



ATTO Technology, Inc.

---

**ATTO ConfigTool  
for  
ThunderStream SC 3808**

## **ATTO Technology, Inc.**

155 CrossPoint Parkway  
Amherst, New York 14068 USA

[www.attotech.com](http://www.attotech.com)

Tel (716) 691-1999

Fax (716) 691-9353

Sales support: [sls@attotech.com](mailto:sls@attotech.com)

Technical support: Monday -- Friday, 8am-6pm EST

[techsupp@attotech.com](mailto:techsupp@attotech.com) (716) 691-1999 x242

© 2012 ATTO Technology, Inc. All rights reserved. All brand or product names are trademarks of their respective holders. No part of this manual may be reproduced in any form or by any means without the express written permission of ATTO Technology, Inc.

# Contents

---

<b>1 ATTO ConfigTool .....</b>	<b>1</b>
Pre-Installation	
Installation	
Using the ATTO ConfigTool	
Navigating the ConfigTool	
Driver update	
System Reboot	
<b>2 ThunderStream SC NVRAM Settings .....</b>	<b>9</b>
<b>3 RAID Configuration .....</b>	<b>10</b>
Preliminary Steps	
Setting up DVRAID	
Customizing a RAID setup	
Creating a Hot Spare Pool	
Modifying RAID Groups	
Replacing a faulted drive	
Setup RAID Notification	
Using SCSI Enclosure Services (SES)	
Identifying Drives	
Selecting All Drives in the Same Enclosure	
Monitoring Health	
Monitoring S.M.A.R.T. Data	
Enabling or Disabling S.M.A.R.T. Monitoring	
Checking S.M.A.R.T. Status	
Filtering S.M.A.R.T. Attributes	
S.M.A.R.T. Notifications	
Test Drive Performance	
RAID Media Scan Feature	
Starting a Media Scan	
Media Scan Status	
Viewing the Scan Report	
Pausing or Resuming Media Scan	
Scheduling Media Scan	
Viewing Scheduled Tasks	
Automatic Cancellation of Media Scan	
Data Recovery for OFFLINE RAID Groups	
Definitions	
Failure Scenario Table	
Drive Replacement on a Failure Condition	
Recovering from Failed Drives	
File System Repair Tools	
Command Line Interface	
<b>4 SNMP Configuration .....</b>	<b>49</b>
Definitions	
Details	
Configuration Options	
Control Buttons	
Troubleshooting	
<b>5 Troubleshoot the ATTO ConfigTool .....</b>	<b>56</b>

**Appendix A CLI provides an ASCII-based interface .....i**

- CLI error messages
- CLI summary
- CLI command explanations

**Appendix B Glossary .....xvi**

# 1 ATTO ConfigTool

---

The ATTO ConfigTool is a utility program that displays information about installed storage controllers, drivers and devices and provides a mechanism to configure installed storage controllers in local and remote hosts.

The ATTO ConfigTool provides a central interface for managing local and remote hosts with ATTO storage controllers installed. Features include:

- The names of ATTO storage controllers installed in the host
- Information about the devices attached to ATTO storage controllers
- Information about the drivers and firmware controlling the storage controllers, including version information

You may also use the ATTO ConfigTool to:

- Update the flash image when a new version is released by ATTO
- Modify the NVRAM settings (refer to the product-specific NVRAM information sections in this manual)
- Manage RAID groups
- Configure RAID notifications
- Revert to default factory settings

The factory settings should provide excellent performance for a wide range of applications. However, some applications may benefit from modification of the storage controller NVRAM settings which tune the storage controller for a specific performance range.

## Pre-Installation

---

The ATTO ConfigTool is a free utility available on CD (if supplied with your storage controller) or via the ATTO website at [www.attotech.com](http://www.attotech.com). This application supports a variety of operating systems. Please visit the ATTO website for specific operating system support information and downloads.

