an XML file. Technical support 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.

Parameter	Description
dbmDatabase	<p>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:</p> <pre>dbmDatabase="C:\Array Backups \DBMbackup_03302010.dbm"</pre> <p>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.</p>

## 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>1T0nnnnnnnnnABCD</A>
<B>1T0nnnnnnnnnABCD</B>
</Controllers>
<Header>
<ImageVersion>1</ImageVersion>
<TimeStamp>1269388283</TimeStamp>
</Header>
<Trailer>
<CRC>nnnnnnnnnn</CRC><
/Trailer>
</DbmImageMetadata>
```

## Minimum firmware level

7.75

## Related references

[Load storage array DBM database](#) on page 223

[Save storage array DBM database](#) on page 273

[Show storage array DBM database](#) on page 420

[Load storage array DBM database](#) on page 223

[Save storage array DBM database](#) on page 273

[Show storage array DBM database](#) on page 420

## 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 technical support 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 technical support.

## 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: <pre>file="C:\Program Files\CLI\logs\sadiag.zip"</pre> 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 .zip 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`.

## Related references

[Reset storage array diagnostic data](#) on page 251

## 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 technical support 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.
file	<p>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:</p> <pre>file="C:\Program Files\CLI\logs\events.txt"</pre> <p>The default name of the file that contains the contents of the Major Event Log is <code>major-event-log.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>
count	The number of events that you want to save to a file. If you do not enter a value for the count, all events matching the type of event you specify are saved to the file. If you enter a value for the count, only that number of events (starting with the last event entered) are saved to the file. Use integer values.
forceSave	<p>The parameter to force saving the debug events you may have filtered from view to a file. To force saving the events, set this parameter to <code>TRUE</code>. The default value is <code>FALSE</code>.</p> <p><b>Note:</b> Using this parameter saves debug events that have been filtered out of the SANtricity Event Viewer.</p>

## Minimum firmware level

7.77 add these parameters:

- warningEvents
- infoEvents
- debugEvents
- forceSave

## Related references

[Using curl commands with the E2800 storage array](#) on page 39

[Clear storage array event log](#) on page 101

[Set event alert filtering](#) on page 316

[Show blocked events](#) on page 384

[Test alerts](#) on page 484

## 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)
- Power supplies

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 technical support 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	<p>The statistics that are collected are all statistics from the controller start-of-day. Enclose the parameter in square brackets ([ ]).</p>

Parameter	Description
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\sup\ibstat.csv"</code> The default name of the file that contains the InfiniBand performance statistics is <code>ib-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 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 statistics 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

**Related references**

- [Activate iSCSI initiator](#) on page 72
- [Create iSCSI initiator](#) on page 131
- [Delete iSCSI initiator](#) on page 184
- [Diagnose controller iSCSI host cable](#) on page 193
- [Reset storage array iSCSI baseline](#) on page 252
- [Set iSCSI initiator](#) on page 324
- [Set iSCSI target properties](#) on page 326
- [Show current iSCSI sessions](#) on page 390
- [Show storage array negotiation defaults](#) on page 423
- [Show storage array unconfigured iSCSI initiators](#) on page 424
- [Start iSCSI DHCP refresh](#) on page 453
- [Stop storage array iSCSI session](#) on page 479

**Save storage array iSER statistics**

This command saves the iSER performance statistics of the storage array to a file.

**Syntax**

```
save storageArray iSerStatistics [(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: <pre>file="C:\Program Files\CLI\logs\iserstat.csv"</pre> The default name of the file that contains the iSCSI performance statistics is <code>iser-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 iSER baseline statistics since the controller start-of-day, the time at the start-of-day is the default baseline time.

**Minimum firmware level**

8.20

**Related references**

- [Reset storage array iSER baseline](#) on page 253
- [Set iSER target](#) on page 327



## 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 <code>performanceStatistics.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 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\rlscent.csv"  The default name of the file that contains the RLS counts is <code>read-link-status.csv</code> . You can use any file name, but you must use the <code>.csv</code> 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	<p>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:</p> <pre>file="C:\Program Files\CLI\logs\sasphy.csv"</pre> <p>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.</p>

**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	<p>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:</p> <pre>file="C:\Program Files\CLI\logs\socstat.csv"</pre> <p>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.</p>

**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	<p>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:</p> <pre>file="C:\Program Files\CLI\logs\state.txt"</pre> <p>You can use any file name, but you must use the .txt extension.</p> <p>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.</p>

## Save storage array support data

This command saves the support-related information of the storage array to a file. The following table lists the type of support data that you can collect.

Type of data	File name and description
Automatic Load Balancing statistics	<pre>alb-statistics-A.txt</pre> <pre>alb-statistics-B.txt</pre> <p>This file, one per controller, specify various statistics related to the Automatic Load Balancing feature that allow further analysis via offline analysis tool.</p>
Core dump metadata	<pre>all-coredump.xml</pre> <p>This file contains core dump metadata for the array.</p> <p><b>Note:</b> Starting with the 8.25 release, <code>core-dump-info.xml</code> has been renamed to <code>all-coredump.xml</code>.</p>
AutoSupport status	<pre>autosupport.xml</pre> <p>This file specifies the current status of ASUP for the given array.</p>
The AutoSupport history file	<pre>autosupport-history.xml</pre> <p>This file specifies all ASUP messages, both standard and AutoSupport OnDemand messages, for the given array. Each entry in the table corresponds to an ASUP message currently sitting in one of the repositories for a given array.</p>

Type of data	File name and description
AutoSupport logs	<p>asup-transmission-logs.txt</p> <p>This file contains all the transmitted AutoSupport messages.</p>
Web Server Client Info (E2800 only)	<p>client-info.txt</p> <p>This file specifies various Web Services client information</p> <p><b>Restriction:</b> When you use the RESTful API or SANtricity System Manager to collect the support bundle, this file is present. This file is not included when collecting the support bundle via the CLI.</p>
Drive cabling topology	<p>connection.txt</p> <p>This file contains information about the connections between the drive tray environmental services module (ESM) and the controller pair.</p>
Cumulative statistics bundles	<p>cumulative-drive-vol-stats.xml</p> <p>This file contains cumulative statistics for the controllers.</p>
Misbehaving drive log	<p>controller-drive-error-event-log.txt</p> <p>This file contains the storage array controller drive-side error and event log messages.</p>
Controller health image metadata	<p>core-dump-info.xml</p> <p>This file contains DPL controller health image information.</p> <p><b>Note:</b> Starting with the 8.25 release, core-dump-info.xml has been renamed to all-coredump.xml.</p>
DOM 0 journal files for controller B (E2800 only)	<p>dom0-complete-journal-A.7z</p> <p>This file, one per controller, contains detailed Linux system logs. This includes logs for NetApp Linux packages and standard Linux utilities and services.</p> <p><b>Restriction:</b> When you use the RESTful API or SANtricity System Manager to collect the support bundle, this file is present. This file is not included when collecting the support bundle via the CLI.</p>
DOM 0 journal files for controller B (E2800 only)	<p>dom0-complete-journal-B.7z</p> <p>This file, one per controller, contains detailed Linux system logs. This includes logs for NetApp Linux packages and standard Linux utilities and services.</p> <p><b>Restriction:</b> When you use the RESTful API or SANtricity System Manager to collect the support bundle, this file is present. This file is not included when collecting the support bundle via the CLI.</p>

Type of data	File name and description
Miscellaneous DOM 0 log files for controller A (E2800 only)	<p>dom0-misc-logs-A.7z</p> <p>This file, one per controller, contains system logs that cannot be contained with the Journal. This includes the serial log for the RAID Application and debug logs for the Hypervisor.</p> <p><b>Restriction:</b> When you use the RESTful API or SANtricity System Manager to collect the support bundle, this file is present. This file is not included when collecting the support bundle via the CLI.</p>
Miscellaneous DOM 0 log files for controller B (E2800 only)	<p>dom0-misc-logs-B.7z</p> <p>This file, one per controller, contains system logs that cannot be contained with the Journal. This includes the serial log for the RAID Application and debug logs for the Hypervisor.</p> <p><b>Restriction:</b> When you use the RESTful API or SANtricity System Manager to collect the support bundle, this file is present. This file is not included when collecting the support bundle via the CLI.</p>
Drive command aging timeout values	<p>drive-command-aging-timeout.txt</p> <p>This file contains the default values and current values for the command aging timeout field for every drive.</p>
Drive health logs	<p>drive-health-data.bin</p> <p>This file contains various drive information related to the drive's health.</p> <p><b>Note:</b> This file is a binary file and will require an offline parser to convert to human readable format.</p>
Drive performance analyzer data	<p>drive-performance-log.txt</p> <p>This file contains drive performance data that helps you identify any drives that are performing below expectations.</p>
Enterprise Management Window configuration	<p>emwdata_v04.bin</p> <p>This file contains the EMW configuration data store file.</p> <p><b>Note:</b> In support bundles for the E2800, this file isn't present.</p>
Tray event logs	<p>expansion-tray-log.txt</p> <p>ESM event logs.</p>
Failed repository analysis	<p>failed-repository-analysis.txt</p> <p>This file contains the failed repository analysis information.</p>
Features of the storage array	<p>feature-bundle.txt</p> <p>This file contains a list of the number of volumes, drives, and drive trays allowed in the storage array and a list of the features available and their limits.</p>
Firmware inventory	<p>firmware-inventory.txt</p> <p>This file contains a list of all of the firmware versions for all of the components in the storage array.</p>
Infiniband interface statistics (Infiniband only)	<p>ib-statistics.csv</p> <p>This file contains the Infiniband interface statistics.</p>

Type of data	File name and description
I/O path statistics	<p><code>io-path-statistics.7z</code></p> <p>This file contains raw performance data for each controller that can be used to analyze application performance issues.</p>
IOC dump info for the host interface chip	<p><code>ioc-dump-info.txt</code></p> <p>This file contains IOC dump information for the host interface chip.</p>
IOC dump logs for the host interface chip	<p><code>ioc-dump.gz</code></p> <p>This file contains the log dump from the host interface chip on the controller. The file is compressed in gz format. The zip file is saved as a file inside of the Customer Support Bundle.</p>
iSCSI connections (iSCSI only)	<p><code>iscsi-session-connections.txt</code></p> <p>This file contains a list of all of the current iSCSI sessions.</p>
iSCSI statistics (iSCSI only)	<p><code>iscsi-statistics.csv</code></p> <p>This file contains statistics for the Ethernet media access control (MAC), Ethernet Transmission Control Protocol (TCP)/Internet Protocol (IP), and iSCSI target.</p>
iSER interface statistics (iSER only)	<p><code>iser-statistics.csv</code></p> <p>This file contains the statistics for the host interface card that runs iSER.</p>
Major event log	<p><code>major-event-log.txt</code></p> <p>This file contains a detailed list of events that occur on the storage array. The list is stored in reserved areas on the disks in the storage array. The list records configuration events and component failures in the storage array.</p>
Manifest file	<p><code>manifest.xml</code></p> <p>This file contains a table that describes the files included in the archive file and the collected status of each of those files.</p>
Storage management software runtime information	<p><code>msh-runtime-info.txt</code></p> <p>This file contains the storage management software application runtime information. Contains the JRE version currently used by the storage management software.</p>
NVS RAM data	<p><code>nvsram-data.txt</code></p> <p>This controller file specifies the default settings for the controllers.</p>
Object bundle	<p><code>object-Bundle</code></p> <p>This bundle contains a detailed description of the status of your storage array and its components, which was valid at the time that the file was generated.</p>
Summary performance statistics	<p><code>perf-stat-daily-summary-a.csv</code>  <code>perf-stat-daily-summary-b.csv</code></p> <p>This file contains various controller performance statistics, one file per controller.</p>

Type of data	File name and description
Persistent reservations and registrations	<p><code>persistent-reservations.txt</code></p> <p>This file contains a detailed list of volumes on the storage array with persistent reservations and registrations.</p>
Storage management software user preferences	<p><code>pref-01.bin</code></p> <p>This file contains the user preference persistent data store.</p> <p><b>Note:</b> In support bundles for the E2800, this file isn't present.</p>
Recovery Guru procedures	<p><code>recovery-guru-procedures.html</code></p> <p>This file contains a detailed list of all of the recovery guru topics that are issued in response to problems detected on the storage array.</p>
Recovery profile	<p><code>recovery-profile.csv</code></p> <p>This file contains a detailed description of the latest recovery profile record and historical data.</p>
SAS PHY error logs	<p><code>sas-phy-error-logs.csv</code></p> <p>This file contains the error information for SAS PHY.</p>
State capture data	<p><code>state-capture-data.txt</code></p> <p>This file contains a detailed description of the current state of your storage array.</p>
Storage array configuration	<p><code>storage-array-configuration.cfg</code></p> <p>This file contains a detailed description of the logical configuration on your storage array.</p>
Storage array profile	<p><code>storage-array-profile.txt</code></p> <p>This file contains a description of all of the components and properties of a storage array.</p>
Trace buffer contents	<p><code>trace-buffers.7z</code></p> <p>This file contains the contents of the controllers' trace buffers that are used to record debug information.</p>
Tray capture data	<p><code>tray-component-state-capture.7z</code></p> <p>If your tray contains drawers, the diagnostic data is archived in this zipped file. The Zip file contains a separate text file for each tray that contains drawers. The Zip file is saved as a file inside of the Customer Support Bundle.</p>
Unreadable sectors	<p><code>unreadable-sectors.txt</code></p> <p>This file contains a detailed list of all of the unreadable sectors that have been logged to the storage array.</p>
Web Services Trace Log (E2800 only)	<p><code>web-server-trace-log-A.7z</code> <code>web-server-trace-log-B.7z</code></p> <p>This file, one per controller, contains Web Services trace buffers that are used to record debug information.</p> <p><b>Restriction:</b> When you use the RESTful API or SANtricity System Manager to collect the support bundle, this file is present. This file is not included when collecting the support bundle via the CLI.</p>

Type of data	File name and description
Workload capture analytics log file	wlc-analytics-a.lz4 wlc-analytics-b.lz4  This file, one per controller, contains computed key workload characteristics such as LBA histogram, read/write ratio and I/O throughput across all active volumes.
X-header data file	x-header-data.txt  This AutoSupport message header consists of plain text key-value pairs; which include information about the array and message type.

### Syntax

```
save storageArray supportData file="filename" [force=(TRUE | FALSE)]
```

### Parameters

Parameter	Description
file	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:  file="C:\Program Files\CLI\logs\supdat.7z"

### 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

7.86 requires that the file name uses a .7z extension.

8.30 adds information for the E2800 storage array.

### Related references

[Using curl commands with the E2800 storage array](#) on page 39

[Capture or view an AutoSupport log](#) on page 89

[Display AutoSupport bundle collection settings](#) on page 198

[Display AutoSupport message collection schedule](#) on page 200

[Disable AutoSupport at the EMW level SMcli version](#) on page 197

[Enable or disable AutoSupport at the EMW management domain level \(SMcli only\)](#) on page 210

[Enable or disable the AutoSupport OnDemand feature at the EMW management domain level \(SMcli only\)](#) on page 211

[Enable or disable the AutoSupport OnDemand Remote Diagnostics feature at the EMW management domain level \(SMcli only\)](#) on page 213

[Reset AutoSupport message collection schedule](#) on page 247

[Specify the AutoSupport delivery method](#) on page 438

[Test the AutoSupport configuration](#) on page 485



## 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

### Related references

[Download environmental card firmware](#) on page 204

[Download tray configuration settings](#) on page 207

## 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. You collect system and configuration information in support bundles to transmit the information to technical support for analysis and diagnostics. You can create two types of schedules:

- Daily/Weekly - Use this schedule when you need to frequently collect data to monitor new configurations or troubleshoot an existing configuration.
- Monthly/Yearly - Use this schedule when you are collecting data for evaluating the long term operation of your storage array.

### Syntax for creating a daily/weekly schedule

```
SMcli -supportBundle schedule enable
(all | storageArrayName)
startTime=HH:MM
startDate=MM:DD:YYYY
endDate=MM:DD:YYYY
daysOfWeek=[Sunday Monday Tuesday Wednesday Thursday Friday Saturday]
data=pathName
```

### Syntax for creating a monthly/yearly schedule

```
SMcli -supportBundle schedule enable
(all | storageArrayName)
startTime=HH:MM
startDate=MM:DD:YYYY
endDate=MM:DD:YYYY
months=[January February March April May June July August September
October
November December]
weekNumber=(First|Second|Third|Fourth|Last)
dayOfWeek=(Sunday|Monday|Tuesday|Wednesday|Thursday|Friday|Saturday)|
onDays=[1-31] |
data=pathName
```

### Syntax for removing a schedule

**Note:** This command completely removes the schedule for the storage array. To start saving support bundle information again, you will need to create a new schedule.

```
SMcli -supportBundle schedule disable
(all|storageArrayName)
```

### Parameters

The following table lists the parameters for a daily/weekly schedule.

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.
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.
startTime	The time of a day that you want the collection of a support bundle to start. The format for entering the time is <i>HH:MM</i> , where <i>HH</i> is the hour and <i>MM</i> is the minute past the hour. Use a 24-hour clock. 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.
startDate	A specific date on which you want to start collecting support bundle data. The format for entering the date is <i>MM:DD:YY</i> . The default date is the current date.
endDate	A specific date on which you want to stop collecting support bundle data. The format for entering the date is <i>MM:DD:YY</i> . The default is none. If you use this parameter, it must be at least one week after the start date and cannot occur in the past.

Parameter	Description
daysOfWeek	<p>Specific day or 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 one or more days of week. For example:</p> <pre>daysOfWeek=[Tuesday]</pre> <pre>daysOfWeek=[Monday Wednesday Friday]</pre> <p>Enclose the days of the week in square brackets ([ ]). If you enter more than one day, separate each day with a space.</p>
data	<p>The file path and the file name to which you want to save the support bundle data. For example:</p> <pre>file="C:\Program Files\CLI\sup\data.txt"</pre> <p><b>Note:</b> You can use any file extension.</p> <p>Enclose the file path and name in double quotation marks (" ").</p>

The following table lists the parameters for a monthly/yearly schedule.

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.
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.
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>

Parameter	Description
months	<p>Specific month or months of the year on which you want to collect support bundle data. Use this parameter when you want to collect support bundle data on one or more months of a year. For example:</p> <pre>months=[June]</pre> <pre>months=[January April July October]</pre> <p>Enclose the month in square brackets ([ ]). If you enter more than one month, separate each month with a space.</p>
weekNumber	<p>A week in the month during which you want to collect support bundle data. For example:</p> <pre>weekNumber=first</pre>
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:</p> <pre>dayOfWeek=Wednesday</pre>
onDays	<p>Specific day or days in a month on which you want to collect support bundle data. For example:</p> <pre>onDays=[15]</pre> <pre>onDays=[7 21]</pre> <p>Enclose the day in square brackets ([ ]). If you enter more than one day, separate each day with a space.</p> <p><b>Note:</b> You cannot use the <code>onDays</code> parameter with either the <code>weekNumber</code> parameter or the <code>dayOfWeek</code> parameter.</p>
data	<p>The file path and the file name to which you want to save the support bundle data. For example:</p> <pre>file="C:\Program Files\CLI\sup\data.txt"</pre> <p><b>Note:</b> You can use any file extension.</p> <p>Enclose the file path and name in double quotation marks (" ").</p>

The following table lists the parameters for removing a schedule.

Parameter	Description
disable	<p>Disables the automatic collection of support bundles and deletes any previously defined schedules immediately.</p> <p><b>Note:</b> Disabling a schedule also deletes the schedule.</p>

Parameter	Description
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.

## 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.

Following are examples of the usage of this command. The first example is of a daily/weekly schedule that meets these criteria for collecting support bundle data:

- The storage array name is `DevKit4`
- The collection start time is `02:00` (2:00 in the morning)
- The start date is `05:01:2013` (May 1, 2013)
- Data will be collected on Monday and Friday of each week
- This schedule does not have an end date, and can be stopped only by running the `SMcli - supportBundle schedule disable` command

```
SMcli -supportBundle schedule enable DevKit4 startTime=02:00
startDate=05:01:2013 endDate=05:10:2014 daysOfWeek=[Monday Friday]
```

The second example is of a monthly/yearly schedule that meets these criteria for collecting support bundle data:

- The storage array name is `eng_stor1`
- The collection start time is `04:00` (4:00 in the morning)
- The start date is `05:01:2013` (May 1, 2013)
- Data will be collected in March, April, May, June, and August
- Data will be collected on the first and twenty-first days of each month
- This schedule does not have an end date, and can be stopped only by running the `SMcli - supportBundle schedule disable` command

```
SMcli -supportBundle schedule enable eng_stor1 startTime=04:00
startDate=05:01:2013 months=[March April May June August] onDays=[1 21]
```

The third example is of a monthly/yearly schedule that meets these criteria for collecting support bundle data:

- The storage array name is `firmware_2`
- The collection start time is `22:00` (10:00 at night)
- The start date is `05:01:2013` (May 1, 2013)
- Data will be collected in March, April, May, June, and August

- Data will be collected on Friday of the first week of each month
- This schedule will end on 05:10:2014 (May 10, 2014)

```
SMcli -supportBundle schedule enable firmware_2 startTime=22:00
startDate=05:01:2013 endDate=05:10:2014 months=[March April May June
August]
weekNumber=First dayOfWeek=Friday
```

### 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]
[volume="repos_xxxx" increaseRepositoryCapacity
(repositoryVolume="repos_xxxx")]
[syncInterval=integer (minutes | hours | days)]
[warningSyncThreshold=integer (minutes | hours | days)]
[warningRecoveryThreshold=integer (minutes | hours | days)]
[userLabel="New_asyncMirrorGroupName"]
[warningThresholdPercent=percentValue]
[role=(primary | secondary)]
[autoResync=(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 or numbers, you must enclose the name in double quotation marks (" ") inside square brackets.
volume	The name of an asynchronous mirror group repository volume for which you want to increase the capacity. The name of a repository volume is comprised of two parts: <ul style="list-style-type: none"> <li>• The term repos</li> <li>• A four digit numerical identifier that the storage management software assigns to the repository volume name</li> </ul> Enclose the name of the repository volume in double quotation marks (" ").

Parameter	Description
<code>syncInterval</code>	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.
<code>warningSyncThreshold</code>	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.
<code>warningRecoveryThreshold</code>	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.  <b>Note:</b> You must set the Recovery Point Threshold to be twice as large as the synchronization interval threshold.
<code>userLabel</code>	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 (" ").
<code>warningThresholdPercent</code>	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.
<code>role</code>	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> .
<code>autoResync</code>	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: <ul style="list-style-type: none"> <li>• <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.</li> <li>• <code>disabled</code> – Automatic resynchronization is turned off. To resynchronize the primary volume and the secondary volume, you must run the <code>start asyncMirrorGroup</code> command.</li> </ul>

## 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.

## Minimum firmware level

7.84

## Related references

- [Add volume to asynchronous mirror group](#) on page 80
- [Cancel asynchronous mirror group role reversal](#) on page 88
- [Check asynchronous mirror group consistency](#) on page 91
- [Create asynchronous mirror group](#) on page 104
- [Delete asynchronous mirror group](#) on page 177
- [Remove volume from asynchronous mirror group](#) on page 239
- [Reset asynchronous mirror group statistics](#) on page 246
- [Resume asynchronous mirror group](#) on page 255
- [Save asynchronous mirror group statistics](#) on page 265
- [Show asynchronous mirror groups](#) on page 382
- [Show asynchronous mirror group synchronization progress](#) on page 383
- [Start asynchronous mirroring synchronization](#) on page 440
- [Suspend asynchronous mirror group](#) on page 481
- [Test asynchronous mirror group connectivity](#) on page 484

## 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
consistencyGroupName	The name of the consistency group for which you are setting properties. Enclose the consistency group name in double quotation marks (" ") inside square brackets ([ ]).



Parameter	Description
<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

**Related references**

[Add member to consistency group](#) on page 78  
[Create consistency group](#) on page 107  
[Create consistency group snapshot image](#) on page 115  
[Create consistency group snapshot volume](#) on page 116  
[Create consistency group snapshot volume mapping](#) on page 120  
[Delete consistency group](#) on page 178  
[Delete consistency group snapshot image](#) on page 179  
[Delete consistency group snapshot volume](#) on page 180  
[Remove member volume from consistency group](#) on page 237  
[Resume consistency group snapshot volume](#) on page 256  
[Set consistency group snapshot volume](#) on page 298  
[Show consistency group](#) on page 384  
[Show consistency group snapshot image](#) on page 386  
[Start consistency group snapshot rollback](#) on page 443  
[Stop consistency group snapshot rollback](#) on page 466  
[Stop consistency group snapshot volume](#) on page 468  
[Stop pending snapshot images on consistency group](#) on page 471

**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"
```

**Parameters**

Parameter	Description
cgSnapVolume	The name of the consistency group volume that you want to rename. Enclose the name of the consistency group snapshot volume in square brackets ([ ]).
userLabel	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

**Related references**

[Add member to consistency group](#) on page 78  
[Create consistency group](#) on page 107  
[Create consistency group snapshot image](#) on page 115

[Create consistency group snapshot volume](#) on page 116  
[Create consistency group snapshot volume mapping](#) on page 120  
[Delete consistency group](#) on page 178  
[Delete consistency group snapshot image](#) on page 179  
[Delete consistency group snapshot volume](#) on page 180  
[Remove member volume from consistency group](#) on page 237  
[Resume consistency group snapshot volume](#) on page 256  
[Set consistency group attributes](#) on page 296  
[Show consistency group](#) on page 384  
[Show consistency group snapshot image](#) on page 386  
[Start consistency group snapshot rollback](#) on page 443  
[Stop consistency group snapshot rollback](#) on page 466  
[Stop consistency group snapshot volume](#) on page 468  
[Stop pending snapshot images on consistency group](#) on page 471

## 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 [portLabel] 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.

Parameter	Description
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 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 <b>(0–255).(0–255).(0–255).(0–255)</b>
IPv6RouterAddress	The IP address of IPv6 router that connects two or more logical subnets. The address format for the IPv6 router is <b>(0–FFFF):(0–FFFF):(0–FFFF):(0–FFFF):(0–FFFF):(0–FFFF):(0–FFFF):(0–FFFF)</b> .
iscsiHostPort	This parameter enables you to set options for the iSCSI ports on the controller. Enter the iSCSI port label or number and then select the options for that port. See the following for more information: <i>"Identifying an iSCSI host port label"</i>
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 <code>TRUE</code> . To turn off the remote login feature, set this parameter to <code>FALSE</code> .
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 <code>on</code> . To turn off the Service Action Allowed indicator light, set this parameter to <code>off</code> .

### Identifying an iSCSI host port label

You must specify a label for the host port. Follow these steps to specify the host port label:

1. If you do not know the port label for the iSCSI host port, run the `show controller` command.
2. In the Host interface section of the results, locate the host port you want to select.  
**Note:** The port label is the complete value returned for the `Port` field.
3. Enclose the entire value of the port label in both quotes and square brackets: `["portLabel"]`. For example, if the port label is `ch 2`, specify the iSCSI host port as follows:

```
iscsiHostPort["ch 2"]
```

**Note:** If you are using a Windows command line and the label contains a pipe (`|`), the character should be escaped (using `^`); otherwise, it will be interpreted as a command. For example, if the port label is `e0b|0b`, specify the iSCSI host port as follows:

```
iscsiHostPort["e0b^|0b"]
```

**Note:** For backward compatibility, the `iscsiPortNumber`, enclosed by braces `[ ]` rather than quotes and braces `" "` can still be used for E2700, E5600, or EF560 controllers (and other previous

generations of E-Series or EF-Series controllers). For those controllers, valid values for `iscsiPortNumber` are as follows:

- For controllers with integrated host ports, the numbering is 3, 4, 5, or 6.
- For controllers with host ports on a host interface card only, the numbering is 1, 2, 3, or 4.

An example of the prior syntax is as follows:

```
iscsiHostPort[3]
```

### 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 technical support 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 removes the `bootp` parameter, and adds the new Ethernet port options and the new iSCSI host port options.

7.50 moves 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.

8.10 revises the identification method for iSCSI host ports.

#### Related references

[Show controller](#) on page 387

## Set controller DNS settings

The `set controller DNSServers` command updates the Domain Name System (DNS) settings for a controller. DNS is used to resolve fully qualified domain names for the controllers and a Network Time Protocol (NTP) server.

#### Before you begin

Make sure that:

- A DNS server is configured.
- You know the addresses of a primary DNS server and a backup DNS server, in case the primary DNS server fails. These addresses can be IPv4 addresses or IPv6 addresses.

#### Context

**Important:** You need to send this command to both controllers. This command is controller-specific.

**Note:** If you already configured the array's management ports with DHCP (Dynamic Host Configuration Protocol), and you have one or more DNS or NTP servers associated with the DHCP setup, then you do not need to manually configure DNS/NTP. In this case, the array should have already obtained the DNS/NTP server addresses automatically.

#### Syntax

```
set controller[a] DNSServers=(auto|(Address1 [Address2]))
```

#### Parameters

Parameter	Description
DNSServers	<p>This parameter configures the DNS servers for the controller. Specify <code>auto</code> to use the DNS servers provided by DHCP. Otherwise, use a whitespace separated list of one or two DNS servers.</p> <p><b>Note:</b> The management Ethernet ports on the array can support IPv4 or IPv6 protocols simultaneously. The addresses may be an IPv4 address or an IPv6 address. If specifying more than one address, the address types do not need to match.</p> <p>If more than one address is specified, they will be used in the order they are specified (first is primary, second is backup).</p>
AddressX	IPv4Address IPv6Address

#### Examples

```
set controller[a] DNSServers=auto;
set controller[a] DNSServers=(192.168.1.1);
```



```
set controller[b] DNSServers=(192.168.1.1 192.168.1.105);
set controller[b] DNSServers=(2001:0db8:85a3:0000:0000:8a2e:0370:7334
192.168.1.1)
```

### Minimum firmware level

8.25

### Related references

[Set controller NTP settings](#) on page 305

## Set controller NTP settings

The `set controller NTPServers` command sets the NTP settings for the controller so the controller can automatically synchronize the clocks with an external host using SNTP (Simple Network Time Protocol). When this feature is enabled, the controller periodically queries the configured NTP server, then uses the results to update its internal time-of-day clock. If only one controller has NTP enabled, the alternate controller will periodically synchronize its clock with the controller that has NTP enabled. If neither controller has NTP enabled, the controllers periodically synchronize their clocks with each other.

### Context

**Note:** This command is controller-specific. It is not necessary to configure NTP on both controllers; however, doing so improves the array's ability to stay synchronized during hardware or communication failures.

**Important:** If you configure NTP using a fully qualified domain name, you must also configure DNS on those controller(s) for the operation to succeed.

See [Set controller DNS Settings](#) on page 304 for more information.

### Syntax

```
set controller(a|b) NTPServers=(disabled|auto|(Address1 [Address2]))
```

### Parameters

Parameter	Description
controller	The controller for which you want to change the NTP settings. Valid controller identifiers are a or b.

Parameter	Description
NTPServers	<p>This parameter configures the NTP servers for the controller.</p> <ul style="list-style-type: none"> <li>Specify <code>disabled</code> to turn off NTP support.</li> <li>Specify <code>auto</code> to use the NTP server addresses provided by a DHCP server.</li> </ul> <p><b>Note:</b> This option should only be used if at least one management port on the controller is set to get interface parameters via DHCP, and at least one NTP server is configured on the DHCP server.</p> <ul style="list-style-type: none"> <li>Otherwise, specify a whitespace separated list of one or two NTP servers. The addresses may be a domain name, IPv4 address, or IPv6 address. If specifying more than one address, the address types do not need to match. If more than one address is specified, they will be used in the order they are specified (first is primary, second is backup).</li> </ul> <p><b>Note:</b> Surround the NTP server name with quotes, as shown in the Examples section.</p>
Address	<p>"DomainName" IPv4Address IPv6Address</p> <p><b>Note:</b> If a domain name is specified, a DNS server must also be configured to allow the controller to resolve the server IP address.</p>

### Examples

```
set controller[a] NTPServers=disabled;
set controller[a] NTPServers=auto;
set controller[a] NTPServers=("0.pool.ntp.org" "1.us.pool.ntp.org");
set controller[b] NTPServers=(192.168.1.105 "1.us.pool.ntp.org");
set controller[b] NTPServers=(2001:0db8:85a3:0000:0000:8a2e:0370:7334);
```

### Minimum firmware level

8.25

### Related references

[Set controller DNS settings](#) on page 304

## 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.)

**Note:** This command is valid only for E2700, E5600, or EF560 (and other previous generations of E-Series or EF-Series controllers). The E2800 controller tray has a single Attention indicator that will illuminate only when 1) there is a failure, and 2) the component that failed can be safely removed.

## Syntax

```
set controller [(a| b)]
serviceAllowedIndicator=(on | off)
```

## Parameters

Parameter	Description
controller	The controller that has the Service Action Allowed indicator light that you want to turn on or turn off. 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 controller firmware returns a syntax error.
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 on . To turn off the Service Action Allowed indicator light, set this parameter to off .

## Minimum firmware level

6.14

## Related references

- [Clear drive channel statistics](#) on page 98
- [Diagnose controller iSCSI host cable](#) on page 193
- [Enable controller data transfer](#) on page 208
- [Reset controller](#) on page 248
- [Save controller NVSRAM](#) on page 267
- [Save drive channel fault isolation diagnostic status](#) on page 267
- [Set controller](#) on page 299
- [Set controller DNS settings](#) on page 304
- [Set controller NTP settings](#) on page 305
- [Set drive channel status](#) on page 313
- [Set host channel](#) on page 321
- [Show controller](#) on page 387
- [Show controller NVSRAM](#) on page 390
- [Show drive channel statistics](#) on page 395
- [Start controller trace](#) on page 444
- [Start drive channel fault isolation diagnostics](#) on page 447
- [Start drive channel locate](#) on page 449
- [Start iSCSI DHCP refresh](#) on page 453
- [Stop drive channel fault isolation diagnostics](#) on page 469
- [Stop drive channel locate](#) on page 470

## Set disk pool

This command sets the attributes associated with a disk pool based on the specified parameters.

## Syntax

```
set (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>	The name of the disk pool for which you are setting attributes. Enclose the disk pool name in square brackets ([ ]). If the disk pool name has special characters or consists only of numbers, you must enclose the disk pool name in double quotation marks (" ") inside square brackets.
<code>diskPools</code>	The names of several disk pools for which you want to set attributes. Enter the names of the disk pools using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Separate each of the names with a space.</li> </ul> If the disk pool names have special characters or numbers, enter the names using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Enclose each of the names in double quotation marks (" ").</li> <li>• Separate each of the names with a space.</li> </ul>
<code>allDiskPools</code>	This parameter sets attributes for all of the disk pools in the storage array.
<code>reservedDriveCount</code>	This parameter reserves space on every drive in the disk pool, to be used exclusively for reconstruction of failed drives. Each unit in this count represents the capacity to reconstruct one failed drive to the remainder of the disk pool.
<code>warningThreshold</code>	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. Valid values are from 0 to 100. The default value is 50. Setting this parameter to 0 disables warning alerts. If you set this to <code>default</code> , the warning alert threshold value is determined by the controller firmware.

Parameter	Description
<code>criticalThreshold</code>	<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>
<code>criticalPriority</code>	<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>.</p> <p>The default value is <code>highest</code>.</p>
<code>degradedPriority</code>	<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>.</p> <p>The default value is <code>high</code>.</p>
<code>backgroundPriority</code>	<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>.</p> <p>The default value is <code>low</code>.</p>
<code>userLabel</code>	<p>The new name that you want to give the disk pool. Enclose the disk pool name 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 Capacity Reduction (DCR)
- 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/DCR

### Minimum firmware level

7.83

### Related references

[Set disk pool \(modify disk pool\)](#) on page 310

[Show disk pool](#) on page 392

## Set disk pool (modify disk pool)

This command adds capacity to a disk pool (Dynamic Capacity Expansion or DCE) 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
<code>diskPool</code>	The name of the disk pool that you want to modify. Enclose the disk pool name in square brackets ( [ ] ). If the disk pool name has special characters or consists only of numbers, you must enclose the name in double quotation marks ( " " ) inside square brackets.
<code>addDrives</code>	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 24 . Enclose the tray ID value, the drawer ID value, and the slot ID value in square brackets ( [ ] ).  <b>Note:</b> This feature has a 12-drive limit.
<code>addCapacity</code>	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 <code>bytes</code> , <code>KB</code> , <code>MB</code> , <code>GB</code> , or <code>TB</code> .
<code>owner</code>	The controller that owns the disk pool. 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.

## 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

## Related references

[Set disk pool](#) on page 307

[Show disk pool](#) on page 392

## 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 ([ ]).
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> .  For information about using <code>forceOnWarning</code> , see the Notes.

### 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

### Related references

[Set tray identification](#) on page 364

[Set tray service action allowed indicator](#) on page 365

## 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
<code>driveChannel</code>	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 ([ ]).
<code>status</code>	The condition of the drive channel. You can set the drive channel status to <code>optimal</code> or <code>degraded</code> .

### 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.

### Related references

[Clear drive channel statistics](#) on page 98

[Show drive channel statistics](#) on page 395

[Start drive channel fault isolation diagnostics](#) on page 447

[Start drive channel locate](#) on page 449

[Stop drive channel fault isolation diagnostics](#) on page 469

[Stop drive channel locate](#) on page 470

## Set drive hot spare

This command assigns or deassigns one or more drives as a hot spare. Only drives in an Unassigned state can be assigned as a hot spare. After deassigning a hot spare drive, the drive returns to an Unassigned state.

### 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 24 . Enclose the tray ID value, the drawer ID value, and the slot ID value in square brackets ([ ]).
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.

### Related references

[Set drive service action allowed indicator](#) on page 314

[Set drive state](#) on page 315

## 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
drive or drives	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 24 . Enclose the tray ID value, the drawer ID value, and the slot ID value in square brackets ([ ]).
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.

## Related references

[Set drive hot spare](#) on page 314

[Set drive state](#) on page 315

[Start drive locate](#) on page 450

[Stop drive locate](#) on page 470

## 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.

Optionally, you can copy the data from the drive before you fail it.

### Syntax

```
set drive [trayID,drawerID,slotID]
operationalState=failed
[copyDrive]
```

### Parameter

Parameter	Description
drive	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 24 . Enclose the tray ID value, the drawer ID value, and the slot ID value in square brackets ([ ]).
copyDrive	This parameter copies the data from the drive that you are setting to the Failed state. The drive is failed after the data is copied. This parameter is only applicable to assigned drives.  <b>Note:</b> The data is copied to the Global Hot Spare in a volume group, or redistributed to the entire disk pool in a disk pool environment.

### 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.

8.20 adds the `copyDrive` parameter.

### Related references

[Set drive hot spare](#) on page 314

[Set drive service action allowed indicator](#) on page 314

[Replace drive](#) on page 244

[Reset drive](#) on page 249

[Stop drive replace](#) on page 471

## Set event alert filtering

The `set event alert` commands manage alert event notification by either disabling or enabling notification related to a specific alertable event. To prevent notification about a specific alertable event, you *block* it. To enable notification about a specific alertable event, you *unblock* it.

**Restriction:** This command is obsolete for the E2800 storage array. Use the RESTful API, SANtricity System Manager, or see [Using curl commands with the E2800 storage array](#) on page 39.

When you define your storage array, you can configure alerts and define how the event alerts are managed. If you have configured the storage array to send alerts, a notification is sent to a designated recipient when an alertable event occurs. That notification can be one or all of the following types:

- email
- syslog
- SNMP trap notifications

The set event alert CLI commands work on a single storage array. When you run the commands on a storage array, only that storage array is affected by the commands. Other storage arrays that have not had the CLI command run against them have the default behavior.

**Note:** Blocking an event alert does not prevent the event from being posted to the system event log. All events continue to be posted to the event log.

**Note:** Non-alertable events cannot be made alertable using this command.

### Syntax to block an event alert

```
set blockEventAlert eventType
```

### Syntax to unblock an event alert

```
set unBlockEventAlert eventType
```

### Parameters

Parameter	Description
<i>eventType</i>	This parameter is the integer value for the event. Enter the event value in a hexadecimal format, for example, 0x280D. Always begin the hexadecimal value with <b>0x</b> to indicate that it is in hexadecimal format. If you do not use 0x, the value is interpreted as a decimal and converted to a hexadecimal value, before applying the block or unblock command. This can cause an incorrect event to be blocked or unblocked.  An error is displayed if you enter an invalid event.

### Minimum firmware level

8.10

### Related references

[Using curl commands with the E2800 storage array](#) on page 39

[Show blocked events](#) on page 384

[Test alerts](#) on page 484

[Clear storage array event log](#) on page 101

[Save storage array events](#) on page 276

## Set FIPS drive security identifier

This command is used to reset a FIPS drive back to the original manufacturer settings. This may be necessary if the drive is locked due to an invalid or missing security key file or unknown pass phrase. All existing drive data will be erased.

**Warning:** This command resets a FIPS drive back to the original manufacturer settings and erases all existing drive data. This operation is completely irreversible.

To erase an FDE drive, use the `start secureErase` command.

### Syntax

```
set drive [trayID,drawerID,slotID] securityID="string"
```

### Parameters

Parameter	Description
drive	The location of the drive that you want to set the security ID. 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 24 . Enclose the tray ID value, the drawer ID value, and the slot ID value in square brackets ([ ]).
securityID	The security ID of the drive to erase, in string form. This string can have a maximum of 32 characters. The form of the security ID will vary by manufacturer.  <b>Note:</b> To find the security ID, remove the drive and read the security ID on the canister label.

### 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.

### Example

```
set drive [1,31,4] securityID="V2zdvtLTdGwdWyTALGHns";
```

### Minimum firmware level

8.25

### Related references

[Start FDE secure drive erase](#) on page 455

[Start FDE secure drive erase](#) on page 455

## 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
drive or drives	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 24 . Enclose the tray ID value, the drawer ID value, and the slot ID value in square brackets ([ ]).
allDrives	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.

### Related references

[Download drive firmware](#) on page 202

[Replace drive](#) on page 244

[Revive drive](#) on page 262  
[Save drive channel fault isolation diagnostic status](#) on page 267  
[Save drive log](#) on page 268  
[Set drive hot spare](#) on page 314  
[Set drive service action allowed indicator](#) on page 314  
[Set drive state](#) on page 315  
[Set FIPS drive security identifier](#) on page 318  
[Show drive](#) on page 393  
[Show drive download progress](#) on page 396  
[Show drive performance statistics](#) on page 396  
[Start drive channel fault isolation diagnostics](#) on page 447  
[Start drive initialize](#) on page 449  
[Start drive locate](#) on page 450  
[Start drive reconstruction](#) on page 451  
[Start FDE secure drive erase](#) on page 455  
[Stop drive channel fault isolation diagnostics](#) on page 469  
[Stop drive locate](#) on page 470

## 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
host	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 or numbers, you must enclose the host name in double quotation marks (" ") inside square brackets.
hostGroup	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 defaultGroup option is the host group that contains the host to which the volume is mapped.
userLabel	The new host name. Enclose the host name in double quotation marks (" ").
hostType	The index label or number of the host type for the host port. Use the show storageArray hostTypeTable 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

### Related references

- [Create host](#) on page 127
- [Create host group](#) on page 128
- [Create host port](#) on page 129
- [Delete host](#) on page 182
- [Delete host group](#) on page 182
- [Delete host port](#) on page 183
- [Set host channel](#) on page 321
- [Set host group](#) on page 322
- [Set host port label](#) on page 323
- [Show host ports](#) on page 398
- [Show storage array host topology](#) on page 421

## 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.

**Related references**[.save drive channel fault isolation diagnostic status](#) on page 267[.set controller](#) on page 299[.set controller DNS settings](#) on page 304[.set controller NTP settings](#) on page 305[.set controller service action allowed indicator](#) on page 306[.set drive channel status](#) on page 313[.show controller](#) on page 387[.show controller NVSRAM](#) on page 390[.show drive channel statistics](#) on page 395[.start controller trace](#) on page 444[.start drive channel fault isolation diagnostics](#) on page 447[.start drive channel locate](#) on page 449[.start iSCSI DHCP refresh](#) on page 453[.stop drive channel fault isolation diagnostics](#) on page 469[.stop drive channel locate](#) on page 470[.create host](#) on page 127[.create host group](#) on page 128[.create host port](#) on page 129[.delete host](#) on page 182[.delete host group](#) on page 182[.delete host port](#) on page 183[.set host](#) on page 320[.set host group](#) on page 322[.set host port label](#) on page 323[.show host ports](#) on page 398[.show storage array host topology](#) on page 421**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 or numbers, you must enclose the host group name in double quotation marks (" ") inside square brackets.
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

## Related references

- [Create host](#) on page 127
- [Create host group](#) on page 128
- [Create host port](#) on page 129
- [Delete host](#) on page 182
- [Delete host group](#) on page 182
- [Delete host port](#) on page 183
- [Set host](#) on page 320
- [Set host channel](#) on page 321
- [Set host port label](#) on page 323
- [Show host ports](#) on page 398
- [Show storage array host topology](#) on page 421

## Set host port label

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] 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 name has special characters or numbers, enclose the host port name in double quotation marks (" ") inside square brackets.
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**

You can use any combination of alphanumeric characters, hyphens, and underscores for the user label. User labels can have a maximum of 30 characters.

**Minimum firmware level**

6.10

**Related references**

- [Create host](#) on page 127
- [Create host group](#) on page 128
- [Create host port](#) on page 129
- [Delete host](#) on page 182
- [Delete host group](#) on page 182
- [Delete host port](#) on page 183
- [Set host](#) on page 320
- [Set host channel](#) on page 321
- [Set host group](#) on page 322
- [Show host ports](#) on page 398
- [Show storage array host topology](#) on page 421

**Set iSCSI initiator**

This command sets the attributes for an iSCSI initiator.

**Syntax**

```
set iscsiInitiator (["iscsiID"] | <"wwID">)
  (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 (" ") inside square brackets ([ ]).

Parameter	Description
<code>iscsiInitiator</code>	The World Wide Identifier (WWID) of the initiator for which you want to set attributes. Enclose the WWID in double quotation marks (" ") inside angle brackets (<>).
<code>userLabel</code>	The new name that you want to use for the iSCSI initiator. Enclose the new name in double quotation marks (" ").
<code>host</code>	The name of the new host to which the host port is connected. Enclose the host name in double quotation marks (" ").
<code>chapSecret</code>	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

## Related references

[Activate iSCSI initiator](#) on page 72

[Create iSCSI initiator](#) on page 131

[Delete iSCSI initiator](#) on page 184

[Diagnose controller iSCSI host cable](#) on page 193

[Reset storage array iSCSI baseline](#) on page 252

[Save storage array iSCSI statistics](#) on page 279

[Set iSCSI target properties](#) on page 326

[Show current iSCSI sessions](#) on page 390

[Show storage array negotiation defaults](#) on page 423

[Show storage array unconfigured iSCSI initiators](#) on page 424

[Start iSCSI DHCP refresh](#) on page 453

[Stop storage array iSCSI session](#) on page 479

## 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
iscsiTarget	The iSCSI target for which you want to set properties. Enclose the <i>userLabel</i> in double quotation marks (" "). You must also enclose the <i>userLabel</i> in square brackets ( [ ] ) if the user label is a target alias or angle brackets (< >) if the user label is an iSCSI Qualified Name (IQN).
authenticationMethod	The means of authenticating your iSCSI session.
chapSecret	The security key that you want to use to authenticate a peer connection.
targetAlias	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

## Related references

- [Activate iSCSI initiator](#) on page 72
- [Create iSCSI initiator](#) on page 131
- [Delete iSCSI initiator](#) on page 184
- [Diagnose controller iSCSI host cable](#) on page 193
- [Reset storage array iSCSI baseline](#) on page 252
- [Save storage array iSCSI statistics](#) on page 279
- [Set iSCSI initiator](#) on page 324
- [Show current iSCSI sessions](#) on page 390
- [Show storage array negotiation defaults](#) on page 423
- [Show storage array unconfigured iSCSI initiators](#) on page 424
- [Start iSCSI DHCP refresh](#) on page 453
- [Stop storage array iSCSI session](#) on page 479

## Set iSER target

This command defines properties for an iSER target.

### Syntax