Oracle® Java version 1.5, or later, is required to use the ATTO ConfigTool. Visit <http://java.com> for the latest Java updates for Linux and Windows. The latest Java runtime for Mac OS X can be obtained through Software Update.



### Note

*Note: The GNU version of the JVM does not work with the ATTO installer. You must use Oracle Java.*

To install the ATTO ConfigTool:

- OS X - mount the .dmg file

# Installation

---

The ATTO ConfigTool includes two components: a GUI Application and a system service. Either, or both, of these components can be installed on a host, depending on the functionality desired. This creates several different use cases, described below:

A host with both the GUI Application and the system service installed will be able to discover and manage local and remote hosts. This is the most comprehensive management capability. During installation of the ATTO ConfigTool select **FULL** installation for this coverage level.

- A host with only the GUI Application installed will be able to discover and manage remote hosts. This host will not be able to discover and manage any ATTO storage controllers installed on the same host. During installation of the ATTO ConfigTool select **APPLICATION ONLY** for this coverage level.
- A host with only the system service installed can be discovered and managed remotely, but it cannot be managed locally (the GUI Application is required for local management). It is common to have multiple hosts with only the system service installed being managed by a single host remotely. To set up a host without any local management capabilities, refer to the OS specific instructions below.

## Instructions for FULL or APPLICATION ONLY installation:

- 1 Launch the ATTO ConfigTool Installer
  - OS X - Run the ConfigTool\_4xx file
- 2 Select FULL or APPLICATION ONLY as desired
- 3 Follow the on-screen instructions

## Instructions for SYSTEM SERVICE ONLY installation:

- 1 Launch the ATTO ConfigTool system service Installer
  - OS X - Run the attocfgd.mpkg file in the Service folder
- 2 Follow the on-screen instructions

# Using the ATTO ConfigTool


To use the ATTO ConfigTool, locate and double-click the application icon. The main page has three windows: **Device Listing**, **Configuration Options** and **Status**. See Exhibit 1.0-1. The **Device Listing** window at the left of the display page lists local and remote hosts found on the network, as well as currently connected devices. You are required to login to manage any host. Once you login, the **device tree** will expand to reveal additional details on connected storage controllers.

The **Configuration Options** window in the right window pane provides information and options for a device highlighted in the device listing. If you highlight a device in the **Device Listing**, tabs and panels display for that device. The following chart lists the tabs displayed for each device type in the device listing tree.

Tree node	Tabs displayed
Network	Host Tab
Host	Basic Info, Notifications, SNMP
Storage Controller	Basic Info, Flash, RAID, RAID CLI, Tasks
Channel	Basic Info, NVRAM
Devices	Basic Info, Flash, SES

The **Status** window provides general information about host and storage controller settings.

Exhibit 1.0-1 The ATTO ConfigTool opening page.



# Navigating the ConfigTool



## CAUTION

Save system data prior to installing or changing hardware configurations.

## Host Discovery

The ATTO ConfigTool is able to discover hosts automatically, as well as manually. A host will be discovered and displayed in the **Device Listing** only if the system service is installed and running on the host. A host that only has the GUI Application installed will not be displayed in the **Device Listing**.

### Automatic Discovery


Hosts on the local subnet are automatically discovered and placed into the **Device Listing**.

### Manual Discovery

A host can be found manually by using the Network menu and selecting "Find host". The IP address or hostname of the remote host is required. If the host is already in the **Device Listing**, it is not added again.

## Host Login Status

The GUI Application must login into each host before that host can be managed or configured.



The GUI Application indicates the login status for a host as follows:

- A host that is not logged in will display the host icon and the text as gray and italicized.
- A host that is logged in will display the host icon with a green background and black non-italicized text.
- A host that was logged in but the host is rebooting will display the host icon and text as gray and italicized with a strike through.

## Login to a Host

A login is started when the host's Device Listing tree is expanded. A login dialogue box displays and you must enter a username and password. The username and password of an administrator for the host is required to proceed, and is checked against the user credentials stored on the host.