```
set iserTarget ["userLabel"]
authenticationMethod=(none | chap) |
chapSecret=securityKey |
targetAlias="userLabel"
```

### Parameters

Parameter	Description
iserTarget	The iSER 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 square brackets ([ ]) if the user label is a target alias or angle brackets (< >) if the user label is an iSCSI Qualified Name (IQN).
authenticationMethod	The means of authenticating your iSCSI session.
chapSecret	The security key that you want to use to authenticate a peer connection.
targetAlias	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

8.20

### Related references

[Reset storage array iSER baseline](#) on page 253

[Save storage array iSER statistics](#) on page 280

## 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
snapVolume	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"> <li>• Use an existing repository volume: name</li> <li>• Create a new repository volume when you run this command</li> </ul> <p>The name of an existing repository volume is comprised of two parts:</p> <ul style="list-style-type: none"> <li>• The term repos</li> <li>• A four digit numerical identifier that the storage management software assigns to the repository volume name</li> </ul> <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"> <li>• An integer value that represents a percentage of the base volume capacity</li> <li>• A decimal fraction value that represents a percentage of the base volume capacity</li> <li>• A specific size for the repository volume. Size is defined in units of bytes , KB , MB , GB , or TB .</li> </ul> <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

### Related references

- [Create read-only snapshot volume](#) on page 145
- [Create snapshot volume](#) on page 156
- [Delete snapshot volume](#) on page 187
- [Rename snapshot volume](#) on page 242
- [Resume snapshot volume](#) on page 259
- [Revive snapshot volume](#) on page 263
- [Set snapshot volume media scan](#) on page 339
- [Set snapshot volume repository volume capacity](#) on page 340
- [Set SSD cache for a volume](#) on page 342
- [Show snapshot volumes](#) on page 403
- [Stop snapshot volume](#) on page 474

## 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
errorAction	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.)

Parameter	Description
password	The password for the storage array. Enclose the password in double quotation marks (" ").
userRole	<p>Defines the user role for the password. The user role can be either:</p> <ul style="list-style-type: none"> <li>admin – Enables you to view and modify storage array configurations</li> <li>monitor – Enables you to view storage array configurations and monitor storage array health conditions, but not modify the configuration</li> </ul> <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>
performanceMonitorInterval	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.
performanceMonitorIterations	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 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 (purge) the snapshot images ( <code>purgeSnapImages</code> ). The default action is <code>purgeSnapImages</code> .
<code>repositoryFullLimit</code>	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.
<code>autoDeleteLimit</code>	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.
<code>rollbackPriority</code>	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

**Related references**

[Create snapshot group](#) on page 146

[Delete snapshot group](#) on page 185

[Revive snapshot group](#) on page 263

[Set snapshot group media scan](#) on page 333

[Set snapshot group repository volume capacity](#) on page 334

[Set snapshot group schedule](#) on page 336

[Set SSD cache for a volume](#) on page 342

[Show snapshot group](#) on page 399

[Stop snapshot group pending snapshot images](#) on page 472

## 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 ([ ]).
mediaScanEnabled	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.)
redundancyCheckEnabled	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> .

### 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

### Related references

[Create snapshot group](#) on page 146

[Delete snapshot group](#) on page 185

[Revive snapshot group](#) on page 263

[Set snapshot group attributes](#) on page 331

[Set snapshot group repository volume capacity](#) on page 334

[Set snapshot group schedule](#) on page 336

[Set SSD cache for a volume](#) on page 342

[Show snapshot group](#) on page 399

[Stop snapshot group pending snapshot images](#) on page 472

## 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 for which you want to increase capacity or decrease capacity. 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"> <li>• Use an existing repository volume: name</li> <li>• Create a new repository volume when you run this command</li> </ul> <p>The name of an existing repository volume consists of two parts:</p> <ul style="list-style-type: none"> <li>• The term repos</li> <li>• A four digit numerical identifier that the storage management software assigns to the repository volume name</li> </ul> <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"> <li>• An integer value that represents a percentage of the base volume capacity</li> <li>• A decimal fraction value that represents a percentage of the base volume capacity</li> <li>• A specific size for the repository volume. Size is defined in units of bytes , KB , MB , GB , or TB .</li> </ul> <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

### Related references

[Create snapshot group](#) on page 146

[Delete snapshot group](#) on page 185

[Revive snapshot group](#) on page 263

[Set snapshot group attributes](#) on page 331

[Set snapshot group media scan](#) on page 333

[Set snapshot group schedule](#) on page 336

[Set SSD cache for a volume](#) on page 342

[Show snapshot group](#) on page 399

[Stop snapshot group pending snapshot images](#) on page 472

## Set snapshot group schedule

This command defines 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 ([ ]).
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 .



Parameter	Description
schedule	<p>Use this parameter to schedule a snapshot operation. You can use one of these options for setting a schedule for a snapshot operation:</p> <ul style="list-style-type: none"> <li>• <code>immediate</code></li> <li>• <code>startDate</code></li> <li>• <code>scheduleDay</code></li> <li>• <code>startTime</code></li> <li>• <code>scheduleInterval</code></li> <li>• <code>endDate</code></li> <li>• <code>timesPerDay</code></li> <li>• <code>timeZone</code></li> <li>• <code>scheduleDate</code></li> <li>• <code>month</code></li> </ul> <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 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.

### Related references

[Create snapshot group](#) on page 146

[Delete snapshot group](#) on page 185

[Revive snapshot group](#) on page 263

[Set snapshot group attributes](#) on page 331

[Set snapshot group media scan](#) on page 333

[Set snapshot group repository volume capacity](#) on page 334

[Set SSD cache for a volume](#) on page 342

[Show snapshot group](#) on page 399

[Stop snapshot group pending snapshot images](#) on page 472

## 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
<code>snapVolume</code>	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 ( [ ] ).
<code>mediaScanEnabled</code>	The setting to turn on or turn off media scan for the snapshot 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.)

Parameter	Description
redundancyCheckEnabled	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> .

### 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

### Related references

- [Create read-only snapshot volume](#) on page 145
- [Create snapshot volume](#) on page 156
- [Delete snapshot volume](#) on page 187
- [Rename snapshot volume](#) on page 242
- [Resume snapshot volume](#) on page 259
- [Revive snapshot volume](#) on page 263
- [Set read-only snapshot volume to read/write volume](#) on page 328
- [Set snapshot volume repository volume capacity](#) on page 340
- [Set SSD cache for a volume](#) on page 342
- [Show snapshot volumes](#) on page 403
- [Stop snapshot volume](#) on page 474

## 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 ([ ]).

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"> <li>• Use an existing repository volume: name</li> <li>• Create a new repository volume when you run this command</li> </ul> <p>The name of an existing repository volume is comprised of two parts:</p> <ul style="list-style-type: none"> <li>• The term repos</li> <li>• A four digit numerical identifier that you assign to the repository volume name</li> </ul> <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"> <li>• An integer value that represents a percentage of the base volume capacity</li> <li>• A decimal fraction value that represents a percentage of the base volume capacity</li> <li>• A specific size for the repository volume. Size is defined in units of bytes , KB , MB , GB , or TB .</li> </ul> <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 parentheses.</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

**Related references**

- [Create read-only snapshot volume](#) on page 145
- [Create snapshot volume](#) on page 156
- [Delete snapshot volume](#) on page 187
- [Rename snapshot volume](#) on page 242
- [Resume snapshot volume](#) on page 259
- [Revive snapshot volume](#) on page 263
- [Set read-only snapshot volume to read/write volume](#) on page 328
- [Set snapshot volume media scan](#) on page 339
- [Set SSD cache for a volume](#) on page 342
- [Show snapshot volumes](#) on page 403
- [Stop snapshot volume](#) on page 474

**Set 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 one of these:

- Standard volume, identified by either a name or a WWID
- Snapshot volume
- Consistency group snapshot volume

**Syntax applicable to a standard volume**

```
set volume ([volumeName] | <"WWID">) 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	The name of the standard volume for which you want to turn on or turn off the SSD cache. Enclose the name of the volume in square brackets ([ ]). If the name of the volume has special characters or numbers, you must enclose the name of the volume in double quotation marks (" ") inside of square brackets ([ ]).

Parameter	Description
volume	The World Wide Identifier (WWID) of the standard volume for which you want to turn on or turn off the SSD cache. Enclose the WWID in double quotation marks ( " ") inside angle brackets (< >).
snapVolume	The name of the snapshot volume for which you want to turn on or turn off the SSD cache. Enclose the name of snapshot volume in double quotation marks ( " ") inside of square brackets ( [ ] ).
cgSnapVolume	The name of the consistency group snapshot volume for which you want to turn on or turn off the SSD cache. Enclose the name of the consistency group snapshot volume 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

## Related references

- [Create snapshot group](#) on page 146
- [Delete snapshot group](#) on page 185
- [Revive snapshot group](#) on page 263
- [Set snapshot group attributes](#) on page 331
- [Set snapshot group media scan](#) on page 333
- [Set snapshot group repository volume capacity](#) on page 334
- [Set snapshot group schedule](#) on page 336
- [Show snapshot group](#) on page 399
- [Stop snapshot group pending snapshot images](#) on page 472
- [Create read-only snapshot volume](#) on page 145
- [Create snapshot volume](#) on page 156
- [Delete snapshot volume](#) on page 187
- [Rename snapshot volume](#) on page 242
- [Resume snapshot volume](#) on page 259
- [Revive snapshot volume](#) on page 263
- [Set read-only snapshot volume to read/write volume](#) on page 328
- [Set snapshot volume media scan](#) on page 339
- [Set snapshot volume repository volume capacity](#) on page 340
- [Show snapshot volumes](#) on page 403
- [Stop snapshot volume](#) on page 474
- [Add drives to SSD cache](#) on page 77
- [Change SSD cache application type](#) on page 90

[Create SSD cache](#) on page 159  
[Delete SSD cache](#) on page 188  
[Remove drives from SSD cache](#) on page 235  
[Rename SSD cache](#) on page 242  
[Resume SSD cache](#) on page 260  
[Show SSD cache](#) on page 406  
[Show SSD cache statistics](#) on page 407  
[Start SSD cache locate](#) on page 457  
[Start SSD cache performance modeling](#) on page 458  
[Stop SSD cache locate](#) on page 475  
[Stop SSD cache performance modeling](#) on page 476  
[Suspend SSD cache](#) on page 482

## Set storage array

**Note:** With firmware version 8.10 the `cacheFlushStop` parameter is deprecated and will be removed in a future release of storage management software.

This command defines the properties of the storage array.

### Syntax

```

set storageArray
[autoSupportConfig (enable | disable)] |
[cacheBlockSize=cacheBlockSizeValue]
[cacheFlushStart=cacheFlushStartSize]
[cacheFlushStop=cacheFlushStopSize]
[defaultHostType=(hostTypeName |
hostTypeIdentifier)]
[failoverAlertDelay=delayValue]
[mediaScanRate=(disabled | 1-30)]
[userRole=(admin | monitor)]
[password="password"]
[userLabel="storageArrayName"]
[isnsRegistration=(TRUE | FALSE)]
[autoLoadBalancingEnable=(TRUE | FALSE)]
  
```

### Parameters

Parameter	Description
<code>autoSupportConfig</code>	<p>The setting for automatically collecting support data each time the firmware detects a critical MEL event. This parameter has these values:</p> <ul style="list-style-type: none"> <li><code>enable</code> – Turns on the collection of support data</li> <li><code>disable</code> – Turns off the collection of support data</li> </ul> <p>Use this parameter under these conditions only:</p> <ul style="list-style-type: none"> <li>The AutoSupport (ASUP) feature is not available on the storage array</li> <li>To support configurations on storage arrays that previously used this parameter</li> </ul> <p><b>Note:</b> You cannot use this parameter with the ASUP commands.</p>



Parameter	Description
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	<p><b>Note:</b> This parameter is deprecated and will be removed in a future release.</p> <p>The percentage of unwritten data in the cache that stops a cache flush. 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.</p>
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 (" ").
userRole	<p>Defines the user role for the password. The user role can be either:</p> <ul style="list-style-type: none"> <li><code>admin</code> – Enables you to view and modify storage array configurations</li> <li><code>monitor</code> – Enables you to view storage array configurations and monitor storage array health conditions, but not modify the configuration</li> </ul> <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>
userLabel	The name for the storage array. Enclose the storage array name in double quotation marks (" ").

Parameter	Description
<code>isnsRegistration</code>	<p>This parameter lists the storage array as an iSCSI target on the iSNS server. To register the storage array on the iSNS server, set this parameter to <code>TRUE</code> . To remove the storage array from the iSNS server, set this parameter to <code>FALSE</code> .</p> <p><b>Note:</b> You cannot use the <code>isnsRegistration</code> parameter with any other parameter when running the <code>set storageArray</code> command.</p> <p>For more information about iSNS registration, refer to the <code>set storageArray isnsRegistration</code> command.</p>
<code>autoLoadBalancingEnable</code>	<p>This parameter enables auto-load balancing on the controller. To enable the option, set this parameter to <code>TRUE</code> . To disable the option, set this parameter to <code>FALSE</code> .</p>

### Notes

Except for the `isnsRegistration`, 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

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.

### 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 storage array before you can create a security key for encrypted full disk encryption drives.

### Auto-load Balancing

Auto-load balancing description here.

### 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.

7.86 removes the `alarm` parameter because it is no longer used.

8.10 deprecates the `cacheFlushStop` parameter.

8.20 updates the `cacheBlockSize` parameter to add the `cacheBlockSizeValue` of 4 (4 KB).

8.30 adds the `autoLoadBalancingEnable` parameter.

#### Related references

[Set storage array iSNS registration](#) on page 351

[Set storage array iSNS server IPv4 address](#) on page 353

[Set storage array iSNS server IPv6 address](#) on page 354

[Set storage array iSNS server listening port](#) on page 354

[Set storage array iSNS server refresh](#) on page 355

[Start storage array iSNS server refresh](#) on page 461

## 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 controller health image allow overwrite

**Note:** With firmware version 8.20 the `coreDumpAllowOverWrite` parameter is deprecated and will be removed in a future release of storage management software. For best operation use the `controllerHealthImageAllowOverWrite` parameter.

This command sets a flag on a controller to allow a new controller health image to overwrite an existing controller health image.

## Syntax

```
set storageArray controllerHealthImageAllowOverWrite
```

## Parameters

None.

## Minimum firmware level

7.86

8.20 replaces the `coreDumpAllowOverWrite` parameter with the `controllerHealthImageAllowOverWrite` parameter.

## Related references

[Save storage array controller health image](#) on page 272

[Show storage array controller health image](#) on page 419

[Start storage array controller health image](#) on page 460

## Set storage array to enable or disable Automatic Load Balancing

This command enables or disables the Automatic Load Balancing feature.

## Syntax

```
set storageArray autoLoadBalancingEnable=(TRUE | FALSE)
```

## Parameter

Parameter	Description
<code>storageArray</code>	Indicates that this parameter works on a storage array.
<code>autoLoadBalancingEnable</code>	This parameter turns on or turns off Automatic Load Balancing on the controller. Set the parameter to <code>TRUE</code> to turn on Automatic Load Balancing. Set the parameter to <code>FALSE</code> to turn off Automatic Load Balancing.

## What is Automatic Load Balancing?

The Automatic Load Balancing feature provides improved I/O resource management by reacting dynamically to load changes over time and automatically adjusting volume controller ownership to correct any load imbalance issues when workloads shift across the controllers.

The workload of each controller is continually monitored and, with cooperation from the multipath drivers installed on the hosts, can be automatically brought into balance whenever necessary. When workload is automatically re-balanced across the controllers, the storage administrator is relieved of the burden of manually adjusting volume controller ownership to accommodate load changes on the storage array.

When Automatic Load Balancing is enabled, it performs the following functions:

- Automatically monitors and balances controller resource utilization.
- Automatically adjusts volume controller ownership when needed, thereby optimizing I/O bandwidth between the hosts and the storage array.

## Enabling and disabling Automatic Load Balancing

Automatic Load Balancing is enabled by default on all storage arrays that ship with SANtricity OS (controller software) 8.30 or later. Linux, Windows, and VMware multipath drivers can use the Automatic Load Balancing feature. If you upgrade your controller from SANtricity OS (controller software) 8.25 or earlier to 8.30 or later, Automatic Load Balancing is disabled by default on the storage array.

You might want to disable Automatic Load Balancing on your storage array for the following reasons:

- You do not want to automatically change a particular volume's controller ownership to balance workload.
- You are operating in a highly tuned environment where load distribution is purposefully set up to achieve a specific distribution between the controllers.

In SANtricity Storage Manager, select the **Storage Array > Configuration > Automatic Load Balancing** menu option to enable or disable the Automatic Load Balancing feature for an individual storage array.

In SANtricity System Manager, select **Settings > System**, scroll down to the Additional Settings section, click the Enable/Disable Automatic Load Balancing link, and select the Enable/Disable automatic load balancing checkbox to enable or disable the feature for an individual storage array.

## Examples

```
SMcli -n Array1 -c "set storageArray autoLoadBalancingEnable=true;"
SMcli completed successfully.
```

## Minimum firmware level

8.30

## Related references

[Save auto-load balancing statistics](#) on page 270

## 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 <code>TRUE</code> to turn on Echo Request messages. Set the parameter to <code>FALSE</code> 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 registration

This command enables you to list a storage array on either an IPv4 or an IPv6 Internet Storage Name Service (iSNS) server.

### Syntax for IPv4

```
set storageArray isnsRegistration=(TRUE | FALSE)
isnsIPv4ConfigurationMethod=[static | dhcp]
isnsIPv4Address=ipAddress
[isnsListeningPort]
```

### Syntax for IPv6

```
set storageArray isnsRegistration=(TRUE | FALSE)
isnsIPv6Address=ipAddress
[isnsListeningPort]
```

## Parameters

Parameters	Description
isnsRegistration	<p>The means of registering the iSCSI target on the iSNS server. Set the parameter to <code>TRUE</code> to list an iSCSI target.</p> <p>When you set this parameter to <code>TRUE</code> you must also use these parameters for IPV4 configurations:</p> <ul style="list-style-type: none"> <li>isnsIPv4ConfigurationMethod</li> <li>isnsIPv4Address</li> </ul> <p>When you set this parameter to <code>TRUE</code> you must also use these parameters for IPV6 configurations:</p> <ul style="list-style-type: none"> <li>isnsIPv6Address</li> </ul> <p>Optionally, you can also use the <code>isnsListeningPort</code> parameter to define the port monitor and manage traffic to the server.</p> <p>To remove the registration for the storage array from the iSNS server set this parameter to <code>FALSE</code> .</p>

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 IPv4 address used to connect to 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> .
<code>isnsIPv6Address</code>	The IPv6 address used to connect to the iSNS server.
<code>isnsListeningPort</code>	<p>The port number 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.</p> <p>The listening port resides on the server and performs these activities:</p> <ul style="list-style-type: none"> <li>• Monitors incoming client connection requests</li> <li>• Manages traffic to the server</li> </ul> <p>When a client requests a network session with a server, the listener receives the actual request. If the client information matches the listener information, then the listener grants a connection to the database server.</p>

## 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

## Related references

- [Set storage array](#) on page 344
- [Set storage array iSNS server IPv4 address](#) on page 353
- [Set storage array iSNS server IPv6 address](#) on page 354
- [Set storage array iSNS server listening port](#) on page 354
- [Set storage array iSNS server refresh](#) on page 355
- [Start storage array iSNS server refresh](#) on page 461



## 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
iSNSIPv4ConfigurationMethod	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> .
iSNSIPv4Address	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

### Related references

- [Set storage array](#) on page 344
- [Set storage array iSNS registration](#) on page 351
- [Set storage array iSNS server IPv6 address](#) on page 354
- [Set storage array iSNS server listening port](#) on page 354
- [Set storage array iSNS server refresh](#) on page 355
- [Start storage array iSNS server refresh](#) on page 461

## Set storage array iSNS server IPv6 address

This command sets the IPv6 address for the iSNS server.

### Syntax

```
set storageArray isnsIPv6Address=ipAddress
```

### Parameter

Parameters	Description
isnsIPv6Address	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

### Related references

- [Set storage array](#) on page 344
- [Set storage array iSNS registration](#) on page 351
- [Set storage array iSNS server IPv4 address](#) on page 353
- [Set storage array iSNS server listening port](#) on page 354
- [Set storage array iSNS server refresh](#) on page 355
- [Start storage array iSNS server refresh](#) on page 461

## Set storage array iSNS server listening port

This command sets the iSNS server listening port.

### Syntax

```
set storageArray isnsListeningPort=listeningPortIPAddress
```

**Parameter**

Parameter	Description
<code>isnsListeningPort</code>	<p>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 53205 .</p> <p>The listening port resides on the server and performs these activities:</p> <ul style="list-style-type: none"> <li>• Monitors incoming client connection requests</li> <li>• Manages traffic to the server</li> </ul> <p>When a client requests a network session with a server, the listener receives the actual request. If the client information matches the listener information, then the listener grants a connection to the database server.</p>

**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

**Related references**

- [Set storage array](#) on page 344
- [Set storage array iSNS registration](#) on page 351
- [Set storage array iSNS server IPv4 address](#) on page 353
- [Set storage array iSNS server IPv6 address](#) on page 354
- [Set storage array iSNS server refresh](#) on page 355
- [Start storage array iSNS server refresh](#) on page 461
- [Set storage array iSNS server IPv4 address](#) on page 353
- [Set storage array iSNS server IPv6 address](#) on page 354
- [Set storage array iSNS server refresh](#) on page 355
- [Start storage array iSNS server refresh](#) on page 461

**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

**Related references**

- [Set storage array](#) on page 344
- [Set storage array iSNS registration](#) on page 351
- [Set storage array iSNS server IPv4 address](#) on page 353
- [Set storage array iSNS server IPv6 address](#) on page 354
- [Set storage array iSNS server listening port](#) on page 354
- [Start storage array iSNS server refresh](#) on page 461
- [Set storage array iSNS server IPv4 address](#) on page 353
- [Set storage array iSNS server IPv6 address](#) on page 354
- [Set storage array iSNS server listening port](#) on page 354
- [Start storage array iSNS server refresh](#) on page 461

**Set storage array controller battery learn cycle**

A learn cycle is an automatic cycle for calibrating the smart battery gauge. 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. The `set storageArray learnCycleDate` command sets the learn cycle for the battery backup unit in a single controller in the array, or for the battery backup unit in both controllers.

**Context**

**Important:** To set individual learn cycles for each controller in the storage array, send this command to both controllers.

**Important:** The learn cycles are not linked together and so a disruption in one controller battery learn cycle will not impact the learn cycle for the battery in the other controller.

**Note:** A learn cycle might take several hours to complete.

**Syntax**

```
set storageArray learnCycleDate (controller[(a| b)] )
(daysToNextLearnCycle=numberOfDays |
day=dayOfTheWeek) time=HH:MM
```

## Parameters

Parameter	Description
controller	The controller for which you want to specify a battery learn cycle. 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 controller firmware returns a syntax error.  <b>Note:</b> Omitting this parameter sets the learn cycle for both controller batteries in a dual-controller array.
daysToNextLearnCycle	Valid values are 0 through 7 , where 0 is immediately and 7 is in seven days. The daysToNextLearnCycle parameter takes place up to seven days after the next scheduled learn cycle.
day	Valid values for the day parameter include the days of the week ( Sunday , Monday , Tuesday , Wednesday , Thursday , Friday , and Saturday ). Setting the day causes the next learn cycle to be scheduled on the specified day, after the currently scheduled learn cycle.
time	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 .

## Example

```
set storageArray learnCycleDate controller [a] daysToNextLearnCycle=4
time=08:30;
```

## 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

8.30 - added the controller parameter.

## 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
redundancyMode	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 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 controller 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.

**Note:** The management software visually shows the first position on top, and the other positions in relative order beneath with the last position on the bottom.

**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
unnamedDiscoverySession	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

**Related references**

[Stop storage array iSCSI session](#) on page 479

**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`.

**Note:** Starting with firmware 8.10, the `writeOrder` parameter and the `writeMode` parameter cannot be used on storage arrays that have either an E2700 controller or an E5500 controller. The parameters are still valid for storage arrays that have an E2600 controller or an E5400 controller.

**Syntax**

```
set syncMirror (localVolume [volumeName] |
localVolumes [volumeName1 ... volumeNameN])
[role=(primary | secondary)]
[syncPriority=(highest | high | medium | low | lowest)]
[autoResync=(enabled | disabled)]
[writeOrder=(preserved | notPreserved)]
[writeMode=(synchronous | asynchronous)]
[force=(TRUE | FALSE)]
```

**Parameters**

Parameter	Description
localVolume	The name of the local volume for which you want to define properties. Enclose the primary volume name in square brackets ([ ]). If the primary volume name has special characters or numbers, you must enclose the primary volume name in double quotation marks (" ") inside square brackets.



Parameter	Description
<code>localVolumes</code>	<p>The names of several local volumes for which you want to define properties. Enter the names of the volumes using these rules:</p> <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ( [ ] ).</li> <li>• Separate each of the names with a space.</li> </ul> <p>If the volume names have special characters or consist only of numbers, enter the names using these rules:</p> <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ( [ ] ).</li> <li>• Enclose each of the names in double quotation marks ( " " ).</li> <li>• Separate each of the names with a space.</li> </ul>
<code>role</code>	<p>The setting for the local 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.</p>
<code>syncPriority</code>	<p>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> .</p>
<code>autoResync</code>	<p>The settings for automatic resynchronization between the primary volumes and the secondary volumes of a remote-mirrored pair. This parameter has these values:</p> <ul style="list-style-type: none"> <li>• <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.</li> <li>• <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.</li> </ul>
<code>writeOrder</code>	<p>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> .</p>
<code>writeMode</code>	<p>This parameter defines how the primary volume writes to the secondary volume. Valid values are <code>synchronous</code> or <code>asynchronous</code> .</p>
<code>force</code>	<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><b>Note:</b> Use the <code>force</code> parameter only when using the <code>role</code> parameter.</p>

## 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 write modes 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

8.10 the `writeOrder` parameter and the `writeMode` parameter cannot be used on storage arrays that have either an E2700 controller or an E5500 controller.