- In Mac OS X, the login information is checked against any user in the Administrators group.

If three successive login attempts fail, the GUI Application will be locked out from the host for three minutes. All login attempts that occur during the lockout period will receive an error status. GUI Applications in other hosts will not be affected by the lockout period.

Once a login is established, the icon of the host is highlighted green and the text becomes black with normal font.



### Note

*There is no logout facility.*

## Select a Host

Select any host by clicking on it in the **Device Listing** window. The following tabs display in the **Configuration Options** window:

- The **Basic Info** tab displays information about the booted operating system on that host. It also shows scheduled reboot information and provides the ability to install a driver on that host. See [Exhibit 1.0-2 on page 6](#).
- The **Notification** tab allows you to set up notification of certain events in the ThunderStream storage controller. Refer to [Setup RAID Notification on page 17](#) or see [Exhibit 3.0-6 on page 19](#).

## Select a Storage Controller

The following tabs display in the **Configuration Options** window when you select a specific storage controller in the **Device Listing** window.

- The **Basic Info** tab provides basic information about the device currently highlighted in the device listing. See [Exhibit 1.0-2 on page 6](#).
- The **Flash** tab provides information about the current revision of flash loaded on the highlighted storage controller. See [Exhibit 1.0-4 on page 7](#). Click on the Browse button at the bottom of the tab to search for new flash files on your host. Click on the Update button to initiate the firmware flashing process using the previously selected firmware.
- The **RAID** tab displays information about the drive inventory, existing RAID groups and Hot Spare devices. From the RAID tab you can create, modify and delete RAID groups. See [Exhibit 3.0-0 on page 10](#).
- The **RAID CLI** tab allows experienced users to enter RAID Command Line Interface commands to the SAS RAID storage controller.
- The **Tasks** tab displays information about tasks that are scheduled to run. In this tab, users can reschedule a task or remove the scheduled task.

## Select a Channel

The following tabs display in the **Configuration Options** window when you select a specific channel in the **Device Listing** window:

- The **Basic Info** tab displays PCI information for the selected channel.
- The **NVRAM** tab displays the NVRAM parameters of the selected channel. For:
  - ThunderStream SC - see [Chapter 2 page 9](#)See, also, [Exhibit 1.0-5 on page 8](#).



### Note

*ATTO storage controllers are designed to operate properly using factory settings. Entering invalid or incorrect settings when using an NVRAM configuration utility such as the ATTO ConfigTool may cause your storage controller to function incorrectly.*

## Select a Device

The following tabs display in the **Configuration Options** window when you select a specific device in the **Device Listing** window:

- The **Basic Info** tab displays information about the selected device.
- The **Flash** tab provides a tool to update the firmware of the selected device. The flash tool is only available for devices that support flash update.
- The **SES** tab displays SES (SCSI Enclosure Services) status information, such as power supplies and fans, for SES devices.

## About panel

The **About** panel, selected from the **About** menu item in the **Help** menu, is an informational page which displays a list of components installed for the ATTO ConfigTool, the tool's version number and ATTO

## Driver update

A storage controller driver can be installed or upgraded on a host by going to the **Basic Info** tab for that host, choosing a driver package in the driver update section and clicking update. The driver package is the .dmg (Mac OS X) file available on the ATTO product CD or downloaded from the ATTO website. If a reboot is required, the GUI will prompt once the install completes. If a driver was installed that

has no matching storage controller, the GUI will prompt to shutdown and install the storage controller.



### Note

*It is important to keep firmware and drivers up-to-date for optimal performance. Refer to the product release notes (PRNs) on the ATTO web site download page for additional information.*

## System Reboot

A host can be rebooted using its **Basic info** tab. Select "Restart" and optionally provide a message. The host will reboot in one minute. The **Message** area will be updated to reflect the fact that the host has a scheduled reboot and display the time the reboot will occur.