### Related references

[Check synchronous mirroring status](#) on page 95

[Create synchronous mirroring](#) on page 162

[Diagnose synchronous mirroring](#) on page 196

[Remove synchronous mirroring](#) on page 238

[Show synchronous mirroring volume candidates](#) on page 426

## 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][hostReportingPolicy=(standard |
thin)]
```

### Parameters

Parameter	Description
volume	The name of the volume for which you want to define properties. Enclose the volume name in square brackets ([ ]). If the volume name has special characters or numbers, you must enclose the volume name in double quotation marks (" ") inside square brackets.
volume	The World Wide Identifier (WWID) of the volume for which you want to define properties. Enclose the WWID in double quotation marks (" ") inside angle brackets (<>).

Parameter	Description
volumes	<p>The names of several volumes for which you want to define properties. All of the volumes will have the same properties. Enter the names of the volumes using these rules:</p> <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Separate each of the names with a space.</li> </ul> <p>If the volume names have special characters or consist only of numbers, enter the names using these rules:</p> <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Enclose each of the names in double quotation marks (" ").</li> <li>• Separate each of the names with a space.</li> </ul>
newCapacity	<p>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 <code>bytes</code>, <code>KB</code>, <code>MB</code>, <code>GB</code>, or <code>TB</code>.</p> <p>The minimum virtual capacity is 32 MB. The maximum virtual capacity is 63 TB.</p>
repositoryMaxCapacity	<p>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.</p>
repositoryExpansionPolicy	<p>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.</p>
warningThresholdPercent	<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>
addRepositoryCapacity	<p>This parameter allocates capacity from the free extent of the disk pool. If insufficient space is available the command fails.</p>
hostReportingPolicy	<p>This parameter changes the way in which the thin volume is reported to the host. Valid values are <code>standard</code> or <code>thin</code>.</p>

## 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 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

### Related references

[Initialize thin volume](#) on page 221

[Set volume attributes for a volume in a disk pool](#) on page 366

[Set volume attributes for a volume in a volume group](#) on page 371

[Show thin volume](#) on page 428

## 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	<p>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 (" ").</p> <p><b>Note:</b> Precede the serial number with SN and a space. For example:</p> <pre>SN SV23802522</pre>

Parameter	Description
id	<p>The value for the controller tray ID, controller-drive tray ID, or the drive tray ID. Tray ID values are 0 through 9 or 10 through 99.</p> <p><b>Note:</b> You do not need to enclose the tray ID value in parentheses.</p> <p><b>Note:</b> Tray values of 00 through 09 are rejected by the controller. Use 0 through 9 to specify those values.</p>

### 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.

### Related references

[Set drawer service action allowed indicator](#) on page 312

[Set tray service action allowed indicator](#) on page 365

## 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 ([ ]).

Parameter	Description
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> .

### 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.

### 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;"
```

### Minimum firmware level

6.14 adds these parameters:

- `powerFan`
- `interconnect`

6.16 adds these parameters:

- `tray`
- `esm`

7.60 adds the identifiers `top` and `bottom`.

### Related references

[Set drawer service action allowed indicator](#) on page 312

## Set volume attributes for a volume in 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 = (trayID1,drawerID1,slotID1 ... trayIDn,drawerIDn,slotIDn) |
cacheFlushModifier=cacheFlushModifierValue |
cacheReadPrefetch = (TRUE | FALSE) |
cacheWithoutBatteryEnabled=(TRUE | FALSE) |
mediaScanEnabled=(TRUE | FALSE) |
mirrorCacheEnabled=(TRUE | FALSE) |
preReadRedundancyCheck = (TRUE | FALSE) |
readCacheEnabled=(TRUE | FALSE) |
redundancyCheckEnabled = (TRUE | FALSE) |
segmentSize = segmentSizeValue
userLabel = userLabelValue
writeCacheEnabled=(TRUE | FALSE) |
dataAssuranceDisabled=(TRUE | FALSE)

```

## Parameters

Parameter	Description
allVolumes	This parameter sets the properties for all of volumes in the storage array.
volume	The name of the volume for which you want to define properties. Enclose the volume name in square brackets ([ ]). If the volume name has special characters or numbers, you must enclose the volume name in double quotation marks (" ") inside square brackets.
volume	The World Wide Identifier (WWID) of the volume for which you want to define properties. Enclose the WWID in double quotation marks (" ") inside angle brackets (<>).  <b>Note:</b> When running this command, do not use colon separators in the WWID.
volumes	The names of several volumes for which you want to define properties. All of the volumes will have the same properties. Enter the names of the volumes using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Separate each of the names with a space.</li> </ul> If the volume names have special characters or numbers, enter the names using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Enclose each of the names in double quotation marks (" ").</li> <li>• Separate each of the names with a space.</li> </ul>

Parameter	Description
<code>addCapacity</code>	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> .  This parameter is not valid for thin volumes.
<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>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>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>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. Use this parameter only if you want to change the volume owner.
<code>preReadRedundancyCheck</code>	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 <code>TRUE</code> . To turn off pre-read redundancy checking, set this parameter to <code>FALSE</code> .  <b>Note:</b> Do not use this parameter on non-redundant volumes, such as RAID 0 volumes.
<code>readCacheEnabled</code>	The setting to turn on or turn off the read cache. To turn on the read cache, set this parameter to <code>TRUE</code> . To turn off the read cache, set this parameter to <code>FALSE</code> .
<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>userLabel</code>	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>dataAssuranceDisabled</code>	<p>The setting to turn on or 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><b>Note:</b> This option is only valid for enablement if the drives support DA.</p> <p>To remove data assurance from a volume that supports data assurance, set this parameter to <code>TRUE</code>. To return a volume to supporting data assurance, set this parameter to <code>FALSE</code>.</p> <p><b>Note:</b> If you remove data assurance from a volume, you cannot reset data assurance for that volume.</p> <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"> <li>1. Remove the data from the volume.</li> <li>2. Delete the volume.</li> <li>3. Recreate a new volume with the properties of the deleted volume.</li> <li>4. Set data assurance for the new volume.</li> <li>5. Move the data to the new volume.</li> </ol>

## 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.
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

7.83

## Related references

[Create volume in disk pool](#) on page 172

[Delete volume from disk pool](#) on page 190

[Increase capacity of volume in disk pool or volume group](#) on page 219

[Set volume attributes for a volume in a volume group](#) on page 371

[Show thin volume](#) on page 428

[Show volume](#) on page 429

## Set volume attributes for a volume in a volume group

This command defines the properties for volumes 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
allVolumes	This parameter sets the properties for all of the volumes in the storage array.
volume	The name of the volume for which you want to define properties. Enclose the volume name in square brackets ([ ]). If the volume name has special characters or numbers, you must enclose the volume name in double quotation marks (" ") inside square brackets.
volume	The World Wide Identifier (WWID) of the volume for which you want to define properties. Enclose the WWID in double quotation marks (" ") inside angle brackets (< >).  <b>Note:</b> When running this command, do not use colon separators in the WWID.
volumes	The names of several volumes for which you want to define properties. All of the volumes will have the same properties. Enter the names of the volumes using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Separate each of the names with a space.</li> </ul> If the volume names have special characters or numbers, enter the names using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Enclose each of the names in double quotation marks (" ").</li> <li>• Separate each of the names with a space.</li> </ul>
cacheFlushModifier	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.

Parameter	Description
cacheWithoutBatteryEnabled	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> .
mediaScanEnabled	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.)
mirrorCacheEnabled	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> .
modificationPriority	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> .
owner	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. Use this parameter only if you want to change the volume owner.
preReadRedundancyCheck	The setting to turn on or turn off preread redundancy checking. Turning on preread redundancy checking verifies the consistency of RAID redundancy data for the stripes containing the read data. Preread redundancy checking is performed on read operations only. To turn on preread redundancy checking, set this parameter to <code>TRUE</code> . To turn off preread redundancy checking, set this parameter to <code>FALSE</code> .  <b>Note:</b> 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 <code>TRUE</code> . To turn off the read cache, set this parameter to <code>FALSE</code> .
writeCacheEnabled	The setting to turn on or turn off the write cache. To turn on the write cache, set this parameter to <code>TRUE</code> . To turn off the write cache, set this parameter to <code>FALSE</code> .
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> .

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><b>Note:</b> This option is only valid for enablement if the drives support DA.</p> <p>To remove data assurance from a volume that supports data assurance, set this parameter to <code>TRUE</code>.</p> <p><b>Note:</b> If you remove data assurance from a volume, you cannot reset data assurance for that volume.</p> <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"> <li>1. Remove the data from the volume.</li> <li>2. Delete the volume.</li> <li>3. Recreate a new volume with the properties of the deleted volume.</li> <li>4. Set data assurance for the new volume.</li> <li>5. Move the data to the new volume.</li> </ol>
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 24. 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>
redundancyCheckEnabled	<p>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>.</p>
segmentSize	<p>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.</p>

Parameter	Description
<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.
<code>.25</code>	Data is flushed after 250 ms.
<code>.5</code>	Data is flushed after 500 ms.
<code>.75</code>	Data is flushed after 750 ms.

Value	Description
1	Data is flushed after 1 s.
1.5	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).
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.

**Important:** Do not set the value of the `cacheFlushModifier` parameter above 10 seconds. An exception is for testing purposes. After running any tests in which you have set the values of the `cacheFlushModifier` parameter above 10 seconds, return the value of the `cacheFlushModifier` parameter to 10 or fewer seconds.

### 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.

8.10 corrects the values for the `cacheFlushModifier` parameter in the cache flush table.

## 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 snapshot volume copy pairs.

### Syntax

```
set volumeCopy target [targetName]
source [sourceName]
[copyPriority=(highest | high | medium | low | lowest)]
[targetReadOnlyEnabled=(TRUE | FALSE)]
[copyType=(online | offline)]
```

### Parameters

Parameter	Description
<code>target</code>	The name of the volume to which the data will be copied. Enclose the target volume name in square brackets ([ ]). If the target volume name has special characters or numbers, you must enclose the target volume name in double quotation marks (" ") inside square brackets.
<code>source</code>	The name of the volume from which the data will be copied. Enclose the source volume name in square brackets ([ ]). If the source volume name has special characters or numbers, you must enclose the source volume name in double quotation marks (" ") inside square brackets.
<code>copyPriority</code>	The priority that the volume copy 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
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 <code>FALSE</code> . To prevent writing to the target volume, set this parameter to <code>TRUE</code> .
copyType	<b>Note:</b> This parameter has been deprecated.

## Notes

When you use this command, you can specify one or more of the optional parameters.

## Minimum firmware level

5.40

7.83 removes the `copyType=(online | offline)` parameter.

## Related references

[Create volume copy](#) on page 165

[Show volume copy](#) on page 431

## Set volume group

Use the `set volumeGroup` command to add drives, change the RAID level, or change volume group ownership from one controller to another.

This command defines the properties for a volume group.

**Note:** You must perform only one operation (add drives, change the RAID level, or change volume group ownership) per command. You cannot perform more than one operation with a single command.

## Syntax

```
set volumeGroup [volumeGroupName]
[addDrives=(trayID1,drawerID1,slotID1 ... trayIDn,drawerIDn,slotIDn) ]|
[raidLevel=(0 | 1 | 3 | 5 | 6) ]|
[owner=(a|b)]
```

## Parameters

Parameter	Description
volumeGroup	The name of the volume group for which you want to set properties. Enclose the volume group name in square brackets ([ ]).
addDrives	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 24 . Enclose the tray ID value, the drawer ID value, and the slot ID value in parentheses.  <b>Note:</b> Only 1 or 2 drives can be added in a single operation.

Parameter	Description
raidLevel	The RAID level for the volume group. Valid values are 0 , 1 , 3 , 5 , or 6 .
owner	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 only one 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

7.10 adds RAID 6 capability.

7.30 removes the `availability` parameter.

7.60 adds the `drawerID` user input.

### Related references

[Set volume group forced state](#) on page 380

[Show volume group](#) on page 433

## 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
<code>volumeGroup</code>	The name of the volume group that you want to place in a Forced state. Enclose the volume group name 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

### Related references

[Set volume group](#) on page 378

## Set volume mapping

This command defines the logical unit number (LUN) mapping between a volume and 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 volume for which you want to define the logical unit number. Enclose the volume name in square brackets ([ ]). If the volume name has special characters or numbers, you must enclose the volume name in double quotation marks (" ") inside square brackets.
volume	The World Wide Identifier (WWID) of the volume for which you want to define the logical unit number. Enclose the WWID in double quotation marks (" ") inside 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  <b>Note:</b> The logical unit number you specify must not already be in use, and must be within the range supported by the host Operating System. There will be no error if you map the volume to a logical unit number not supported by the host, but the host won't be able to access the 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 to which the volume is mapped. Enclose the host group name in double quotation marks (" "). <code>defaultGroup</code> 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

7.83 adds snapshot volume for disk pools.

## Related references

[Clear volume reservations](#) on page 102

*Remove volume LUN mapping* on page 240

*Show volume reservations* on page 437

## Show asynchronous mirror groups

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] |
    asyncMirrorGroups [asyncMirrorGroupName1 ... asyncMirrorGroupNameN])
[summary]
```

### Parameter

Parameter	Description
allAsyncMirrorGroups	Use this parameter if you want to display the properties for all asynchronous mirror groups.
asyncMirrorGroup	The name of an asynchronous mirror group for which you want to show configuration information and the progress of periodic data synchronization. Enclose the asynchronous mirror group name in square brackets ([ ]). If the asynchronous mirror group name has special characters or numbers, you must enclose the asynchronous mirror group name in double quotation marks (" ") inside square brackets.
asyncMirrorGroups	The names of several asynchronous mirror groups for which you want to retrieve information. Enter the names of the asynchronous mirror groups using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Separate each of the names with a space.</li> </ul> If the asynchronous mirror group names have special characters or numbers, enter the names using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Enclose each of the names in double quotation marks (" ").</li> <li>• Separate each of the names with a space.</li> </ul>
summary	This parameter displays a concise list of information about the synchronization progress of one or more asynchronous mirror groups.

### Minimum firmware level

7.84

## Related references

- [Add volume to asynchronous mirror group](#) on page 80
- [Clear asynchronous mirroring fault](#) on page 96
- [Create asynchronous mirror group](#) on page 104
- [Delete asynchronous mirror group](#) on page 177
- [Remove incomplete asynchronous mirrored pair from asynchronous mirror group](#) on page 236
- [Remove volume from asynchronous mirror group](#) on page 239
- [Reset asynchronous mirror group statistics](#) on page 246
- [Save asynchronous mirror group statistics](#) on page 265
- [Set asynchronous mirror group](#) on page 294
- [Show asynchronous mirror group synchronization progress](#) on page 383
- [Start asynchronous mirroring synchronization](#) on page 440
- [Suspend asynchronous mirror group](#) on page 481
- [Test asynchronous mirror group connectivity](#) on page 484

## 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]
[summary]
```

## Parameters

Parameter	Description
<code>asyncMirrorGroup</code>	The name of an asynchronous mirror group for which you want to show synchronization of the asynchronous mirror group between the local and remote storage array. Enclose the asynchronous mirror group name in square brackets ([ ]). If the asynchronous mirror group name has special characters or numbers, you must enclose the asynchronous mirror group name in double quotation marks (" ") inside square brackets.
<code>synchronizationProgress</code>	This parameter displays the periodic synchronization progress of the asynchronous mirror group.
<code>summary</code>	This parameter displays a concise list of information about the synchronization of the asynchronous mirror group between the local and remote storage array.

## Minimum firmware level

7.84

**Related references**

[Resume asynchronous mirror group](#) on page 255

[Start asynchronous mirroring synchronization](#) on page 440

[Suspend asynchronous mirror group](#) on page 481

[Test asynchronous mirror group connectivity](#) on page 484

**Show blocked events**

This command returns a list of events that are currently blocked by the `set blockEventAlert` command. The events in this list are the events that will not send any notifications that you have configured using the alert terminals and parameters; that is, the email, syslog, and trap notifications.

**Restriction:** This command is obsolete for the E2800 storage array. Use the RESTful API, SANtricity System Manager, or see [Using curl commands with the E2800 storage array](#) on page 39.

**Syntax**

```
show blockedEventAlertList
```

**Example**

The output has one line for each event that is blocked, listing the type of event in a hexadecimal format followed by a description of the event. Following is an example of the output:

```
Executing Script...
0x280D Enclosure Failure
0x282B Tray Redundancy Lost
Script execution complete.
```

**Parameters**

None.

**Minimum firmware level**

8.10

**Related references**

[Using curl commands with the E2800 storage array](#) on page 39

[Set event alert filtering](#) on page 316

[Test alerts](#) on page 484

[Clear storage array event log](#) on page 101

[Save storage array events](#) on page 276

**Show consistency group**

This command returns information about one or more consistency groups.

**Syntax**

```
show (allConsistencyGroups | consistencyGroup [consistencyGroupName] |
consistencyGroups [consistencyGroupName1 ... consistencyGroupNameN])
[(summary | schedule)]
```



**Parameters**

Parameter	Description
allConsistencyGroups	This setting returns information about all of the consistency groups in the storage array.
consistencyGroup	The name of the consistency group for which you are retrieving information. Enclose the consistency group name in square brackets ([ ]). If the consistency group name has special characters or numbers, you must enclose the consistency group name in double quotation marks (" ") inside square brackets.
consistencyGroups	The names of several consistency groups for which you are retrieving information. All of the volumes will have the same properties. Enter the names of the volumes using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Separate each of the names with a space.</li> </ul> If the volume names have special characters or consist only of numbers, enter the names using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Enclose each of the names in double quotation marks (" ").</li> <li>• Separate each of the names with a space.</li> </ul>
summary	This setting returns a concise list of information about the consistency groups.
schedule	This parameter returns information about any schedules for a consistency group.

**Minimum firmware level**

7.83

**Related references**

- [Add member to consistency group](#) on page 78
- [Create consistency group](#) on page 107
- [Create consistency group snapshot image](#) on page 115
- [Create consistency group snapshot volume](#) on page 116
- [Create consistency group snapshot volume mapping](#) on page 120
- [Delete consistency group](#) on page 178
- [Delete consistency group snapshot image](#) on page 179
- [Delete consistency group snapshot volume](#) on page 180
- [Remove member volume from consistency group](#) on page 237
- [Resume consistency group snapshot volume](#) on page 256
- [Set consistency group attributes](#) on page 296
- [Set consistency group snapshot volume](#) on page 298
- [Show consistency group snapshot image](#) on page 386
- [Start consistency group snapshot rollback](#) on page 443
- [Stop consistency group snapshot rollback](#) on page 466

[Stop consistency group snapshot volume](#) on page 468

[Stop pending snapshot images on consistency group](#) on page 471

## 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"> <li>The name of the consistency group</li> <li>An identifier for the snapshot image in the consistency group.</li> </ul> <p>The identifier for the snapshot image can be one of these:</p> <ul style="list-style-type: none"> <li>An integer value that is the sequence number of the snapshot in the consistency group.</li> <li>NEWEST - Use this option when you want to show the latest snapshot image created in the consistency group.</li> <li>OLDEST - Use this option when you want to show the earliest snapshot image created in the consistency group.</li> </ul> <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 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

### Related references

- [Add member to consistency group](#) on page 78
- [Create consistency group](#) on page 107
- [Create consistency group snapshot image](#) on page 115
- [Create consistency group snapshot volume](#) on page 116
- [Create consistency group snapshot volume mapping](#) on page 120
- [Delete consistency group](#) on page 178
- [Delete consistency group snapshot image](#) on page 179
- [Delete consistency group snapshot volume](#) on page 180
- [Remove member volume from consistency group](#) on page 237
- [Resume consistency group snapshot volume](#) on page 256
- [Set consistency group attributes](#) on page 296
- [Set consistency group snapshot volume](#) on page 298
- [Show consistency group](#) on page 384
- [Start consistency group snapshot rollback](#) on page 443
- [Stop consistency group snapshot rollback](#) on page 466
- [Stop consistency group snapshot volume](#) on page 468
- [Stop pending snapshot images on consistency group](#) on page 471

## 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.

### Related references

[Show controller NVSRAM](#) on page 390

## 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.

### Syntax

```
show (allControllers | controller [(a|b)])
NVSRAM [hostType=hostTypeIndexLabel | host="hostName"]
```

### Parameters

Parameter	Description
<code>allControllers</code>	The setting to return information about both controllers in the storage array.
<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>hostType</code>	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.
<code>host</code>	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 technical support to learn what regions of the NVSRAM you can modify.

### Related references

[Save controller NVSRAM](#) on page 267

## 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" | <"iqn">] iscsiSessions
```

```
show iscsiTarget ["targetName" | <"iqn">] iscsiSessions
```

## Parameters

Parameter	Description
iscsiInitiator	<p>The name of the iSCSI initiator for which you want to obtain session information.</p> <p>If the iSCSI initiator uses a label or an alias, enclose the iSCSI initiator label or alias in double quotation marks ( " ") inside square brackets ( [ ] ).</p> <p>If the iSCSI initiator uses an iSCSI qualified name (IQN), enclose the IQN in double quotation marks ( " ") inside angle brackets (&lt;&gt;).</p>
iscsiTarget	<p>The name of the iSCSI target for which you want to obtain session information.</p> <p>If the iSCSI target uses a label or an alias, enclose the iSCSI target label or alias in double quotation marks ( " ") inside square brackets ( [ ] ).</p> <p>If the iSCSI target uses an iSCSI qualified name (IQN), enclose the IQN in double quotation marks ( " ") inside angle brackets (&lt;&gt;).</p>

## 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.

An initiator name can be any combination of alpha or numeric characters from 1 to 30 characters in length. An IQN can be up to 255 characters in length and has this format:

```
iqn.yyyy-mm.naming-authority:unique name
```

## Minimum firmware level

7.10

## Related references

- [Activate iSCSI initiator](#) on page 72
- [Create iSCSI initiator](#) on page 131
- [Delete iSCSI initiator](#) on page 184
- [Diagnose controller iSCSI host cable](#) on page 193
- [Reset storage array iSCSI baseline](#) on page 252
- [Save storage array iSCSI statistics](#) on page 279

*Set iSCSI initiator* on page 324

*Set iSCSI target properties* on page 326

*Show storage array negotiation defaults* on page 423

*Show storage array unconfigured iSCSI initiators* on page 424

*Start iSCSI DHCP refresh* on page 453

*Stop storage array iSCSI session* on page 479

## 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 used capacity, free capacity, and percentage full
- The current owner (the controller in slot A or the controller in slot B )
- The drive media (Fibre Channel, SATA, or SAS)
- The drive media type (HDD or SSD)
- The drive interface type (Fibre Channel, iSCSI, Infiniband, SAS)
- Tray loss protection (yes or no)
- Secure capable - indicates whether the disk pool is composed of all secure-capable drives. Secure-capable drives can be either FDE drives or FIPS drives.
- Secure - Indicates whether the disk pool has drive security turned on (this is called secure enabled).
- The remaining capacity notification settings (critical and early warning)
- The priority settings
- 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 ([ ]). If the disk pool name has special characters or numbers, you must enclose the disk pool name in double quotation marks (" ") inside square brackets.

## Notes

Use this command to show the disk pool content of the storage array profile.



**Minimum firmware level**

7.83

**Related references**[Set disk pool](#) on page 307[Set disk pool \(modify disk pool\)](#) on page 310**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. 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.

Parameter	Description
<code>driveMediaType</code>	<p>The type of drive media for which you want to retrieve information. The following values are valid types of drive media:</p> <ul style="list-style-type: none"> <li>• <code>HDD</code> – Use this option when you have hard drives in the drive tray.</li> <li>• <code>SSD</code> – Use this option when you have solid state drives in the drive tray.</li> <li>• <code>unknown</code> – Use this option if you are not sure what types of drive media are in the drive tray.</li> <li>• <code>allMedia</code> – Use this option when you want to use all types of drive media that are in the drive tray.</li> </ul>
<code>driveType</code>	<p>The type of drive for which you want to retrieve information. You cannot mix drive types. Valid drive types are :</p> <ul style="list-style-type: none"> <li>• <code>fibre</code></li> <li>• <code>SATA</code></li> <li>• <code>SAS</code></li> </ul> <p>If you do not specify a drive type, the command defaults to <code>fibre</code> .</p>
<code>drive or drives</code>	<p>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 24 . Enclose the tray ID values, the drawer ID values, and the slot ID values in parentheses.</p>
<code>summary</code>	<p>The setting to return the status, the capacity, the data transfer rate, the product ID, and the firmware version for the specified drives.</p>

## 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.

**Related references**

[Show drive download progress](#) on page 396

[Show drive performance statistics](#) on page 396

**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.
<code>driveChannels</code>	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.
<code>allDriveChannels</code>	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.

**Related references**

[Clear drive channel statistics](#) on page 98

## 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

### Related references

[Download drive firmware](#) on page 202

## Show drive performance statistics

For each drive in the storage array, this command returns the following information:

- The storage array containing 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 24 . 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

## Related references

- [Download drive firmware](#) on page 202
- [Replace drive](#) on page 244
- [Revive drive](#) on page 262
- [Save drive channel fault isolation diagnostic status](#) on page 267
- [Save drive log](#) on page 268
- [Set drive hot spare](#) on page 314
- [Set drive service action allowed indicator](#) on page 314
- [Set drive state](#) on page 315
- [Set FIPS drive security identifier](#) on page 318
- [Set foreign drive to native](#) on page 319
- [Show drive](#) on page 393

- [Show drive download progress](#) on page 396
- [Start drive channel fault isolation diagnostics](#) on page 447
- [Start drive initialize](#) on page 449
- [Start drive locate](#) on page 450
- [Start drive reconstruction](#) on page 451
- [Start FDE secure drive erase](#) on page 455
- [Stop drive channel fault isolation diagnostics](#) on page 469
- [Stop drive locate](#) on page 470

## 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

### Related references

- [Create host](#) on page 127
- [Create host group](#) on page 128
- [Create host port](#) on page 129
- [Delete host](#) on page 182
- [Delete host group](#) on page 182
- [Delete host port](#) on page 183
- [Set host](#) on page 320
- [Set host channel](#) on page 321
- [Set host group](#) on page 322
- [Set host port label](#) on page 323
- [Show storage array host topology](#) on page 421

## Show replaceable drives

This command shows all replaceable drives in a storage array.

For each replaceable drive in the storage array, this command returns the following information:

- The tray location and the slot location
- The volume group name to which the drive belongs
- The World Wide Name (WWN)
- The status of the drive

### Syntax

```
show replaceableDrives
```

#### Example output

```
Replaceable drive at Tray 0, Slot 3
    Volume group: 2
    World-wide name: 5000c50028785aff0000000000000000
    Status: Removed
Replaceable drive at Tray 0, Slot 23
    Volume group: 2
    World-wide name: 5000c500095d46df00000000000000000
    Status: Removed
Replaceable drive at Tray 0, Slot 24
    Volume group: 2
    World-wide name: 5000c500287854d70000000000000000
    Status: Removed
```

## Show snapshot group

This command returns 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.
snapGroup	The name of the snapshot group for which you want to retrieve information. Enclose the snapshot group name in square brackets ([ ]). If the snapshot group name has special characters or numbers, you must enclose the snapshot group name in double quotation marks (" ") inside square brackets.

Parameter	Description
snapGroups	<p>The names of several snapshot groups for which you want to retrieve information. Enter the names of the snapshot groups using these rules:</p> <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ( [ ] ).</li> <li>• Separate each of the names with a space.</li> </ul> <p>If the snapshot group names have special characters or consist only of numbers, enter the names using these rules:</p> <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ( [ ] ).</li> <li>• Enclose each of the names in double quotation marks ( " " ).</li> <li>• Separate each of the names with a space.</li> </ul>
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

### Related references

[Create snapshot group](#) on page 146

[Delete snapshot group](#) on page 185

[Revive snapshot group](#) on page 263

[Set snapshot group attributes](#) on page 331

[Set snapshot group media scan](#) on page 333

[Set snapshot group repository volume capacity](#) on page 334

[Set snapshot group schedule](#) on page 336

[Set SSD cache for a volume](#) on page 342

[Stop snapshot group pending snapshot images](#) on page 472

## 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 (allSnapImages | snapImage ["snapImageName"] |
snapImages ["snapImageName1" ... "snapImageNameN"]])
[summary]
```

### Parameters

Parameter	Description
allSnapImages	The parameter to return information about all of the snapshot images in the storage array.

Parameter	Description
snapImage	<p>The name of a snapshot image for which you want to retrieve information. The name of a snapshot image is comprised of two parts:</p> <ul style="list-style-type: none"> <li>• The name of the snapshot group</li> <li>• An identifier for the snapshot image in the snapshot group</li> </ul> <p>The identifier for the snapshot image can be one of these:</p> <ul style="list-style-type: none"> <li>• An integer value that is the sequence number of the snapshot in the snapshot group.</li> <li>• <b>NEWEST</b> - Use this option when you want to show the latest snapshot image created in the snapshot group.</li> <li>• <b>OLDEST</b> - Use this option when you want to show the earliest snapshot image created in the snapshot group.</li> </ul> <p>Enclose the snapshot image name in double quotation marks ( " ") inside square brackets ( [ ] ).</p>
snapImages	<p>The names of several snapshot images for which you want to retrieve information. Enter the names of the snapshot images using these rules:</p> <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ( [ ] ).</li> <li>• Enclose each of the names in double quotation marks ( " " ).</li> <li>• Separate each of the names with a space.</li> </ul>
summary	<p>This parameter returns a concise list of information about the snapshot images.</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 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

## Related references

- [Create snapshot image](#) on page 155
- [Delete snapshot image](#) on page 186
- [Resume snapshot image rollback](#) on page 258
- [Start snapshot image rollback](#) on page 456
- [Stop snapshot image rollback](#) on page 473

## 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	The name of a snapshot volume about which you want to retrieve information. Enclose the snapshot volume name in double quotation marks (" ") inside of square brackets ([ ]).
snapVolumes	The names of several snapshot volumes for which you want to retrieve information. Enter the names of the snapshot volumes using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Enclose each of the names in double quotation marks (" ").</li> <li>• Separate each of the names with a space.</li> </ul>
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

### Related references

- [Create read-only snapshot volume](#) on page 145
- [Create snapshot volume](#) on page 156
- [Delete snapshot volume](#) on page 187
- [Rename snapshot volume](#) on page 242
- [Resume snapshot volume](#) on page 259
- [Revive snapshot volume](#) on page 263
- [Set read-only snapshot volume to read/write volume](#) on page 328
- [Set snapshot volume media scan](#) on page 339
- [Set snapshot volume repository volume capacity](#) on page 340
- [Set SSD cache for a volume](#) on page 342
- [Stop snapshot volume](#) on page 474

## Show SNMP communities

This command returns information about the Simple Network Management Protocol (SNMP) communities defined for the storage array. SNMP communities are collections of devices, such as routers, switches, printers, storage devices, grouped together to manage and monitor the devices.

Information returned includes the following:

- The total number of communities
- The total number of trap destinations
- The names of the SNMP communities

### Syntax

```
show (allSnmpCommunities |
snmpCommunity communityName="snmpCommunityName" |
snmpCommunities ["snmpCommunityName1" ... "snmpCommunityNameN"])
[summary]
```

### Parameters

Parameter	Description
allSnmpCommunities	This parameter returns information about all of the SNMP communities in the storage array.

Parameter	Description
communityName	The name of the SNMP community for which you want to retrieve information. Enclose the SNMP community name in double quotation marks (" ").
snmpCommunities	The names of several SNMP communities for which you want to retrieve information. Enter the names of the SNMP communities using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Enclose each of the names in double quotation marks (" ").</li> <li>• Separate each of the names with a space.</li> </ul>
summary	This parameter returns a list of the total number of SNMP communities and the total number of SNMP trap destinations. When you use this parameter, all of the detail information is omitted.

### Notes

SNMP supports one or more *communities* to which managers and managed devices belong. SNMP requests contain a *community string* in the data packets passed over the network that acts as a crude password. SNMP agents can reject GET or SET requests with an unrecognized community string. A community string is also included in TRAP notifications sent from agent to manager as well.

The embedded SNMP agent supports both IPV4 and IPV6.

This command returns SNMP community information similar to this example.

```
SNMP COMMUNITIES-----
SUMMARY
  Total SNMP Communities: 2
  Total SNMP Trap Destinations: 1
DETAILS
  SNMP Community: TestComm1
  SNMP Permission: Read Only
  Associated Trap Destination:
  Trap Receiver IP Address: 10.113.173.133
  Send Authentication Failure Traps: true
  SNMP Community: Test2
  SNMP Permission: Read Only
  Associated Trap Destination:
```

### Show SNMP MIB II system group variables

This command returns information about the Simple Network Management Protocol (SNMP) system variables. The system variables are maintained in a Management Information Base II (MIB-II) database.

### Syntax

```
show snmpSystemVariables
```

### Parameters

None.

**Notes**

This command returns SNMP system variables information similar to this example.

```
SYSTEM VARIABLES
  Sys Name: NONE
  Sys Contact: NONE
  Sys Location: NONE
```

**Show SSD cache**

This command displays information about the SSD cache.

**Syntax**

```
show ssdCache [ssdCacheName]
```

**Parameter**

Parameter	Description
ssdCache	The name of the SSD cache for which you want to retrieve information. Enclose the SSD cache name in square brackets ([ ]). If the SSD cache name has special characters or numbers, you must enclose the SSD cache name 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

**Related references**

- [Add drives to SSD cache](#) on page 77
- [Change SSD cache application type](#) on page 90
- [Create SSD cache](#) on page 159
- [Delete SSD cache](#) on page 188

[Remove drives from SSD cache](#) on page 235  
[Rename SSD cache](#) on page 242  
[Resume SSD cache](#) on page 260  
[Set SSD cache for a volume](#) on page 342  
[Show SSD cache statistics](#) on page 407  
[Start SSD cache locate](#) on page 457  
[Start SSD cache performance modeling](#) on page 458  
[Stop SSD cache locate](#) on page 475  
[Stop SSD cache performance modeling](#) on page 476  
[Suspend SSD cache](#) on page 482

## 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 name of the SSD cache for which you want to retrieve information. Enclose the SSD cache name in square brackets ([ ]). If the SSD cache name has special characters or numbers, you must enclose the SSD cache name in double quotation marks (" ") inside square brackets.
ssdCacheStatistics	An optional parameter to the show ssdCache command that indicates you want to retrieve cache statistics.
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 .
file	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.  Enclose the file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\logs\ssdcachestats.csv" . You can use any file name but you must use the .csv extension.

### 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 percent 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 percent. If this number is less than 100 percent, 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 percent, 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 percent 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 hit.  
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 retrieve 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 retrieve 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

### Related references

- [Add drives to SSD cache](#) on page 77
- [Change SSD cache application type](#) on page 90
- [Create SSD cache](#) on page 159
- [Delete SSD cache](#) on page 188
- [Remove drives from SSD cache](#) on page 235
- [Rename SSD cache](#) on page 242
- [Resume SSD cache](#) on page 260
- [Set SSD cache for a volume](#) on page 342
- [Show SSD cache](#) on page 406
- [Start SSD cache locate](#) on page 457

*Start SSD cache performance modeling* on page 458

*Stop SSD cache locate* on page 475

*Stop SSD cache performance modeling* on page 476

*Suspend SSD cache* on page 482

## 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	<p>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.</p> <p><b>Note:</b> The profile parameter returns detailed information about the storage array. The information covers several screens on a display. You might need to increase the size of your display buffer to see all the information. Because this information is so detailed, you might want to save the output of this parameter to a file.</p> <p>Use the following command to save the profile output to a file:</p> <pre>c:\...\smX\client&gt;smcli 123.45.67.88 123.45.67.89 -c "show storagearray profile;" -o "c:\folder\storagearray profile.txt"</pre>
autoSupportConfig	<p>The parameter to return information about the current state of the operation to automatically collect support data. The following information is returned:</p> <ul style="list-style-type: none"> <li>• Whether the operation is enabled or disabled</li> <li>• The location of the folder where the support data file is located</li> </ul>
batteryAge	<p>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. The information for both batteries is shown.</p>

Parameter	Description
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.
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 <code>longRunningOperation</code> parameter returns this information: <ul style="list-style-type: none"> <li>• Name of the volume group or volume</li> <li>• Long running operation</li> <li>• Status</li> <li>• % complete</li> <li>• Time left</li> </ul>
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
HOST TYPE                ALUA/AVT STATUS  ASSOCIATED INDEXS
AIX MPIO                  Disabled         9
AVT_4M                    Enabled          5
Factory Default          Disabled         0
HP-UX                     Enabled          15
Linux (ATTO)              Enabled          24
Linux (DM-MP)             Disabled         6
Linux (Pathmanager)      Enabled          25
Mac OS                    Enabled          22
ONTAP                     Disabled         4
SVC                       Enabled          18
Solaris (v11 or Later)   Enabled          17
Solaris (version 10 or earlier) Disabled         2
VMWare                    Enabled          10 (Default)
Windows                   Enabled          1
```

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
<code>asyncMirror</code>	Trial available
<code>syncMirror</code>	Trial available/Deactivated
<code>thinProvisioning</code>	Trial available
<code>driveSlotLimit</code>	Enabled (12 of 192 used)
<code>snapImage</code>	Enabled (0 of 512 used) - Trial version
<code>expires m/d/y</code>	
<code>snapshot</code>	Enabled (1 of 4 used)
<code>storagePartition</code>	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

Right.

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 REMAINING 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, the information returned has been expanded to show the status of the features in the storage array.

**Related references**

[Disable storage array feature](#) on page 197

[Enable storage array feature](#) on page 214

[Show storage array auto configuration](#) on page 414

[Show storage array DBM database](#) on page 420

[Show storage array host topology](#) on page 421

[Show storage array LUN mappings](#) on page 422

[Show storage array negotiation defaults](#) on page 423

[Show storage array ODX setting](#) on page 423

[Show storage array power information](#) on page 424

[Show storage array unconfigured iSCSI initiators](#) on page 424

[Show storage array unreadable sectors](#) on page 425

**Show storage array auto configuration**

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)
secureDrives=(fips | fde)]
```

**Parameters**

Parameter	Description
driveType	<p>The type of drives that you want to use for the storage array. 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"> <li>• fibre</li> <li>• SATA</li> <li>• SAS</li> </ul> <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"> <li>• HDD – Use this option when you have hard drives in the drive tray.</li> <li>• SSD – Use this option when you have solid state drives in the drive tray.</li> <li>• unknown – Use if you are not sure what types of drive media are in the drive tray.</li> <li>• allMedia – Use this option when you want to use all types of drive media that are in the drive tray.</li> </ul> <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>
hotSpareCount	<p>The number of hot spares that you want in the storage array. Use integer values.</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>
cacheReadPrefetch	<p>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> .</p>

Parameter	Description
<code>securityType</code>	<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"> <li>• <code>none</code> – The volume group and volumes are not secure.</li> <li>• <code>capable</code> – The volume group and volumes are capable of having security set, but security has not been enabled.</li> <li>• <code>enabled</code> – The volume group and volumes have security enabled.</li> </ul>
<code>secureDrives</code>	<p>The type of secure drives to use in the volume group. These settings are valid:</p> <ul style="list-style-type: none"> <li>• <code>fips</code> – To use FIPS compliant drives only.</li> <li>• <code>fde</code> – To use FDE compliant drives.</li> </ul> <p><b>Important:</b> Use this parameter along with the <code>securityType</code> parameter. If you specify <code>none</code> for the <code>securityType</code> parameter, the value of the <code>secureDrives</code> parameter is ignored, because non-secure volume groups do not need to have secure drive types specified.</p>

## 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 groups

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/drawer loss protection, refer to the following tables for additional criteria:

Level	Criteria for Tray Loss Protection	Minimum number of trays required
Disk Pool	The disk pool contains no more than two drives in a single tray.	6
RAID 6	The volume group contains no more than two drives in a single tray.	3
RAID 3 or RAID 5	Each drive in the volume group is located in a separate tray.	3
RAID 1	Each drive in a RAID 1 pair must be located in a separate tray.	2
RAID 0	Cannot achieve Tray Loss Protection.	Not applicable

Level	Criteria for Drawer Loss Protection	Minimum number of drawers required
Disk Pool	The disk pool includes drives from all five drawers and there are an equal number of drives in each drawer. A 60-drive tray can achieve Drawer Loss Protection when the disk pool contains 15, 20, 25, 30, 35, 40, 45, 50, 55, or 60 drives.	5
RAID 6	The volume group contains no more than two drives in a single drawer.	3
RAID 3 or RAID 5	Each drive in the volume group is located in a separate drawer.	3
RAID 1	Each drive in a mirrored pair must be located in a separate drawer.	2
RAID 0	Cannot achieve Drawer Loss Protection.	Not applicable

### Hot spares

With volume groups, a valuable strategy to protect data is to assign available drives in the storage array as hot spare drives. A hot spare is a drive, containing no data, that acts as a standby in the storage array in case a drive fails in a RAID 1, RAID 3, RAID 5, or RAID 6 volume group. The hot spare adds another level of redundancy to the storage array.

Generally, hot spare drives must have capacities that are equal to or greater than the used capacity on the drives that they are protecting. Hot spare drives must be of the same media type, the same interface type, and the same capacity as the drives that they are protecting.

If a drive fails in the storage array, the hot spare is normally substituted automatically for the failed drive without requiring your intervention. If a hot spare is available when a drive fails, the controller uses redundancy data parity to reconstruct the data onto the hot spare. Data evacuation support also allows data to be copied to a hot spare before the software marks the drive "failed."

After the failed drive is physically replaced, you can use either of the following options to restore the data:

When you have replaced the failed drive, the data from the hot spare is copied back to the replacement drive. This action is called copyback.

If you designate the hot spare drive as a permanent member of a volume group, the copyback operation is not needed.

The availability of tray loss protection and drawer loss protection for a volume group depends on the location of the drives that comprise the volume group. Tray loss protection and drawer loss protection might be lost because of a failed drive and the location of the hot spare drive. To make sure that tray loss protection and drawer loss protection are not affected, you must replace a failed drive to initiate the copyback process.

The storage array automatically selects Data Assurance (DA)-capable drives for hot spare coverage of DA-enabled volumes.

Make sure you have DA-capable drives in the storage array for hot spare coverage of DA-enabled volumes. For more information about DA-capable drives, refer to Data Assurance feature.

Secure-capable (FIPS and FDE) drives can be used as a hot spare for both secure-capable and non-secure-capable drives. Non-secure-capable drives can provide coverage for other non-secure-capable drives, and for secure-capable drives if the volume group does not have the security enabled. A FIPS volume group can only use a FIPS drive as a hot spare; however, you can use a FIPS hot spare for non-secure-capable, secure-capable, and secure-enabled volume groups.

If you do not have a hot spare, you can still replace a failed drive while the storage array is operating. If the drive is part of a RAID 1, RAID 3, RAID 5, or RAID 6 volume group, the controller uses redundancy data parity to automatically reconstruct the data onto the replacement drive. This action is called reconstruction.

### **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. Valid values for the `cacheReadPrefetch` parameter are `TRUE` or `FALSE`. The default is `TRUE`.

## Security type

Use the `securityType` parameter to specify the security settings for the storage array.

Before you can set the `securityType` parameter to `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`
- `export storageArray securityKey`
- `import storageArray securityKey`
- `set storageArray securityKey`
- `enable volumeGroup [volumeGroupName] security`
- `enable diskPool [diskPoolName] security`

## Secure drives

Secure-capable drives can be either Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives. Use the `secureDrives` parameter to specify the type of secure drives to use. The values you can use are `fips` and `fde`.

## Example command

```
show storageArray autoConfiguration securityType=capable
secureDrives=fips;
```

## Minimum firmware level

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

7.50 adds the `securityType` parameter.

7.75 adds the `dataAssurance` parameter.

8.25 adds the `secureDrives` parameter.

## Related references

[Autoconfigure storage array](#) on page 82

## Show storage array controller health image

**Note:** With firmware version 8.20 the `coreDump` parameter has been replaced with the `controllerHealthImage` parameter.

This command shows details of the controller health image on the controller cache, if a controller health image is available.

## Syntax

```
show storageArray controllerHealthImage
```

## Parameters

None.

**Minimum firmware level**

7.83

8.20 replaces the `coreDump` parameter with the `controllerHealthImage` parameter.**Related references**[Save storage array controller health image](#) on page 272[Set storage array controller health image allow overwrite](#) on page 348[Start storage array controller health image](#) on page 460**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	LOCATIO N	REVISIO N	ID	GEN NUMBER	STAT US	ACCESS MODE
A	Cache	X.Y	9999 99	9999999	R/W	Optimal
B	Cache	X.Y	9999 99	9999999	R/W	Optimal
N/A	Disk	X.Y	9999 99	9999999	R/W	Optimal

**Minimum firmware level**

7.83

**Related references**[Load storage array DBM database](#) on page 223[Save storage array DBM database](#) on page 273

[Save storage array DBM validator information file](#) on page 274

[Stop storage array configuration database diagnostic](#) on page 478

[Load storage array DBM database](#) on page 223

[Save storage array DBM database](#) on page 273

[Save storage array DBM validator information file](#) on page 274

## 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: 28:37:48:55:55:55:55:55
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
NVRAM HOST TYPE INDEX DEFINITIONS
HOST TYPE                ALUA/AVT STATUS  ASSOCIATED INDEXS
AIX MPIO                  Disabled         9
AVT_4M                    Enabled          5
Factory Default          Disabled         0
HP-UX                     Enabled          15
Linux (ATTO)              Enabled          24
Linux (DM-MP)             Disabled         6
Linux (Pathmanager)      Enabled          25
Mac OS                    Enabled          22
ONTAP                     Disabled         4
SVC                       Enabled          18
Solaris (v11 or Later)   Enabled          17
Solaris (version 10 or earlier) Disabled         2
VMWare                    Enabled          10 (Default)
Windows                   Enabled          1

```

### Minimum firmware level

5.20

**Related references**

[Create host](#) on page 127  
[Create host group](#) on page 128  
[Create host port](#) on page 129  
[Delete host](#) on page 182  
[Delete host group](#) on page 182  
[Delete host port](#) on page 183  
[Set host](#) on page 320  
[Set host channel](#) on page 321  
[Set host group](#) on page 322  
[Set host port label](#) on page 323  
[Show host ports](#) on page 398  
[Show storage array](#) on page 410  
[Show storage array LUN mappings](#) on page 422

**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

**Related references**

[Show storage array](#) on page 410

[Show storage array host topology](#) on page 421

## 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 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

### Related references

- [Activate iSCSI initiator](#) on page 72
- [Create iSCSI initiator](#) on page 131
- [Delete iSCSI initiator](#) on page 184
- [Diagnose controller iSCSI host cable](#) on page 193
- [Reset storage array iSCSI baseline](#) on page 252
- [Save storage array iSCSI statistics](#) on page 279
- [Set iSCSI initiator](#) on page 324
- [Set iSCSI target properties](#) on page 326
- [Show current iSCSI sessions](#) on page 390
- [Show storage array unconfigured iSCSI initiators](#) on page 424
- [Start iSCSI DHCP refresh](#) on page 453
- [Stop storage array iSCSI session](#) on page 479
- [Show storage array](#) on page 410

## Show storage array ODX setting

This command returns the current settings for Offloaded Data Transfer (ODX) and VMware vStorage API Array Architecture (VAAI) on the storage array. The settings can be one of these:

- True – ODX and VAAI are turned on.
- False – ODX and VAAI are turned off.
- Inconsistent – The controllers do not have the same settings.
- Unknown – The setting for ODX or VAAI cannot be determined.

### Syntax

```
show storageArray odxsetting
```

**Parameters**

None.

**Notes**

This command returns the ODX and VAAI information similar to this example.

```
Windows ODX Setting Status
odxEnabled True | False | Inconsistent | Unknown
vaaisEnabled True | False | Inconsistent | Unknown
```

**Minimum firmware level**

8.20

**Related references**

[Enable or disable ODX](#) on page 209

[Enable or disable VAAI](#) on page 214

**Show storage array power information**

This command returns information about the amount of power consumed by the entire storage array and each tray in the storage array.

**Syntax**

```
show storageArray powerInfo
```

**Parameters**

None.

**Notes**

This command returns power consumption information similar to this example.

```
total power drawn: 310 watts
number of trays: 1
tray power input details:
  tray id  power supply serial number  input power
    0      0                      160 watts
    0      1                      150 watts
```

**Minimum firmware level**

8.10

**Related references**

[Show storage array](#) on page 410

**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

## Related references

- [Activate iSCSI initiator](#) on page 72
- [Create iSCSI initiator](#) on page 131
- [Delete iSCSI initiator](#) on page 184
- [Diagnose controller iSCSI host cable](#) on page 193
- [Reset storage array iSCSI baseline](#) on page 252
- [Save storage array iSCSI statistics](#) on page 279
- [Set iSCSI initiator](#) on page 324
- [Set iSCSI target properties](#) on page 326
- [Show current iSCSI sessions](#) on page 390
- [Show storage array negotiation defaults](#) on page 423
- [Start iSCSI DHCP refresh](#) on page 453
- [Stop storage array iSCSI session](#) on page 479
- [Reset storage array iSCSI baseline](#) on page 252
- [Save storage array iSCSI statistics](#) on page 279
- [Stop storage array iSCSI session](#) on page 479

## 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 – 0xxxxxxxxx)
- Drive location (tray t, slot s)
- Drive-relative logical block address (hexadecimal format – 0xxxxxxxxx)
- 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
<code>primary</code>	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 (" ").
<code>remoteStorageArrayName</code>	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

**Related references**

- [Check synchronous mirroring status](#) on page 95
- [Create synchronous mirroring](#) on page 162
- [Diagnose synchronous mirroring](#) on page 196
- [Re-create synchronous mirroring repository volume](#) on page 230
- [Remove synchronous mirroring](#) on page 238
- [Set synchronous mirroring](#) on page 360

**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
volume	The name of the primary volume of the remote mirrored pair for which you want to check synchronization progress. Enclose the volume name in square brackets ([ ]). If the volume name has special characters or numbers, you must enclose the volume name in double quotation marks (" ") inside square brackets.
volumes	The names of the primary volume of the remote mirrored pair for which you want to check synchronization progress. Enter the names of the volumes using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Separate each of the names with a space.</li> </ul> If the volume names have special characters or consist only of numbers, enter the names using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Enclose each of the names in double quotation marks (" ").</li> <li>• Separate each of the names with a space.</li> </ul>

**Minimum firmware level**

5.40

**Related references**

- [Check synchronous mirroring status](#) on page 95
- [Diagnose synchronous mirroring](#) on page 196
- [Resume synchronous mirroring](#) on page 261

[Set synchronous mirroring](#) on page 360

[Start synchronous mirroring synchronization](#) on page 462

[Suspend synchronous mirroring](#) on page 482

## 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	This parameter returns information about all of the thin volumes in the storage array.
volume	The name of a thin volume for which you are retrieving information. Enclose the thin volume name in square brackets ([ ]). If the thin volume name has special characters or numbers, you must enclose the thin volume name in double quotation marks (" ") inside square brackets.
volumes	The names of several thin volumes for which you want to retrieve information. Enter the names of the volumes using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Separate each of the names with a space.</li> </ul> If the volume names have special characters or numbers, enter the names using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Enclose each of the names in double quotation marks (" ").</li> <li>• Separate each of the names with a space.</li> </ul>
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 Type	Start Capacity	End Capacity
MM/DD/YYYY HH:MM:SS	Manual   Automatic	NNNNNNNN bytes	NNNNNNNN bytes

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

### Related references

- [Initialize thin volume](#) on page 221
- [Set thin volume attributes](#) on page 362
- [Show volume](#) on page 429
- [Show volume action progress](#) on page 431
- [Show volume performance statistics](#) on page 435
- [Show volume reservations](#) on page 437

## 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 mirror repository volumes

To see 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	This parameter returns performance statistics about all of the volumes in the storage array.
volume	The name of a volume for which you want to retrieve information. Enclose the volume name in square brackets ([ ]). If the volume name has special characters or numbers, you must enclose the name in double quotation marks (" ") inside square brackets.
volumes	The names of several volumes for which you want to retrieve information. Enter the names of the volumes using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Separate each of the names with a space.</li> </ul> If the volume names have special characters or numbers, enter the names using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Enclose each of the names in double quotation marks (" ").</li> <li>• Separate each of the names with a space.</li> </ul>
summary	The setting to return a concise list of information about the volumes.