Once a reboot is scheduled, another reboot cannot be scheduled. Pressing the **Refresh** button detects any scheduled reboot and displays the reboot information on the **Basic Info** tab. If the GUI Application attempts

to schedule a reboot and there is already one scheduled, the GUI Application will display an error status and the GUI Application is now aware that a scheduled reboot is active. A scheduled reboot can be cancelled using the **Cancel** button.

The host that is scheduled to reboot displays on its console a pop-up message that indicates the system will be rebooting. The format of the pop-up message is dependent upon the operating system.

*Exhibit 1.0-2 The **Basic Info** taExhibit 1.0-2 on page 6b when you choose a **Host** from the **Device Listing**.*




Exhibit 1.0-3 The **Basic Info** tab when a storage controller is chosen from the **Device Listing**.




Exhibit 1.0-4 The **Flash** tab.





Exhibit 1.0-5 The NVRAM information page.



## 2 ThunderStream SC NVRAM Settings

The settings in the NVRAM tab vary depending upon the selected storage controller and the operating system.

ATTO storage controllers are designed to operate properly using factory settings. Entering invalid or incorrect settings when using an NVRAM configuration utility such as the ATTO ConfigTool may cause your storage controller to function incorrectly.



### CAUTION

**Back up system data when installing or changing hardware configurations.**

Use caution when making changes to NVRAM settings and only make changes to those with which you are familiar. Once you have made the desired changes, click **Commit** to save the changes. Changes do not take effect until you reboot the system.

If you do not want to make any changes, you may choose one of the following:

- **Defaults:** restores the storage controller to factory default settings. The **Commit** button must be clicked to save any changes.
- **Restore:** reverts to the NVRAM settings saved the last time the **Commit** button was used. Clicking **Commit** is not necessary.

### SAS Address

Read only

*Displays the SAS address assigned to the storage controller. The value cannot be modified.*

### Heartbeat

Choices: enabled, disabled

Default: enabled

*When enabled, requires the firmware to respond to periodic activity. If the firmware does not respond, the system driver resets the firmware on the storage controller.*

### NCQ

Choices: enabled, disabled

Default: enabled

*When enabled, the storage controller driver sends multiple simultaneous commands to NCQ capable SATA disk drives.*

### Phy Speed

Choices: Auto, 6Gb/s, 3Gb/s, 1.5 Gb/s

Default: Auto

*Allows the user to manually adjust the speed of the PHY.*

### Drive Wait Time

Choices: 1-255 seconds

Default: 3

*Specifies the number of seconds which the driver waits for devices to appear.*

### Device Wait Count

Choices: 1-255 devices

Default: 1

*Specifies the number of devices which must appear in order to cancel the [Drive Wait Time](#) period.*

### Spinup Delay

Choices: 0-20 seconds

Default: 0

*Specifies the number of seconds each SAS PHY waits for disk drives to spin up.*

### Multiplexing

Choices: enabled, disabled

Default: disabled

*When enabled, it allows multiple 3Gb devices to aggregate 6Gb SAS bandwidth. In order to utilize this feature, devices must support Multiplexing and conform to SAS 2.0 compliancy.*

### 3 RAID Configuration

ThunderStream SC storage controllers provides the capability to configure disk storage into RAID groups or Hot Spare drives.

Use the ATTO ConfigTool to set up RAID groups on your ThunderStream SC storage controller in one of the following RAID levels:

- JBOD
- RAID Level 0
- RAID Level 1
- RAID Level 4
- RAID Level 5
- RAID Level 6
- RAID Level 10
- ATTO DVRAID™ (parity redundancy optimized for digital video environments: refer to [Setting up DVRAID](#) on page 11).

ATTO DVRAID is set up automatically by the ThunderStream SC storage controller firmware. All other types of RAID require customized input. See [Exhibit 3.0-0](#).

Each RAID group may be divided into one or more partitions; each partition appears to the host operating system as a virtual disk.