## Related references

[Show thin volume](#) on page 428

[Show volume action progress](#) on page 431

[Show volume performance statistics](#) on page 435

[Show volume reservations](#) on page 437

## 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 a volume for which you want to retrieve information about a long running operation. Enclose the volume name in square brackets ([ ]). If the volume name has special characters or numbers, you must enclose the name in double quotation marks (" ") inside square brackets.

### Minimum firmware level

5.43

### Related references

[Show thin volume](#) on page 428

[Show volume](#) on page 429

[Show volume performance statistics](#) on page 435

[Show volume reservations](#) on page 437

## 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 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 ([ ]).

## Related references

[Create volume copy](#) on page 165

[Set volume copy](#) on page 377

[Show volume copy source candidates](#) on page 432

[Show volume copy target candidates](#) on page 432

## 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 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
```

## 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 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 volume name in square brackets ([ ]). If the volume name has special characters or numbers, you must enclose the volume name in double quotation marks (" ") inside square brackets.

## Show volume group

This command returns this information about a volume group:

- The status (Online or Offline)
- The capacity
- The current owner (the controller in slot A or the controller in slot B )
- The RAID level
- The drive media type (HDD or SSD)
- The drive interface type (Fibre Channel, iSCSI, Infiniband, SAS)
- Tray loss protection (yes or no)
- Secure capable - indicates whether the volume group is composed of all secure-capable drives. Secure-capable drives can be either FDE drives or FIPS drives.
- Secure - Indicates whether the volume group has drive security turned on (this is called secure enabled).
- The associated volumes and free capacity
- The associated drives

## Syntax

```
show volumeGroup [volumeGroupName]
```

## Parameter

Parameter	Description
volumeGroup	The name of the volume group for which you want to show information. Enclose the volume group name in square brackets ([ ]).

## Notes

This command returns volume group information as shown in this example:

```

Name:                SecureGroup

Status:              Optimal
Capacity:            120.000 GB
Current owner:      Controller in slot A

Quality of Service (QoS) Attributes

RAID level:          5
Drive media type:    Hard Disk Drive
Drive interface type: Fibre Channel
Tray loss protection: No
Secure Capable:      No
Secure:              No

Total Volumes:       1
Standard volumes:    1
Repository volumes:  0
Free Capacity:       110.000 GB

Associated drives - present (in piece order)
Total drives present: 5

Tray  Slot
99    1
99    2
99    3
99    4
99    5

```

**Minimum firmware level**

6.10

**Related references**

[Show volume group export dependencies](#) on page 434

[Show volume group import dependencies](#) on page 435

**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 name of the volume group for which you want to show export dependencies. Enclose the volume group name 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

**Related references**[Show volume group import dependencies](#) on page 435[Start volume group export](#) on page 464[Start volume group import](#) on page 464**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 name of the volume group for which you want to show import dependencies. Enclose the volume group name 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 <code>TRUE</code> . To let the drives stay spinning, set this parameter to <code>FALSE</code> .

**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

**Related references**[Show volume group export dependencies](#) on page 434[Start volume group export](#) on page 464[Start volume group import](#) on page 464**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	The name of a volume for which you want to retrieve information. Enclose the volume name in square brackets ([ ]). If the volume name has special characters or numbers, you must enclose the name in double quotation marks (" ") inside square brackets.
volumes	The names of several volumes for which you want to retrieve information. Enter the names of the volumes using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Separate each of the names with a space.</li> </ul> If the volume names have special characters or numbers, enter the names using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Enclose each of the names in double quotation marks (" ").</li> <li>• Separate each of the names with a space.</li> </ul>

## 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", "0.0"
```

## Minimum firmware level

6.10

## Related references

- [Show thin volume](#) on page 428
- [Show volume](#) on page 429
- [Show volume action progress](#) on page 431

## 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 parameter to return performance statistics about all of the volumes in the storage array.
volume	The name of a volume for which you want to retrieve information. Enclose the volume name in square brackets ([ ]). If the volume name has special characters or numbers, you must enclose the name in double quotation marks (" ") inside square brackets.

Parameter	Description
volumes	<p>The names of several volumes for which you want to retrieve information. Enter the names of the volumes using these rules:</p> <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ( [ ] ).</li> <li>• Separate each of the names with a space.</li> </ul> <p>If the volume names have special characters or numbers, enter the names using these rules:</p> <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ( [ ] ).</li> <li>• Enclose each of the names in double quotation marks ( " " ).</li> <li>• Separate each of the names with a space.</li> </ul>

### Minimum firmware level

5.40

### Related references

[Show thin volume](#) on page 428

[Show volume](#) on page 429

## Specify the AutoSupport delivery method

The `SMcli -autoSupportConfig` command sets up the delivery method for sending AutoSupport messages.

### Description

You can create two major types of delivery methods for sending AutoSupport messages:

- **SMTP** - Use this method when you want to use email as the delivery method for sending AutoSupport messages.
- **HTTP/HTTPS** - Use this method when you want to send AutoSupport messages using the HTTP or HTTPS protocols. You need to use HTTPS if you want to use AutoSupport OnDemand or AutoSupport Remote Diagnostics.

### Context

**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.

**Restriction:** This command is obsolete for the E2800 storage array. Use the RESTful API, SANtricity System Manager, or see [Using curl commands with the E2800 storage array](#) on page 39.

### Syntax for SMTP Delivery Method

```
SMcli -autoSupportConfig deliveryMethod=SMTP
mailRelayServer=<mailRelayServerName> senderMail=<senderEmailAddress>
```

### Parameters for the SMTP Delivery Method

Parameter	Description
deliveryMethod	Allows you to specify the delivery method of the AutoSupport messages. Valid choice is SMTP. See the next section of this topic for information about using HTTP or HTTPS, which are also valid.
mailRelayServer	Allows you to specify the mail relay server for the AutoSupport messages.
senderMail	Allows you to specify the sender email address for the AutoSupport messages.

### Syntax for HTTP/HTTPS Delivery Methods

```
SMcli -autoSupportConfig deliveryMethod={HTTP/HTTPS}
      {direct|proxyConfigScript=<proxyConfigScript>| proxyServer
      hostAddress:<address>portNumber=<portNumber>
      [userName=<userName>password=<password>]}
```

### Parameters for the HTTP/HTTPS Delivery Methods

The following table lists the parameters for the HTTP/HTTPS delivery methods.

Parameter	Description
deliveryMethod	Allows you to specify the delivery method of the AutoSupport messages. Valid choices are HTTPS and HTTP (see prior section for use of SMTP, which is also valid).
direct	Allows you to connect directly to the destination technical support systems using the HTTPS or HTTP protocols.
proxyConfigScript	Allows you to specify the location of a Proxy Auto-Configuration (PAC) script file.
proxyServer	Allows you to specify the HTTP(s) proxy server details required for establishing connection with the destination technical support system.
hostAddress	The host IP address of the proxyServer.
portNumber	The port number of the proxy server. This parameter is required when the proxyServer parameter is used.
userName	The user name. This parameter is required.
password	The authentication password. This parameter is required.

### Minimum Firmware Level

8.25

### Examples

```
SMcli -autoSupportConfig deliveryMethod=SMTP
mailRelayServer:server.eng.netapp.com
senderEmail:user@netapp.com
```

```
SMcli -autoSupportConfig deliveryMethod=HTTPS
proxyServer hostAddress:10.117.12.112 portNumber=8080
userName=user password=0987654321
```

## Verification

Send a test message using the `SMcli -autoSupportConfig test` command to verify that your delivery methods are set up correctly.

## Related references

[Using curl commands with the E2800 storage array](#) on page 39

[Capture or view an AutoSupport log](#) on page 89

[Display AutoSupport bundle collection settings](#) on page 198

[Display AutoSupport message collection schedule](#) on page 200

[Disable AutoSupport at the EMW level SMcli version](#) on page 197

[Enable or disable AutoSupport at the EMW management domain level \(SMcli only\)](#) on page 210

[Enable or disable the AutoSupport OnDemand feature at the EMW management domain level \(SMcli only\)](#) on page 211

[Enable or disable the AutoSupport OnDemand Remote Diagnostics feature at the EMW management domain level \(SMcli only\)](#) on page 213

[Reset AutoSupport message collection schedule](#) on page 247

[Save storage array support data](#) on page 283

[Test the AutoSupport configuration](#) on page 485

## Start asynchronous mirroring synchronization

This command starts Asynchronous Mirroring synchronization.

## Syntax

```
start asyncMirrorGroup ["asyncMirrorGroupName"] synchronize
[deleteRecoveryPointIfNecessary]
```

## 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 double quotation marks (" ") inside square brackets ([ " "]).
deleteRecoveryPointIfNecessary	The parameter to delete the recovery point if the recoverable synchronization data has exceeded time threshold for recovery. Recovery point age is measured from the time the data was frozen on the primary storage array.

## Minimum firmware level

7.84

8.10 adds the `deleteRecoveryPointIfNecessary` parameter.



## Related references

- [Resume asynchronous mirror group](#) on page 255
- [Show asynchronous mirror group synchronization progress](#) on page 383
- [Suspend asynchronous mirror group](#) on page 481
- [Test asynchronous mirror group connectivity](#) on page 484

## Specify the AutoSupport delivery method

The `SMcli -autoSupportConfig` command sets up the delivery method for sending AutoSupport messages.

### Description

You can create two major types of delivery methods for sending AutoSupport messages:

- SMTP - Use this method when you want to use email as the delivery method for sending AutoSupport messages.
- HTTP/HTTPS - Use this method when you want to send AutoSupport messages using the HTTP or HTTPS protocols. You need to use HTTPS if you want to use AutoSupport OnDemand or AutoSupport Remote Diagnostics.

### Context

**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.

**Restriction:** This command is obsolete for the E2800 storage array. Use the RESTful API, SANtricity System Manager, or see [Using curl commands with the E2800 storage array](#) on page 39.

### Syntax for SMTP Delivery Method

```
SMcli -autoSupportConfig deliveryMethod=SMTP
mailRelayServer=<mailRelayServerName> senderMail=<senderEmailAddress>
```

### Parameters for the SMTP Delivery Method

Parameter	Description
deliveryMethod	Allows you to specify the delivery method of the AutoSupport messages. Valid choice is SMTP. See the next section of this topic for information about using HTTP or HTTPS, which are also valid.
mailRelayServer	Allows you to specify the mail relay server for the AutoSupport messages.
senderMail	Allows you to specify the sender email address for the AutoSupport messages.

### Syntax for HTTP/HTTPS Delivery Methods

```
SMcli -autoSupportConfig deliveryMethod={HTTP/HTTPS}
{direct|proxyConfigScript=<proxyConfigScript>| proxyServer
hostAddress:<address>portNumber=<portNumber>
[user=<userName>password=<password>]}
```

## Parameters for the HTTP/HTTPS Delivery Methods

The following table lists the parameters for the HTTP/HTTPS delivery methods.

Parameter	Description
deliveryMethod	Allows you to specify the delivery method of the AutoSupport messages. Valid choices are HTTPS and HTTP (see prior section for use of SMTP, which is also valid).
direct	Allows you to connect directly to the destination technical support systems using the HTTPS or HTTP protocols.
proxyConfigScript	Allows you to specify the location of a Proxy Auto-Configuration (PAC) script file.
proxyServer	Allows you to specify the HTTP(s) proxy server details required for establishing connection with the destination technical support system.
hostAddress	The host IP address of the proxyServer.
portNumber	The port number of the proxy server. This parameter is required when the proxyServer parameter is used.
userName	The user name. This parameter is required.
password	The authentication password. This parameter is required.

### Minimum Firmware Level

8.25

### Examples

```
SMcli -autoSupportConfig deliveryMethod=SMTP
mailRelayServer:server.eng.netapp.com
senderEmail:user@netapp.com
```

```
SMcli -autoSupportConfig deliveryMethod=HTTPS
proxyServer hostAddress:10.117.12.112 portNumber=8080
userName=user password=0987654321
```

### Verification

Send a test message using the `SMcli -autoSupportConfig test` command to verify that your delivery methods are set up correctly.

### Related references

[Using curl commands with the E2800 storage array](#) on page 39

[Capture or view an AutoSupport log](#) on page 89

[Display AutoSupport bundle collection settings](#) on page 198

[Display AutoSupport message collection schedule](#) on page 200

[Disable AutoSupport at the EMW level SMcli version](#) on page 197

[Enable or disable AutoSupport at the EMW management domain level \(SMcli only\)](#) on page 210

[Enable or disable the AutoSupport OnDemand feature at the EMW management domain level \(SMcli only\)](#) on page 211

[Enable or disable the AutoSupport OnDemand Remote Diagnostics feature at the EMW management domain level \(SMcli only\)](#) on page 213

[Reset AutoSupport message collection schedule](#) on page 247

[Save storage array support data](#) on page 283

[Test the AutoSupport configuration](#) on page 485

## Start consistency group snapshot rollback

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
<code>cgSnapImage</code>	<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"> <li>The name of the snapshot group</li> <li>An identifier for the snapshot image in the snapshot group.</li> </ul> <p>The identifier for the snapshot image can be one of these:</p> <ul style="list-style-type: none"> <li>An integer value that is the sequence number of the snapshot in the snapshot group.</li> <li><code>NEWEST</code> - Use this option when you want to show the latest snapshot image created in the snapshot group.</li> <li><code>OLDEST</code> - Use this option when you want to show the earliest snapshot image created in the snapshot group.</li> </ul> <p>Enclose the snapshot image name in double quotation marks ( " ") inside square brackets ( [ ] ).</p>
<code>memberVolumeSet</code>	<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 space.</p> <p>When the <code>memberVolumeSet</code> 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

**Related references**

- [Add member to consistency group](#) on page 78
- [Create consistency group](#) on page 107
- [Create consistency group snapshot image](#) on page 115
- [Create consistency group snapshot volume](#) on page 116
- [Create consistency group snapshot volume mapping](#) on page 120
- [Delete consistency group](#) on page 178
- [Delete consistency group snapshot image](#) on page 179
- [Delete consistency group snapshot volume](#) on page 180
- [Remove member volume from consistency group](#) on page 237
- [Resume consistency group snapshot volume](#) on page 256
- [Set consistency group attributes](#) on page 296
- [Set consistency group snapshot volume](#) on page 298
- [Show consistency group](#) on page 384
- [Show consistency group snapshot image](#) on page 386
- [Stop consistency group snapshot rollback](#) on page 466
- [Stop consistency group snapshot volume](#) on page 468
- [Stop pending snapshot images on consistency group](#) on page 471

**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 technical support 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)
file="fileName"
```

## 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"> <li>current – Retrieves the current DQ traces</li> <li>flushed – Retrieves all flushed DQ traces</li> <li>currentFlushed – Retrieves both the current DQ trace and the flushed DQ trace</li> <li>all – Retrieves the current DQ trace, flushed DQ trace, and all platform DQ traces</li> </ul> <p><b>NOTE</b> – 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><b>NOTE</b> – 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. Enclose the file name in double quotation marks (" "). 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

### Related references

- [Clear drive channel statistics](#) on page 98
- [Diagnose controller iSCSI host cable](#) on page 193
- [Enable controller data transfer](#) on page 208
- [Reset controller](#) on page 248
- [Save controller NVSRAM](#) on page 267
- [Save drive channel fault isolation diagnostic status](#) on page 267
- [Set controller](#) on page 299
- [Set controller DNS settings](#) on page 304
- [Set controller NTP settings](#) on page 305
- [Set controller service action allowed indicator](#) on page 306
- [Set drive channel status](#) on page 313
- [Set host channel](#) on page 321
- [Show controller](#) on page 387
- [Show controller NVSRAM](#) on page 390
- [Show drive channel statistics](#) on page 395
- [Start drive channel fault isolation diagnostics](#) on page 447
- [Start drive channel locate](#) on page 449
- [Start iSCSI DHCP refresh](#) on page 453
- [Stop drive channel fault isolation diagnostics](#) on page 469
- [Stop drive channel locate](#) on page 470

## 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 for which you want to locate. Enclose the disk pool name in square brackets ([ ]). If the disk pool name has special characters or numbers, you must enclose the disk pool name in double quotation marks (" ") inside square brackets.

**Minimum firmware level**

7.83

**Related references**

[Stop disk pool locate](#) on page 469

**Start drive channel fault isolation diagnostics**

**Note:** With firmware version 8.10, the `start driveChannel faultDiagnostics` command is deprecated.

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. Enclose the controller identifier in square brackets ([ ]).

Parameter	Description
testDevices	<p>The identifiers of the devices (controllers, environmental services module [ESMs], or drives) that you want to test. You can specify <code>all</code> or enter the specific identifiers for the devices that you want to diagnose.</p> <ul style="list-style-type: none"> <li><code>controller</code> – Identifiers are <code>a</code> or <code>b</code>, where <code>a</code> is the RAID controller module in slot A, and <code>b</code> is the RAID controller module in slot B.</li> <li><code>esms</code> – 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.</li> <li><code>drives</code> – Identifiers are tray ID and slot ID, where tray ID values are 0 to 99 and slot ID values are 0 to 31.</li> </ul>
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.
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

You can enter more than one type of device to test, and you can enter more than one type of test to run.

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 also can stop this command at any time by pressing **Ctrl+C**.

## Minimum firmware level

7.15

## Related references

[Save drive channel fault isolation diagnostic status](#) on page 267

[Stop drive channel fault isolation diagnostics](#) on page 469

[Download drive firmware](#) on page 202

[Replace drive](#) on page 244

[Revive drive](#) on page 262

[Save drive channel fault isolation diagnostic status](#) on page 267

[Save drive log](#) on page 268

[Set drive hot spare](#) on page 314

[Set drive service action allowed indicator](#) on page 314

[Set drive state](#) on page 315



[Set FIPS drive security identifier](#) on page 318  
[Set foreign drive to native](#) on page 319  
[Show drive](#) on page 393  
[Show drive download progress](#) on page 396  
[Show drive performance statistics](#) on page 396  
[Start drive initialize](#) on page 449  
[Start drive locate](#) on page 450  
[Start drive reconstruction](#) on page 451  
[Start FDE secure drive erase](#) on page 455  
[Stop drive channel fault isolation diagnostics](#) on page 469  
[Stop drive locate](#) on page 470

## 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.

### Related references

[Stop drive channel locate](#) on page 470

## Start drive initialize

This command starts a 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 initialize. 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 24 . 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.**Related references**

[Download drive firmware](#) on page 202

[Replace drive](#) on page 244

[Revive drive](#) on page 262

[Save drive channel fault isolation diagnostic status](#) on page 267

[Save drive log](#) on page 268

[Set drive hot spare](#) on page 314

[Set drive service action allowed indicator](#) on page 314

[Set drive state](#) on page 315

[Set FIPS drive security identifier](#) on page 318

[Set foreign drive to native](#) on page 319

[Show drive](#) on page 393

[Show drive download progress](#) on page 396

[Show drive performance statistics](#) on page 396

[Start drive channel fault isolation diagnostics](#) on page 447

[Start drive locate](#) on page 450

[Start drive reconstruction](#) on page 451

[Start FDE secure drive erase](#) on page 455

[Stop drive channel fault isolation diagnostics](#) on page 469

[Stop drive locate](#) on page 470

**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 24 . 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.

## Related references

[Stop drive locate](#) on page 470

## 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 24 . 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.**Related references**

[Download drive firmware](#) on page 202  
[Replace drive](#) on page 244  
[Revive drive](#) on page 262  
[Save drive channel fault isolation diagnostic status](#) on page 267  
[Save drive log](#) on page 268  
[Set drive hot spare](#) on page 314  
[Set drive service action allowed indicator](#) on page 314  
[Set drive state](#) on page 315  
[Set FIPS drive security identifier](#) on page 318  
[Set foreign drive to native](#) on page 319  
[Show drive](#) on page 393  
[Show drive download progress](#) on page 396  
[Show drive performance statistics](#) on page 396  
[Start drive channel fault isolation diagnostics](#) on page 447  
[Start drive initialize](#) on page 449  
[Start drive locate](#) on page 450  
[Start FDE secure drive erase](#) on page 455  
[Stop drive channel fault isolation diagnostics](#) on page 469  
[Stop drive locate](#) on page 470

**Start input output controller (IOC) dump**

This command produces a dump of the IOC log of data transmissions between a host and a controller. Controllers are identified as "a " or "b ". Each controller can have up to four host channels; each channel is uniquely identified by a number from 1 through 4.

**Syntax**

```
start IOCLog [(a1 | a2 | a3 | a4 | b1 | b2 | b3 | b4)]
[overwrite=(TRUE | FALSE)]
```

## Parameters

Parameter	Description
controller-channel identifiers	<p>This parameter specifies the controller and host channel from which to produce the IOC dump. Valid controller identifiers are a or b , where a is the controller in slot A , and b is the controller in slot B . Host channels have numerical identifiers. Enclose the controller identifier and host channel identifier in square brackets ([ ]).</p> <p>Valid controller identifier and host channel values are a1 , a2 , a3 , a4 , b1 , b2 , b3 , b4 .</p> <p>If you do not specify a controller, the storage management software returns a syntax error.</p>
overwrite	<p>This parameter causes the new IOC log dump to overwrite an existing dump. To overwrite the existing dump, set this parameter to TRUE . The default value is FALSE .</p>

## Notes

This command generates a debug log from the IOC of the selected controller and stores the data in a compressed format in a persistent memory buffer on the controller. You can retrieve the data from the debug log using the `save IOCLog` command. The controller returns an error for these conditions:

- The controller platform and HIC do not support an IOC dump.
- The specified controller has an outstanding IOC dump and the overwrite parameter is false.
- The specified controller identifier or channel identifier are out of the valid range.

## Minimum firmware level

8.20

## Related references

[Save input output controller \(IOC\) dump](#) on page 269

## 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 [portLabel] dhcpRefresh
```

## Parameter

Parameter	Description
controller	<p>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.</p>

Parameter	Description
<code>iscsiHostPort</code>	The host port label or the number of the iSCSI host port for which you want to refresh the DHCP parameters. See the following for more information: <a href="#">"Identifying an iSCSI host port label"</a>

### Identifying an iSCSI host port label

You must specify a label for the host port. Follow these steps to specify the host port label:

1. If you do not know the port label for the iSCSI host port, run the `show controller` command.
2. In the Host interface section of the results, locate the host port you want to select.

**Note:** The port label is the complete value returned for the `Port` field.

3. Enclose the entire value of the port label in both quotes and square brackets: `["portLabel"]`. For example, if the port label is `ch 2`, specify the iSCSI host port as follows:

```
iscsiHostPort["ch 2"]
```

**Note:** If you are using a Windows command line and the label contains a pipe (`|`), the character should be escaped (using `^`); otherwise, it will be interpreted as a command. For example, if the port label is `e0b|0b`, specify the iSCSI host port as follows:

```
iscsiHostPort["e0b^|0b"]
```

**Note:** For backward compatibility, the `iscsiPortNumber`, enclosed by braces `[ ]` rather than quotes and braces `[" "]` can still be used for E2700, E5600, or EF560 controllers (and other previous generations of E-Series or EF-Series controllers). For those controllers, valid values for `iscsiPortNumber` are as follows:

- For controllers with integrated host ports, the numbering is 3, 4, 5, or 6.
- For controllers with host ports on a host interface card only, the numbering is 1, 2, 3, or 4.

An example of the prior syntax is as follows:

```
iscsiHostPort[3]
```

### Notes

This operation ends the iSCSI connections for the portal and temporarily brings down the portal.

### Minimum firmware level

7.10

8.10 revises the numbering system for iSCSI host ports.

8.30 revises the identification method for iSCSI host ports in the E2800.

### Related references

[Set controller DNS settings](#) on page 304

[Set controller NTP settings](#) on page 305

[Activate iSCSI initiator](#) on page 72  
[Create iSCSI initiator](#) on page 131  
[Delete iSCSI initiator](#) on page 184  
[Diagnose controller iSCSI host cable](#) on page 193  
[Reset storage array iSCSI baseline](#) on page 252  
[Save storage array iSCSI statistics](#) on page 279  
[Set iSCSI initiator](#) on page 324  
[Set iSCSI target properties](#) on page 326  
[Show current iSCSI sessions](#) on page 390  
[Show storage array negotiation defaults](#) on page 423  
[Show storage array unconfigured iSCSI initiators](#) on page 424  
[Stop storage array iSCSI session](#) on page 479

## Start FDE 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.

**Note:** To erase a FIPS drive when the drive is locked and the security key to unlock it is unavailable, use the `set drive securityID` command.

### Syntax

```
start secureErase (drive [trayID,drawerID,slotID] |
drives [trayID1, drawerID1, slotID1 ... trayIDn,drawerIDn,slotIDn])
```

### Parameters

Parameter	Description
drive or drives	<p>The tray and the slot where the drive resides. 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 24 . Enclose the tray ID value, the drawer ID value, and the slot ID value in square brackets ([ ]).</p> <p><b>Attention:</b> The drives list cannot contain all of the drives in the storage array, or the command will be rejected. To secure erase all drives, run this command twice, specifying the drive list in two separate groups.</p>

### 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

**Related references**

[Set FIPS drive security identifier](#) on page 318

**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"> <li>The name of the snapshot group</li> <li>An identifier for the snapshot image in the snapshot group</li> </ul> <p>The identifier for the snapshot image can be one of these:</p> <ul style="list-style-type: none"> <li>An integer value that is the sequence number of the snapshot in the snapshot group.</li> <li>NEWEST - Use this option when you want to show the latest snapshot image created in the snapshot group.</li> <li>OLDEST - Use this option when you want to show the earliest snapshot image created in the snapshot group.</li> </ul> <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

### Related references

- [Create snapshot image](#) on page 155
- [Delete snapshot image](#) on page 186
- [Resume snapshot image rollback](#) on page 258
- [Show snapshot image](#) on page 401
- [Stop snapshot image rollback](#) on page 473

## Start SSD cache locate

This command identifies the Solid State Disks (SSDs) that are logically grouped together to form the SSD cache by blinking the indicator lights on the SSDs. (Use the `stop ssdCache locate` command to turn off the indicator lights on the drives.)

### Syntax

```
start ssdCache [ssdCacheName] locate
```

### Parameters

Parameter	Description
<code>ssdCache</code>	The name of the SSD cache that you want to locate. Enclose the SSD cache name in square brackets ([ ]). If the SSD cache name contains special characters or consists only of numbers, you must enclose the SSD cache name in double quotation marks (" ") inside square brackets.

### Minimum firmware level

7.84

### Related references

- [Add drives to SSD cache](#) on page 77
- [Change SSD cache application type](#) on page 90
- [Create SSD cache](#) on page 159
- [Delete SSD cache](#) on page 188
- [Remove drives from SSD cache](#) on page 235
- [Rename SSD cache](#) on page 242
- [Resume SSD cache](#) on page 260
- [Set SSD cache for a volume](#) on page 342
- [Show SSD cache](#) on page 406
- [Show SSD cache statistics](#) on page 407
- [Start SSD cache performance modeling](#) on page 458

[Stop SSD cache locate](#) on page 475

[Stop SSD cache performance modeling](#) on page 476

[Suspend SSD cache](#) on page 482

## 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 name of the SSD cache for which you want to model performance. Enclose the SSD cache name in square brackets ([ ]). If the SSD cache name contains special characters or consists only of numbers, you must enclose the SSD cache name 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 the storage management software.

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

### Related references

[Add drives to SSD cache](#) on page 77

[Change SSD cache application type](#) on page 90

[Create SSD cache](#) on page 159

[Delete SSD cache](#) on page 188

[Remove drives from SSD cache](#) on page 235

[Rename SSD cache](#) on page 242

[Resume SSD cache](#) on page 260

[Set SSD cache for a volume](#) on page 342

- [Show SSD cache](#) on page 406
- [Show SSD cache statistics](#) on page 407
- [Start SSD cache locate](#) on page 457
- [Stop SSD cache locate](#) on page 475
- [Stop SSD cache performance modeling](#) on page 476
- [Suspend SSD cache](#) on page 482

## 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"> <li>• <code>disk</code> indicates that data comes directly from the database on the drive</li> <li>• <code>onboard</code> indicates that data comes from the RPA memory location</li> </ul> <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 checks the structural integrity of the database.</p> <p><code>mirror</code> – The tests run using this option vary according to the value of the <code>sourceLocation</code> parameter:</p> <ul style="list-style-type: none"> <li>• When the <code>sourceLocation</code> parameter is set to <code>disk</code>, the peer controller initiates a block check.</li> <li>• When the <code>sourceLocation</code> parameter is set to <code>onboard</code>, the peer controller initiates a record check.</li> </ul> <p><b>Note:</b> 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> <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 technical support.

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`

## Related references

[Stop storage array configuration database diagnostic](#) on page 478

## Start storage array controller health image

**Note:** With firmware version 8.20 the `coreDump` parameter is obsoleted. Use the `controllerHealthImage` parameter.

**Important:** Use this command only under the direction of technical support.

This command produces a storage array controller health image.

## Syntax

```
start storageArray controllerHealthImage controller [(a|b)]
```

## Parameters

Parameter	Description
controller	This parameter specifies the controller from which to produce the controller health image. 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 controllerHealthImage` command to save a copy of the cache data to a host file.

## Minimum firmware level

7.83

8.20 replaces the `coreDump` parameter with the `controllerHealthImage` parameter.

## Related references

[Save storage array controller health image](#) on page 272

[Set storage array controller health image allow overwrite](#) on page 348

[Show storage array controller health image](#) on page 419

## 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

**Related references**[Set storage array](#) on page 344[Set storage array iSNS registration](#) on page 351[Set storage array iSNS server IPv4 address](#) on page 353[Set storage array iSNS server IPv6 address](#) on page 354[Set storage array iSNS server listening port](#) on page 354[Set storage array iSNS server refresh](#) on page 355**Start storage array locate**

This command locates a storage array 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

**Related references**[Stop storage array locate](#) on page 479**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

**Related references**

[Check synchronous mirroring status](#) on page 95

[Remove synchronous mirroring](#) on page 238

[Resume synchronous mirroring](#) on page 261

[Show synchronous mirroring volume synchronization progress](#) on page 427

[Suspend synchronous mirroring](#) on page 482

**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

**Related references**

[Stop tray locate](#) on page 480

**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 name of the volume group that you want to defragment. Enclose the volume group name 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

**Related references**[Show volume group](#) on page 433**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 features from one storage array to another storage array.

**Syntax**

```
start volumeGroup [volumeGroupName] export
```

**Parameter**

Parameter	Description
volumeGroup	The name of the volume group that you want to export. Enclose the volume group name 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

**Related references**[Show volume group export dependencies](#) on page 434[Show volume group import dependencies](#) on page 435[Start volume group import](#) on page 464**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 features from one storage array to another storage array.

## Syntax

```
start volumeGroup [volumeGroupName] import
```

## Parameter

Parameter	Description
volumeGroup	The name of the volume group that you want to import. Enclose the volume group name in square brackets ([ ]).

## Notes

Higher-level volumes that are specifically related to features (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

## Related references

[Show volume group export dependencies](#) on page 434

[Show volume group import dependencies](#) on page 435

[Start volume group export](#) on page 464

## 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 name of the volume group for which you want to locate the drives that belong to that volume group. Enclose the volume group name in square brackets ([ ]).

## Minimum firmware level

6.16

## Related references

[Stop volume group locate](#) on page 481

## 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 a volume that you want to format. Enclose the volume name in square brackets ([ ]). If the volume name has special characters or numbers, you must enclose the name in double quotation marks (" ") inside square brackets.

### Minimum firmware level

6.10

### Related references

[Initialize thin volume](#) on page 221

[Show volume action progress](#) on page 431

## 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"> <li>• The name of the snapshot group</li> <li>• An identifier for the snapshot image in the snapshot group</li> </ul> <p>The identifier for the snapshot image can be one of these:</p> <ul style="list-style-type: none"> <li>• An integer value that is the sequence number of the snapshot in the snapshot group.</li> <li>• <b>NEWEST</b> - Use this option when you want to show the latest snapshot image created in the snapshot group.</li> <li>• <b>OLDEST</b> - Use this option when you want to show the earliest snapshot image created in the snapshot group.</li> </ul> <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 space.</p> <p>When the <code>memberVolumeSet</code> 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

**Related references**

[Add member to consistency group](#) on page 78  
[Create consistency group](#) on page 107  
[Create consistency group snapshot image](#) on page 115  
[Create consistency group snapshot volume](#) on page 116  
[Create consistency group snapshot volume mapping](#) on page 120  
[Delete consistency group](#) on page 178  
[Delete consistency group snapshot image](#) on page 179  
[Delete consistency group snapshot volume](#) on page 180  
[Remove member volume from consistency group](#) on page 237  
[Resume consistency group snapshot volume](#) on page 256  
[Set consistency group attributes](#) on page 296  
[Set consistency group snapshot volume](#) on page 298  
[Show consistency group](#) on page 384  
[Show consistency group snapshot image](#) on page 386  
[Start consistency group snapshot rollback](#) on page 443  
[Stop consistency group snapshot volume](#) on page 468  
[Stop pending snapshot images on consistency group](#) on page 471

**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
cgSnapVolume	The name of the consistency group snapshot volume that you want to stop creating. Enclose the volume name in square brackets ([ ]). If the volume name has special characters or numbers, you must enclose the name in double quotation marks (" ") inside square brackets.

**Minimum firmware level**

7.83

**Related references**

[Add member to consistency group](#) on page 78  
[Create consistency group](#) on page 107  
[Create consistency group snapshot image](#) on page 115  
[Create consistency group snapshot volume](#) on page 116

- [Create consistency group snapshot volume mapping](#) on page 120
- [Delete consistency group](#) on page 178
- [Delete consistency group snapshot image](#) on page 179
- [Delete consistency group snapshot volume](#) on page 180
- [Remove member volume from consistency group](#) on page 237
- [Resume consistency group snapshot volume](#) on page 256
- [Set consistency group attributes](#) on page 296
- [Set consistency group snapshot volume](#) on page 298
- [Show consistency group](#) on page 384
- [Show consistency group snapshot image](#) on page 386
- [Start consistency group snapshot rollback](#) on page 443
- [Stop consistency group snapshot rollback](#) on page 466
- [Stop pending snapshot images on consistency group](#) on page 471

## 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

### Related references

- [Start disk pool locate](#) on page 446

## Stop drive channel fault isolation diagnostics

**Note:** With firmware version 8.10, the `stop driveChannel faultDiagnostics` command is deprecated.

This command stops the drive channel fault isolation diagnostics, which stops the `start drive channel 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 also can stop the `start driveChannel faultDiagnostics` command at any time by pressing **Ctrl+C**.

**Minimum firmware level**

7.15

**Related references**

[Save drive channel fault isolation diagnostic status](#) on page 267

[Start drive channel fault isolation diagnostics](#) on page 447

[Start drive channel fault isolation diagnostics](#) on page 447

**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

**Related references**

[Start drive channel locate](#) on page 449

**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

**Related references**

[Start drive locate](#) on page 450

**Stop drive replace**

This command stops a data copy operation that was started by either a `set drive` command or a `replace drive` command with the `copyDrive` user input.

**Syntax**

```
stop drive replace ([trayID,drawerID,slotID] | <"wwID">)
```

**Parameters**

Parameter	Description
drive	The location of the drive that is being copied before you replace it. 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 24 . Enclose the tray ID value, the drawer ID value, and the slot ID value in square brackets ( [ ] ).
drive	The World Wide Identifier (WWID) of the drive that you want to replace. Enclose the WWID in double quotation marks ( " ") inside angle 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**

8.20

**Related references**

[Replace drive](#) on page 244

[Reset drive](#) on page 249

[Set drive state](#) on page 315

**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
consistencyGroupName	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 ([ ]). If the consistency group name has special characters or numbers, you must enclose the consistency group name in double quotation marks (" ") inside square brackets.

**Minimum firmware level**

7.83

**Related references**

- [Add member to consistency group](#) on page 78
- [Create consistency group](#) on page 107
- [Create consistency group snapshot image](#) on page 115
- [Create consistency group snapshot volume](#) on page 116
- [Create consistency group snapshot volume mapping](#) on page 120
- [Delete consistency group](#) on page 178
- [Delete consistency group snapshot image](#) on page 179
- [Delete consistency group snapshot volume](#) on page 180
- [Remove member volume from consistency group](#) on page 237
- [Resume consistency group snapshot volume](#) on page 256
- [Set consistency group attributes](#) on page 296
- [Set consistency group snapshot volume](#) on page 298
- [Show consistency group](#) on page 384
- [Show consistency group snapshot image](#) on page 386
- [Start consistency group snapshot rollback](#) on page 443
- [Stop consistency group snapshot rollback](#) on page 466
- [Stop consistency group snapshot volume](#) on page 468

**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
snapGroup	The name of the snapshot group for which you want to stop pending snapshot images. Enclose the snapshot group name in square brackets ([ ]). If the snapshot group name has special characters or numbers, you must enclose the snapshot group name in double quotation marks (" ") inside square brackets.
consistencyGroup	The name of the snapshot consistency group for which you want to stop pending snapshot images. Enclose the snapshot consistency group name in square brackets ([ ]). If the snapshot consistency group name has special characters or numbers, you must enclose the snapshot group name in double quotation marks (" ") inside square brackets.

## Minimum firmware level

7.83

## Related references

[Create snapshot group](#) on page 146

[Delete snapshot group](#) on page 185

[Revive snapshot group](#) on page 263

[Set snapshot group attributes](#) on page 331

[Set snapshot group media scan](#) on page 333

[Set snapshot group repository volume capacity](#) on page 334

[Set snapshot group schedule](#) on page 336

[Set SSD cache for a volume](#) on page 342

[Show snapshot group](#) on page 399

## 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"> <li>• The name of the snapshot group</li> <li>• An identifier for the snapshot image in the snapshot group</li> </ul> <p>The identifier for the snapshot image can be one of these:</p> <ul style="list-style-type: none"> <li>• An integer value that is the sequence number of the snapshot in the snapshot group.</li> <li>• <b>NEWEST</b> - Use this option when you want to show the latest snapshot image created in the snapshot group.</li> <li>• <b>OLDEST</b> - Use this option when you want to show the earliest snapshot image created in the snapshot group.</li> </ul> <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

**Related references**

- [Create snapshot image](#) on page 155
- [Delete snapshot image](#) on page 186
- [Resume snapshot image rollback](#) on page 258
- [Show snapshot image](#) on page 401
- [Start snapshot image rollback](#) on page 456

**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

## Related references

- [Create read-only snapshot volume](#) on page 145
- [Create snapshot volume](#) on page 156
- [Delete snapshot volume](#) on page 187
- [Rename snapshot volume](#) on page 242
- [Resume snapshot volume](#) on page 259
- [Revive snapshot volume](#) on page 263
- [Set read-only snapshot volume to read/write volume](#) on page 328
- [Set snapshot volume media scan](#) on page 339
- [Set snapshot volume repository volume capacity](#) on page 340
- [Set SSD cache for a volume](#) on page 342
- [Show snapshot volumes](#) on page 403

## Stop SSD cache locate

This command turns off the indicator lights on the Solid State Disks (SSDs) that were turned on by the `start ssdCache locate` command.

## Syntax

```
stop ssdCache locate
```

## Parameters

None.

## Minimum firmware level

7.84

## Related references

- [Add drives to SSD cache](#) on page 77
- [Change SSD cache application type](#) on page 90

[Create SSD cache](#) on page 159  
[Delete SSD cache](#) on page 188  
[Remove drives from SSD cache](#) on page 235  
[Rename SSD cache](#) on page 242  
[Resume SSD cache](#) on page 260  
[Set SSD cache for a volume](#) on page 342  
[Show SSD cache](#) on page 406  
[Show SSD cache statistics](#) on page 407  
[Start SSD cache locate](#) on page 457  
[Start SSD cache performance modeling](#) on page 458  
[Stop SSD cache performance modeling](#) on page 476  
[Suspend SSD cache](#) on page 482

## 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 name of the SSD cache for which you want to stop modeling performance. Enclose the SSD cache name in square brackets ([ ]). If the SSD cache name has special characters or numbers, you must enclose the SSD cache name 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\ssdcacheperf.csv" You can use any file name, but you must use the .csv 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):

Writes          SSD Reads          HDD Reads          HDD
Overall         Avg.              % of              Avg.              %
Cache Response Response           Response           %
of Response % of
  
```

Capacity	Time	I/Os	Time	I/Os	Time	I/
Os	Time	I/Os				
186 GB	0 ms		0 ms	0.0 %	0 ms	0.0 % 0
ms	0.0 %					
372 GB	0 ms		0 ms	0.0 %	0 ms	0.0 % 0
ms	0.0 %					
557 GB *	0 ms		0 ms	0.0 %	0 ms	0.0 % 0
ms	0.0 %					
558 GB	0 ms		0 ms	0.0 %	0 ms	0.0 % 0
ms	0.0 %					
744 GB	0 ms		0 ms	0.0 %	0 ms	0.0 % 0
ms	0.0 %					
931 GB	0 ms		0 ms	0.0 %	0 ms	0.0 % 0
ms	0.0 %					
1117 GB	0 ms		0 ms	0.0 %	0 ms	0.0 % 0
ms	0.0 %					
1303 GB	0 ms		0 ms	0.0 %	0 ms	0.0 % 0
ms	0.0 %					
1489 GB	0 ms		0 ms	0.0 %	0 ms	0.0 % 0
ms	0.0 %					
1675 GB	0 ms		0 ms	0.0 %	0 ms	0.0 % 0
ms	0.0 %					
1862 GB	0 ms		0 ms	0.0 %	0 ms	0.0 % 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

## Related references

- [Add drives to SSD cache](#) on page 77
- [Change SSD cache application type](#) on page 90
- [Create SSD cache](#) on page 159
- [Delete SSD cache](#) on page 188
- [Remove drives from SSD cache](#) on page 235
- [Rename SSD cache](#) on page 242
- [Resume SSD cache](#) on page 260
- [Set SSD cache for a volume](#) on page 342
- [Show SSD cache](#) on page 406
- [Show SSD cache statistics](#) on page 407
- [Start SSD cache locate](#) on page 457
- [Start SSD cache performance modeling](#) on page 458
- [Stop SSD cache locate](#) on page 475

[Suspend SSD cache](#) on page 482

## 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.

### Related references

[Start storage array configuration database diagnostic](#) on page 459

[Start storage array configuration database diagnostic](#) on page 459

## 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

**Related references**

[Download storage array drive firmware](#) on page 204

**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

**Related references**

[Activate iSCSI initiator](#) on page 72

[Create iSCSI initiator](#) on page 131

[Delete iSCSI initiator](#) on page 184

[Diagnose controller iSCSI host cable](#) on page 193

[Reset storage array iSCSI baseline](#) on page 252

[Save storage array iSCSI statistics](#) on page 279

[Set iSCSI initiator](#) on page 324

[Set iSCSI target properties](#) on page 326

[Show current iSCSI sessions](#) on page 390

[Show storage array negotiation defaults](#) on page 423

[Show storage array unconfigured iSCSI initiators](#) on page 424

[Start iSCSI DHCP refresh](#) on page 453

[Reset storage array iSCSI baseline](#) on page 252

[Save storage array iSCSI statistics](#) on page 279

[Show storage array unconfigured iSCSI initiators](#) on page 424

**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

**Related references**[Start storage array locate](#) on page 462**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

**Related references**[Start tray locate](#) on page 463**Stop volume copy**

This command stops a volume copy operation. This command is valid for 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 or numbers, you must enclose the target volume name in double quotation marks (" ") inside square brackets.
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 or numbers, you must enclose the source volume name in double quotation marks (" ") inside square brackets.

**Related references**[Create volume copy](#) on page 165[Recopy volume copy](#) on page 224



[Remove volume copy](#) on page 239

[Set volume copy](#) on page 377

[Show volume copy](#) on page 431

[Show volume copy source candidates](#) on page 432

[Show volume copy target candidates](#) on page 432

## 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

### Related references

[Start volume group locate](#) on page 465

## 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
<code>asyncMirrorGroup</code>	The name of the asynchronous mirror group for which you want to suspend synchronization of data. Enclose the asynchronous mirror group name in square brackets ([ ]). If the asynchronous mirror group name has special characters or numbers, you 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

**Related references**[Resume asynchronous mirror group](#) on page 255[Show asynchronous mirror group synchronization progress](#) on page 383[Start asynchronous mirroring synchronization](#) on page 440**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 name of the SSD cache that you want to suspend. Enclose the SSD cache name in square brackets ([ ]). If the SSD cache name has special characters or numbers, you must enclose the SSD cache name in double quotation marks (" ") inside square brackets.

**Notes**

To restart caching, use the `resume ssdCache` command.

**Minimum firmware level**

7.84

**Related references**[Add drives to SSD cache](#) on page 77[Change SSD cache application type](#) on page 90[Create SSD cache](#) on page 159[Delete SSD cache](#) on page 188[Remove drives from SSD cache](#) on page 235[Rename SSD cache](#) on page 242[Resume SSD cache](#) on page 260[Set SSD cache for a volume](#) on page 342[Show SSD cache](#) on page 406[Show SSD cache statistics](#) on page 407[Start SSD cache locate](#) on page 457[Start SSD cache performance modeling](#) on page 458[Stop SSD cache locate](#) on page 475[Stop SSD cache performance modeling](#) on page 476**Suspend synchronous mirroring**

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>	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 or numbers, you must enclose the volume name in double quotation marks (" ") inside square brackets.
<code>primaries</code>	The names of several primary volumes for which you want to suspend operations. All of the volumes will have the same properties. Enter the names of the volumes using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Separate each of the names with a space.</li> </ul> If the volume names have special characters or numbers, enter the names using these rules: <ul style="list-style-type: none"> <li>• Enclose all of the names in square brackets ([ ]).</li> <li>• Enclose each of the names in double quotation marks (" ").</li> <li>• Separate each of the names with a space.</li> </ul>
<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 <code>TRUE</code> . For the volumes that are separate, set this parameter to <code>FALSE</code> .

## 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

**Related references**[Remove synchronous mirroring](#) on page 238[Resume synchronous mirroring](#) on page 261[Set synchronous mirroring](#) on page 360[Show synchronous mirroring volume candidates](#) on page 426[Show synchronous mirroring volume synchronization progress](#) on page 427[Start synchronous mirroring synchronization](#) on page 462**Test alerts**

**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

**Restriction:** This command is obsolete for the E2800 storage array. Use the RESTful API, SANtricity System Manager, or see [Using curl commands with the E2800 storage array](#) on page 39.

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

**Related references**[Using curl commands with the E2800 storage array](#) on page 39[Set event alert filtering](#) on page 316[Show blocked events](#) on page 384[Clear storage array event log](#) on page 101[Save storage array events](#) on page 276**Test asynchronous mirror group connectivity**

This command tests possible communication problems between the local storage array and the remote storage array associated with an asynchronous mirror group.

**Syntax**

```
diagnose asyncMirrorGroup [asyncMirrorGroupName]
testID=(all | connectivity | latency | bandwidth | portConnections)
```

**Parameter**

Parameter	Description
<code>asyncMirrorGroup</code>	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 or numbers, you must enclose the asynchronous mirror group name in double quotation marks ( " " ) inside square brackets.
<code>testID</code>	The identifier for the diagnostic test you want to run. The identifier and corresponding tests are as follow: <ul style="list-style-type: none"> <li>• All— Performs all the tests associated with this command.</li> <li>• 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.</li> <li>• 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.</li> <li>• 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.</li> <li>• 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.</li> </ul>

**Minimum firmware level**

7.84

**Related references**

[Check storage array connectivity](#) on page 94

**Test the AutoSupport configuration**

The `SMcli -autoSupportConfig test` command operates on the management domain. The storage management software chooses an Auto-Support capable array in the management domain for sample data to include with the test message.

**Description**

Use this command to test the AutoSupport configuration by sending a test message.

## Context

**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.

**Restriction:** This command is obsolete for the E2800 storage array. Use the RESTful API, SANtricity System Manager, or see [Using curl commands with the E2800 storage array](#) on page 39.

**Note:** The Persistent Monitor should be running on the management station to send scheduled or event-triggered AutoSupport messages. Start the service before using this command.

## Syntax

```
SMcli -autoSupportConfig test reply-toEmail<address>
```

## Parameters

The following table lists the parameters for sending an AutoSupport test message.

Parameter	Description
test	Required parameter for the command. Requires no value.
reply-toEmail	Allows the user to specify the reply-to email address for the AutoSupport messages of SMTP delivery type

## Minimum Firmware Level

8.25

## Example of testing using HTTP/HTTPs delivery

Following are examples of the usage of this command when using an HTTP or HTTPs delivery type, where you will not specify an email-address for delivery.

```
SMcli -autoSupportConfig test
The sample AutoSupport message was successfully sent to the ASUP gateway.

SMcli completed successfully.
```

## Example of testing using SMTP delivery

Following are examples of the usage of this command. When using an SMTP delivery method, you can use the reply-toEmail parameter to specify the destination address.

```
SMcli -autoSupportConfig test reply-toEmail:user@netapp.com
The sample AutoSupport message was successfully sent to the ASUP gateway.

An email has been sent to 'user@netapp.com', please check the mailbox
for delivery confirmation.