You may use the RAID Command Line Interface page from the **RAID CLI** tab in the ATTO ConfigTool to set up or modify various parameters (Refer to [Appendix A](#)). However, the ATTO ConfigTool procedures listed in this chapter are the preferred procedures for setting up RAID configurations for ExpressSAS RAID and ThunderStream SC storage controllers.



Note


The ConfigTool screens are similar for all operating systems.

Exhibit 3.0-0 ConfigTool RAID page.




## Preliminary Steps

- 1 Locate the **ConfigTool** icon in the folder you created during installation (Refer to [ATTO ConfigTool](#) on page 1).
- 2 Double-click on the icon to start the application.
- 3 The **Welcome** screen appears.
- 4 From the left-hand panel **Device Listing**, expand all the elements of the selected hosts in the Network tree.
- 5 Click on **ThunderStream SC 3808**.



- 6 A new set of tabs appears in the right panel. Click on the **RAID** tab.



- 7 The application scans for drives.  
A list of drives appears. Devices are displayed in the top panel and RAID groups and Hot Spares are displayed in the bottom panel.  
If you have not yet defined any RAID groups, selecting the **RAID** tab automatically starts a RAID wizard. The wizard is displayed in the bottom panel.  
If the RAID wizard does not start automatically, select the **RAID Management** menu item at the top of the screen, then select the **Create RAID Group** from the menu presented.
- 8 Select either **Setup DVRAID** (continue with [Setting up DVRAID](#) on page 11) or **Custom RAID setup** (continue with [Customizing a RAID setup](#) on page 12).

## Setting up DVRAID

DVRAID (Digital Video RAID) provides parity redundancy for your data. Optimized for performance for the high data transfer rates required in digital video environments, DVRAID is ATTO Technology proprietary technology which supports the editing of uncompressed 10-bit High Definition (HD) video and multiple streams of real-time, uncompressed Standard

Definition (SD) video. The DVRAID wizard automatically sets up DVRAID using all storage attached to the ExpressSAS RAID or ThunderStream SC storage controller based on the number of available drives. See [Exhibit 3.0-1](#). DVRAID setup cannot be used if any other RAID groups are present.

Exhibit 3.0-1 The DVRAID wizard automatically sets up the number of RAID groups based on the number of available drives. Each RAID group uses one drive as a parity drive.

Available drives	RAID groups created	Drives in each group
6	1	6
7	1	7
8	1	8
12	2	6
14	2	7
16	2	8
24	4	6

If you do not want all storage set up in DVRAID or you do not have the correct number of drives, use

[Customizing a RAID setup.](#)

- 1 After following [Preliminary Steps](#) on page 11, select the **DVRAID** radio button in the ConfigTool RAID wizard.
- 2 Click on **OK**.
- 3 The ThunderStream SC storage controller firmware automatically uses all unassigned disks to create a DVRAID configuration.
- 4 A confirmation dialog box asks you to confirm the configuration you have chosen. Click **Yes**.
- 5 A message box displays while the RAID group is being created.

When the RAID group is complete, the lower panel displays the RAID group(s).

The DVRAID group begins to rebuild.

- 6 Double click on your RAID group in the lower panel to see more detail such as the status of the rebuild.

Wait until the rebuild is complete before sending data to the storage.



**Note**

*A RAID rebuild may take several hours to complete.*

- 7 Click on **OK**.

## Customizing a RAID setup

- 1 After following [Preliminary Steps](#) on page 11, select the **Custom RAID** radio button. See [Exhibit 3.0-0](#) on page 10.
- 2 Select the options to configure the new RAID group (see [Exhibit 3.0-2](#) on page 14):
  - **RAID Group Name:** Assign a name to the RAID group. The name must be unique and no more than 14 characters.
  - **RAID Group Level:** Select a RAID group level from the drop-down box.
  - **RAID Group Interleave:** Select an interleave value. The default value is 128KB.
  - **RAID Group Mirror Count:** Select a mirror count, a copy of the original data stored on a separate disk, for RAID groups that have mirrors.
  - **Initialize:** Select the initialization method for the RAID group. The default is **Advanced**.