SMcli completed successfully.
```

## Related references

[Using curl commands with the E2800 storage array](#) on page 39

[Capture or view an AutoSupport log](#) on page 89

[Display AutoSupport bundle collection settings](#) on page 198  
[Display AutoSupport message collection schedule](#) on page 200  
[Disable AutoSupport at the EMW level SMcli version](#) on page 197  
[Enable or disable AutoSupport at the EMW management domain level \(SMcli only\)](#) on page 210  
[Enable or disable the AutoSupport OnDemand feature at the EMW management domain level \(SMcli only\)](#) on page 211  
[Enable or disable the AutoSupport OnDemand Remote Diagnostics feature at the EMW management domain level \(SMcli only\)](#) on page 213  
[Reset AutoSupport message collection schedule](#) on page 247  
[Save storage array support data](#) on page 283  
[Specify the AutoSupport delivery method](#) on page 438

## Test SNMP trap destination

This command tests the connection and identification of the trap receiver for a specific Simple Network Management Protocol (SNMP) community. This command tests the trap destination by sending a trap message to the trap receiver.

### Syntax

```
start snmpTrapDestination trapReceiverIP=ipAddress
communityName="communityName"
```

### Parameters

Parameter	Description
trapReceiverIP	The IP address of the SNMP manager to which you want to send trap messages.
communityName	The name of the SNMP community for which you want to send trap messages.

### Minimum firmware level

8.30

## Unregister SNMP community

This command deletes an existing Simple Network Management Protocol (SNMP) community that you previously created and registered using the `create snmpCommunity` command. Deleting an SNMP community effectively unregisters the community. Any trap destinations associated with that community are also removed.

### Syntax

```
delete snmpCommunity communityName="snmpCommunityName"
```

### Parameters

Parameter	Description
communityName	The name of the SNMP community that you want to delete. Enclose the SNMP community name in double quotation marks (" ").

**Minimum firmware level**

8.30

**Unregister SNMP trap destination**

This command deletes an existing Simple Network Management Protocol (SNMP) trap destination that you previously created and registered using the `create snmpTrapDestination` command.

**Syntax**

```
delete snmpTrapDestination trapReceiverIP=ipAddress
```

**Parameters**

Parameter	Description
trapReceiverIP	The IP address of the SNMP manager from which you want to stop sending trap messages.

**Minimum firmware level**

8.30

**Update SNMP community**

This command creates a new name for an existing Simple Network Management Protocol (SNMP) community.

**Syntax**

```
set snmpCommunity communityName="snmpCommunityName"  
newCommunityName="newSnmpCommunityName"
```

**Parameters**

Parameter	Description
communityName	The name of the existing SNMP community that you want to rename. Enclose the SNMP community name in double quotation marks (" ").
newCommunityName	The new name that you want to give to the SNMP community. Enclose the SNMP community name in double quotation marks (" ").

**Minimum firmware level**

8.30

**Update SNMP MIB II system group variables**

This command changes the system variables for the Simple Network Management Protocol (SNMP). The following system variables are maintained in a Management Information Base II (MIB-II) database:

- Name of the system
- Name of the system contact
- Location of the system



## Syntax

```
set snmpSystemVariables
[sysName=newSystemName]
[sysContact=contactName]
[sysLocation=systemLocation]
```

## Parameters

Parameter	Description
sysName	The new name that you want to give to the SNMP system. Use standard SNMP and MIB conventions for the system name. Enclose the SNMP system name in square brackets([ ]).
sysContact	The name of contact person for the managed system with information on how to contact this person. Enclose the SNMP contact name in square brackets([ ]).
sysLocation	The physical location of the system, such as "3rd flr". Enclose the SNMP system location in square brackets([ ]).

## Minimum firmware level

8.30

## Update SNMP trap destination

This command turns on or turns off sending authentication failure messages for an SNMP trap destination. Failure occurs when the SNMP agent received a message from an SNMP manager, but the message contained an invalid community name.

## Syntax

```
set snmpTrapDestination trapReceiverIP=ipAddress
[communityName="communityName"]
[sendAuthenticationFailureTraps=(TRUE | FALSE)]
```

## Parameters

Parameter	Description
trapReceiverIP	The IP address of the SNMP manager to which you want to send trap messages.
communityName	The name of the SNMP community for which you want to send trap messages.
sendAuthenticationFailureTraps	This parameter turns on or turns off sending authentication failure messages to an SNMP manager. To send authentication failure messages, set the parameter to <code>TRUE</code> . To prevent sending authentication failure messages, set the parameter to <code>FALSE</code> . The default is <code>TRUE</code> .

## Minimum firmware level

8.30

## Validate storage array security key

The `validate storageArray securityKey` command validates the security key for a storage array that has Full Disk Encryption (FDE) drives to make sure that the back-up security key file is not corrupt. Security key validation uses the pass phrase used to generate the file to decrypt the key and match it with the value maintained in controller memory.

### 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" <b>IMPORTANT</b> – 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

### Related references

[Create storage array security key](#) on page 161

[Export storage array security key](#) on page 218

[Import storage array security key](#) on page 218

[Set storage array security key](#) on page 358

## Changes in the CLI and script commands in this release

This topic provides the following tables that list the changes in the SANtricity Storage Manager 11.30 (firmware 8.30) release of the CLI commands and Script commands.

The following table lists changes to the CLI and Script Commands with this release of SANtricity Storage Manager.

Command	Description of Change
<i>Set storage array to enable or disable automatic load balancing</i> on page 349	Added topic
<i>Set controller service action allowed indicator</i> on page 306	This command is only valid for E2700, E5600, or EF560 (and other previous generations of E-Series or EF-Series controllers). The E2800 controller tray has a single Attention indicator that will only illuminate when 1) there is a failure, and 2) the component that failed can be safely removed.
<i>Start iSCSI DHCP refresh</i> on page 453	Added information about the new way to identify and specify an iSCSI host port in CLI command.
<i>Set controller</i> on page 299	Added information about the new way to identify and specify an iSCSI host port in CLI command.
<i>Diagnose controller iSCSI host cable</i> on page 193	Added information about the new way to identify and specify an iSCSI host port in CLI command.
<i>Activate synchronous mirroring</i> on page 74	The reference, "Slot ID values are 1 to 32" changed to "Slot ID values are 1 to 24". Documentation improvement or correction only
<i>Create volume group</i> on page 167	The reference, "Slot ID values are 1 to 32" changed to "Slot ID values are 1 to 24". Documentation improvement or correction only
<i>Create disk pool</i> on page 121	The reference, "Slot ID values are 1 to 32" changed to "Slot ID values are 1 to 24". Documentation improvement or correction only
<i>Download drive firmware</i> on page 202	The reference, "Slot ID values are 1 to 32" changed to "Slot ID values are 1 to 24". Documentation improvement or correction only
<i>Set foreign drive to native</i> on page 319	The reference, "Slot ID values are 1 to 32" changed to "Slot ID values are 1 to 24". Documentation improvement or correction only
<i>Set volume group</i> on page 378	The reference, "Slot ID values are 1 to 32" changed to "Slot ID values are 1 to 24". Documentation improvement or correction only
<i>Set drive state</i> on page 315	The reference, "Slot ID values are 1 to 32" changed to "Slot ID values are 1 to 24". Documentation improvement or correction only

Command	Description of Change
<i>Set disk pool (modify disk pool)</i> on page 310	The reference, "Slot ID values are 1 to 32" changed to "Slot ID values are 1 to 24". Documentation improvement or correction only
<i>Reset drive</i> on page 249	The reference, "Slot ID values are 1 to 32" changed to "Slot ID values are 1 to 24". Documentation improvement or correction only
<i>Replace drive</i> on page 244	The reference, "Slot ID values are 1 to 32" changed to "Slot ID values are 1 to 24". Documentation improvement or correction only
<i>Reduce disk pool capacity</i> on page 233	The reference, "Slot ID values are 1 to 32" changed to "Slot ID values are 1 to 24". Documentation improvement or correction only
<i>Set storage array learn cycle</i> on page 356	Added the controller parameter. The learn cycles are not linked together and so a disruption in one controller battery learn cycle will not impact the learn cycle for the battery in the other controller.
<i>Initialize thin volume</i> on page 221	Added notes that the maximum physical capacity is 256 TB (change from 64 TB). Added note that Immediate Availability Format (IAF) does not apply to volumes larger than 64 TB.
<i>Register SNMP community</i> on page 234	Added topic.
<i>Register SNMP trap destination</i> on page 234	Added topic.
<i>Show SNMP communities</i> on page 404	Added topic.
<i>Show SNMP MIB II system group variables</i> on page 405	Added topic.
<i>Test SNMP trap destination</i> on page 487	Added topic.
<i>Unregister SNMP community</i> on page 487	Added topic.
<i>Unregister SNMP trap destination</i> on page 488	Added topic.
<i>Update SNMP community</i> on page 488	Added topic.
<i>Update SNMP MIB II system group variables</i> on page 488	Added topic.
<i>Update SNMP trap destination</i> on page 489	Added topic.
<i>Save storage array support data</i> on page 283	Updated the list of artifacts in the support bundle to describe the new files for the E2800, and the obsoleted files for the E2800.

Command	Description of Change
<a href="#">Download Environmental Card firmware</a> on page 204	<p>For the allTrays parameter, specified that it this parameter more specifically means "all compatible" trays, and that the incompatible tray list is determined by the firmware package selected for download.</p> <p>Old Version: This parameter downloads new firmware to all of the trays in the storage array.</p> <p>New Version: This parameter downloads new firmware to all of the compatible trays in the storage array. The firmware package you select determines which trays are compatible. Incompatible trays are skipped. No error message is displayed about incompatible trays.</p>
<a href="#">Capture or view an AutoSupport log</a> on page 89	Added a note that this command is obsolete for the E2800, and that SANtricity System Manager or the RESTful API should be used instead.
<a href="#">Display AutoSupport bundle collection settings</a> on page 198	Added a note that this command is obsolete for the E2800, and that SANtricity System Manager or the RESTful API should be used instead.
<a href="#">Display AutoSupport message collection schedule</a> on page 200	Added a note that this command is obsolete for the E2800, and that SANtricity System Manager or the RESTful API should be used instead.
<a href="#">Disable AutoSupport at the EMW level (SMcli only)</a> on page 197	Added a note that this command is obsolete for the E2800, and that SANtricity System Manager or the RESTful API should be used instead.
<a href="#">Enable AutoSupport at the EMW management domain level (SMcli only)</a> on page 210	Added a note that this command is obsolete for the E2800, and that SANtricity System Manager or the RESTful API should be used instead.
<a href="#">Enable or disable the AutoSupport OnDemand feature at the EMW management domain level (SMcli only)</a> on page 211	Added a note that this command is obsolete for the E2800, and that SANtricity System Manager or the RESTful API should be used instead.
<a href="#">Enable or disable the AutoSupport OnDemand remote diagnostics feature at the EMW management domain level (SMcli only)</a> on page 213	Added a note that this command is obsolete for the E2800, and that SANtricity System Manager or the RESTful API should be used instead.
<a href="#">Reset AutoSupport message collection schedule</a> on page 247	Added a note that this command is obsolete for the E2800, and that SANtricity System Manager or the RESTful API should be used instead.
<a href="#">Test alerts</a> on page 484	Added a note that this command is obsolete for the E2800, and that SANtricity System Manager or the RESTful API should be used instead.
<a href="#">Specify AutoSupport delivery method</a> on page 438	Added a note that this command is obsolete for the E2800, and that SANtricity System Manager or the RESTful API should be used instead.
<a href="#">Test AutoSupport configuration</a> on page 485	Added a note that this command is obsolete for the E2800, and that SANtricity System Manager or the RESTful API should be used instead.

Command	Description of Change
<i>Storage management</i> on page 55	Added topic.
<i>Managing an E2800 storage array</i> on page 56	Added topic.
<i>Managing multiple E2800 storage arrays</i> on page 56	Added topic.
<i>Management of mixed storage arrays</i> on page 57	Added topic.
<i>Event Management</i> on page 58	Added topic.
<i>Event management for E2800 storage arrays</i> on page 60	Added topic.
<i>Event management for the E2700, E5600, and EF560</i> on page 61	Added topic.
<i>AutoSupport Overview</i> on page 62	Added topic.
<i>AutoSupport and AutoSupportOnDemand</i> on page 62	Added topic.
<i>AutoSupport for E2800 storage arrays</i> on page 64	Added topic.
<i>AutoSupport for E2700, E5600, and EF560 storage arrays</i> on page 65	Added topic.
<i>Management Methods in SANtricity 11.30</i> on page 66	Added topic.
<i>Terminology differences in SANtricity 11.30</i> on page 70	Added topic.
<i>Show replaceable drives</i> on page 399	Added topic. This is not new for Lehigh, but had been missed in the 7.10 release. Documentation improvement or correction only
<i>Validate storage array security key</i> on page 490	Added clarification to wording. Documentation improvement or correction only
<i>Save auto-load balancing statistics</i> on page 270	Added topic. Documentation improvement or correction only
<i>Using curl commands with the E2800</i> on page 39	Added topic.
<i>Create asynchronous mirror group</i> on page 104	Rearranged order of parameters to show remotePassword before interfaceType, changed case of the remoteStorageArrayNameWwn parameter to remoteStorageArrayNameWWN, and noted that time-based parameters must have no space inbetween the integer and the interval specified. Documentation improvement or correction only

Command	Description of Change
<i>Set host port label</i> on page 323	Removed the host optional parameter. This command only resets the user label for the host port, not the associated host. Documentation improvement or correction only
Reset storage array battery install date	Obsoleted the command for all SANtricity 11.30 controllers (E2800, E2700, E5600, EF560)
Stop cache backup device diagnostic	Obsoleted the command for all SANtricity 11.30 controllers (E2800, E2700, E5600, EF560)
Start cache backup device diagnostic	Obsoleted the command for all SANtricity 11.30 controllers (E2800, E2700, E5600, EF560)
Stop cache memory diagnostic	Obsoleted the command for all SANtricity 11.30 controllers (E2800, E2700, E5600, EF560)
Start cache memory diagnostic	Obsoleted the command for all SANtricity 11.30 controllers (E2800, E2700, E5600, EF560)
Stop controller diagnostic	Obsoleted the command for all SANtricity 11.30 controllers (E2800, E2700, E5600, EF560)
Start controller diagnostic	Obsoleted the command for all SANtricity 11.30 controllers (E2800, E2700, E5600, EF560)
Stop host interface card diagnostic	Obsoleted the command for all SANtricity 11.30 controllers (E2800, E2700, E5600, EF560)
Start host interface card diagnostic	Obsoleted the command for all SANtricity 11.30 controllers (E2800, E2700, E5600, EF560)
Show cache backup device diagnostic status	Obsoleted the command for all SANtricity 11.30 controllers (E2800, E2700, E5600, EF560)
Show cache memory diagnostic status	Obsoleted the command for all SANtricity 11.30 controllers (E2800, E2700, E5600, EF560)
Show controller diagnostic status	Obsoleted the command for all SANtricity 11.30 controllers (E2800, E2700, E5600, EF560)

## Copyright information

---

Copyright © 1994–2016 NetApp, Inc. All rights reserved. Printed in the U.S.

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.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

**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).



## Trademark information

---

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL and other names are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web.

<http://www.netapp.com/us/legal/netapptmlist.aspx>

## How to send comments about documentation and receive update notifications

---

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

[\*doccomments@netapp.com\*](mailto:doccomments@netapp.com)

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

## Changes in the CLI and script commands in this release

This topic provides the following tables that list the changes in the SANtricity Storage Manager 11.30 (firmware 8.30) release of the CLI commands and Script commands.

The following table lists changes to the CLI and Script Commands with this release of SANtricity Storage Manager.

Command	Description of Change
<i>Set storage array to enable or disable automatic load balancing</i> on page 349	Added topic
<i>Set controller service action allowed indicator</i> on page 306	This command is only valid for E2700, E5600, or EF560 (and other previous generations of E-Series or EF-Series controllers). The E2800 controller tray has a single Attention indicator that will only illuminate when 1) there is a failure, and 2) the component that failed can be safely removed.
<i>Start iSCSI DHCP refresh</i> on page 453	Added information about the new way to identify and specify an iSCSI host port in CLI command.
<i>Set controller</i> on page 299	Added information about the new way to identify and specify an iSCSI host port in CLI command.
<i>Diagnose controller iSCSI host cable</i> on page 193	Added information about the new way to identify and specify an iSCSI host port in CLI command.
<i>Activate synchronous mirroring</i> on page 74	The reference, "Slot ID values are 1 to 32" changed to "Slot ID values are 1 to 24". Documentation improvement or correction only
<i>Create volume group</i> on page 167	The reference, "Slot ID values are 1 to 32" changed to "Slot ID values are 1 to 24". Documentation improvement or correction only
<i>Create disk pool</i> on page 121	The reference, "Slot ID values are 1 to 32" changed to "Slot ID values are 1 to 24". Documentation improvement or correction only
<i>Download drive firmware</i> on page 202	The reference, "Slot ID values are 1 to 32" changed to "Slot ID values are 1 to 24". Documentation improvement or correction only
<i>Set foreign drive to native</i> on page 319	The reference, "Slot ID values are 1 to 32" changed to "Slot ID values are 1 to 24". Documentation improvement or correction only
<i>Set volume group</i> on page 378	The reference, "Slot ID values are 1 to 32" changed to "Slot ID values are 1 to 24". Documentation improvement or correction only
<i>Set drive state</i> on page 315	The reference, "Slot ID values are 1 to 32" changed to "Slot ID values are 1 to 24". Documentation improvement or correction only

Command	Description of Change
<i>Set disk pool (modify disk pool)</i> on page 310	The reference, "Slot ID values are 1 to 32" changed to "Slot ID values are 1 to 24". Documentation improvement or correction only
<i>Reset drive</i> on page 249	The reference, "Slot ID values are 1 to 32" changed to "Slot ID values are 1 to 24". Documentation improvement or correction only
<i>Replace drive</i> on page 244	The reference, "Slot ID values are 1 to 32" changed to "Slot ID values are 1 to 24". Documentation improvement or correction only
<i>Reduce disk pool capacity</i> on page 233	The reference, "Slot ID values are 1 to 32" changed to "Slot ID values are 1 to 24". Documentation improvement or correction only
<i>Set storage array learn cycle</i> on page 356	Added the controller parameter. The learn cycles are not linked together and so a disruption in one controller battery learn cycle will not impact the learn cycle for the battery in the other controller.
<i>Initialize thin volume</i> on page 221	Added notes that the maximum physical capacity is 256 TB (change from 64 TB). Added note that Immediate Availability Format (IAF) does not apply to volumes larger than 64 TB.
<i>Register SNMP community</i> on page 234	Added topic.
<i>Register SNMP trap destination</i> on page 234	Added topic.
<i>Show SNMP communities</i> on page 404	Added topic.
<i>Show SNMP MIB II system group variables</i> on page 405	Added topic.
<i>Test SNMP trap destination</i> on page 487	Added topic.
<i>Unregister SNMP community</i> on page 487	Added topic.
<i>Unregister SNMP trap destination</i> on page 488	Added topic.
<i>Update SNMP community</i> on page 488	Added topic.
<i>Update SNMP MIB II system group variables</i> on page 488	Added topic.
<i>Update SNMP trap destination</i> on page 489	Added topic.
<i>Save storage array support data</i> on page 283	Updated the list of artifacts in the support bundle to describe the new files for the E2800, and the obsoleted files for the E2800.

Command	Description of Change
<a href="#">Download Environmental Card firmware</a> on page 204	<p>For the allTrays parameter, specified that it this parameter more specifically means "all compatible" trays, and that the incompatible tray list is determined by the firmware package selected for download.</p> <p>Old Version: This parameter downloads new firmware to all of the trays in the storage array.</p> <p>New Version: This parameter downloads new firmware to all of the compatible trays in the storage array. The firmware package you select determines which trays are compatible. Incompatible trays are skipped. No error message is displayed about incompatible trays.</p>
<a href="#">Capture or view an AutoSupport log</a> on page 89	Added a note that this command is obsolete for the E2800, and that SANtricity System Manager or the RESTful API should be used instead.
<a href="#">Display AutoSupport bundle collection settings</a> on page 198	Added a note that this command is obsolete for the E2800, and that SANtricity System Manager or the RESTful API should be used instead.
<a href="#">Display AutoSupport message collection schedule</a> on page 200	Added a note that this command is obsolete for the E2800, and that SANtricity System Manager or the RESTful API should be used instead.
<a href="#">Disable AutoSupport at the EMW level (SMcli only)</a> on page 197	Added a note that this command is obsolete for the E2800, and that SANtricity System Manager or the RESTful API should be used instead.
<a href="#">Enable AutoSupport at the EMW management domain level (SMcli only)</a> on page 210	Added a note that this command is obsolete for the E2800, and that SANtricity System Manager or the RESTful API should be used instead.
<a href="#">Enable or disable the AutoSupport OnDemand feature at the EMW management domain level (SMcli only)</a> on page 211	Added a note that this command is obsolete for the E2800, and that SANtricity System Manager or the RESTful API should be used instead.
<a href="#">Enable or disable the AutoSupport OnDemand remote diagnostics feature at the EMW management domain level (SMcli only)</a> on page 213	Added a note that this command is obsolete for the E2800, and that SANtricity System Manager or the RESTful API should be used instead.
<a href="#">Reset AutoSupport message collection schedule</a> on page 247	Added a note that this command is obsolete for the E2800, and that SANtricity System Manager or the RESTful API should be used instead.
<a href="#">Test alerts</a> on page 484	Added a note that this command is obsolete for the E2800, and that SANtricity System Manager or the RESTful API should be used instead.
<a href="#">Specify AutoSupport delivery method</a> on page 438	Added a note that this command is obsolete for the E2800, and that SANtricity System Manager or the RESTful API should be used instead.
<a href="#">Test AutoSupport configuration</a> on page 485	Added a note that this command is obsolete for the E2800, and that SANtricity System Manager or the RESTful API should be used instead.

Command	Description of Change
<i>Storage management</i> on page 55	Added topic.
<i>Managing an E2800 storage array</i> on page 56	Added topic.
<i>Managing multiple E2800 storage arrays</i> on page 56	Added topic.
<i>Management of mixed storage arrays</i> on page 57	Added topic.
<i>Event Management</i> on page 58	Added topic.
<i>Event management for E2800 storage arrays</i> on page 60	Added topic.
<i>Event management for the E2700, E5600, and EF560</i> on page 61	Added topic.
<i>AutoSupport Overview</i> on page 62	Added topic.
<i>AutoSupport and AutoSupportOnDemand</i> on page 62	Added topic.
<i>AutoSupport for E2800 storage arrays</i> on page 64	Added topic.
<i>AutoSupport for E2700, E5600, and EF560 storage arrays</i> on page 65	Added topic.
<i>Management Methods in SANtricity 11.30</i> on page 66	Added topic.
<i>Terminology differences in SANtricity 11.30</i> on page 70	Added topic.
<i>Show replaceable drives</i> on page 399	Added topic. This is not new for Lehigh, but had been missed in the 7.10 release. Documentation improvement or correction only
<i>Validate storage array security key</i> on page 490	Added clarification to wording. Documentation improvement or correction only
<i>Save auto-load balancing statistics</i> on page 270	Added topic. Documentation improvement or correction only
<i>Using curl commands with the E2800</i> on page 39	Added topic.
<i>Create asynchronous mirror group</i> on page 104	Rearranged order of parameters to show remotePassword before interfaceType, changed case of the remoteStorageArrayNameWwn parameter to remoteStorageArrayNameWWN, and noted that time-based parameters must have no space inbetween the integer and the interval specified. Documentation improvement or correction only

Command	Description of Change
<i>Set host port label</i> on page 323	Removed the host optional parameter. This command only resets the user label for the host port, not the associated host. Documentation improvement or correction only
Reset storage array battery install date	Obsoleted the command for all SANtricity 11.30 controllers (E2800, E2700, E5600, EF560)
Stop cache backup device diagnostic	Obsoleted the command for all SANtricity 11.30 controllers (E2800, E2700, E5600, EF560)
Start cache backup device diagnostic	Obsoleted the command for all SANtricity 11.30 controllers (E2800, E2700, E5600, EF560)
Stop cache memory diagnostic	Obsoleted the command for all SANtricity 11.30 controllers (E2800, E2700, E5600, EF560)
Start cache memory diagnostic	Obsoleted the command for all SANtricity 11.30 controllers (E2800, E2700, E5600, EF560)
Stop controller diagnostic	Obsoleted the command for all SANtricity 11.30 controllers (E2800, E2700, E5600, EF560)
Start controller diagnostic	Obsoleted the command for all SANtricity 11.30 controllers (E2800, E2700, E5600, EF560)
Stop host interface card diagnostic	Obsoleted the command for all SANtricity 11.30 controllers (E2800, E2700, E5600, EF560)
Start host interface card diagnostic	Obsoleted the command for all SANtricity 11.30 controllers (E2800, E2700, E5600, EF560)
Show cache backup device diagnostic status	Obsoleted the command for all SANtricity 11.30 controllers (E2800, E2700, E5600, EF560)
Show cache memory diagnostic status	Obsoleted the command for all SANtricity 11.30 controllers (E2800, E2700, E5600, EF560)
Show controller diagnostic status	Obsoleted the command for all SANtricity 11.30 controllers (E2800, E2700, E5600, EF560)