**Advanced** initialization is recommended for new drives because the procedure erases and verifies the drive media. The RAID group is unavailable until initialization is complete.

**Express** initialization performs RAID group setup in the background and the RAID group is immediately available for use.


- 3 Click **Next**.
- 4 Select the disk devices in the top panel and drag them into the device area in the bottom panel. See [Exhibit 3.0-3](#) on page 14.
- 5 If you want the RAID group to be presented as one virtual disk (partition) with the default RAID group properties, click **Finish**.  
If you want to change other parameters from default values, click **Next** and select the desired property.

- **SpeedRead, Auto-Rebuild and Rebuild Priority:** refer to [Change RAID group properties](#) on page 16 for specific information on these parameters.
- **Sector Size:** specifies the sector size that each partition (virtual disk) within the RAID group presents to the host system. The sector size can be set to 512 bytes (default) or 4096 bytes. This parameter can only be set during RAID group creation. The available size is determined by the sector sizes of the physical disks selected for the RAID group. If the sector size for the physical disks is 512 bytes, then either 512 or 4096 is valid. If the sector size of the physical disks is 4096 bytes, then 4096 is the only valid value for this property.  
The selection of a sector size of 4096 bytes provides the capability to create a 16TB virtual disk that is supported on Windows XP 32-bit version.

- 6 If you want more than one virtual disk (partition) click **Next** and select one of the following:
- **leave as one partition**
  - **partition by count**
  - **partition by size**

When you have made all your selections, click **Finish**.

- 7 A confirmation dialogue box asks you to confirm the configuration you have chosen. Click **Yes**.



- 8 The RAID group configuration you have chosen is initialized and completed. The time it takes to initialize the RAID configuration you have chosen depends on the RAID level selected, the capacity of the drives and the initialization method selected.

Exhibit 3.0-2 Selecting the options to configure the new RAID group.





Exhibit 3.0-3 Selecting drives for a new RAID group.



## Creating a Hot Spare Pool

If a member of a RAID group becomes degraded or fails, you lose some redundancy in your RAID group until a new member is rebuilt into the RAID group.

You can set up a Hot Spare Pool with drives of different sizes which are designated as replacements for faulted devices at any time, either before or after creating RAID groups.

A degraded RAID group is automatically rebuilt if a suitable disk is available in the Hot Spare Pool.

- 1 After following the [Preliminary Steps](#) on page 11, find the **Hot Spare** tab in the bottom panel within the **RAID** tab.

- 2 Select the **Hot Spare** tab to show existing members of the Hot Spare Pool.
- 3 To add drives to the Hot Spare Pool, select unallocated drives from the top panel and drag them to the Hot Spare Pool.

To remove a drive from the Hot Spare Pool, select the drive, click on it and click on **Delete Hot Spares**.



### Note

*An unallocated drive or unallocated storage is storage which is not part of a RAID group, not already designated as a Hot Spare or was offline when you set up a RAID group using the ATTO ConfigTool.*

## Modifying RAID Groups

The ATTO ConfigTool interface may be used to replace a failed drive, add capacity to a RAID group, or change a RAID configuration from the current configuration to a new configuration.



### CAUTION

**Data can be compromised or lost when deleting storage or rearranging storage configurations.**

The ATTO ConfigTool interface takes you step by step through many procedures which allow you to modify your storage and RAID configurations. Read all notes and cautions carefully as you go to ensure the best performance and use of your storage. Many of these procedures are only available using unallocated storage.

Begin each process by following the [Preliminary Steps](#) on page 11 and clicking on the desired process in the **RAID Management** menu.



### Note

*An unallocated drive or unallocated storage is storage which is not part of a RAID group, not designated as a Hot Spare or was offline when you set up a RAID group using the ATTO ConfigTool.*

### Expand capacity

Click on **Expand Capacity** in the **RAID Management** menu and follow the on-screen

instructions. Depending on the RAID configuration, you may need to add more than one device.



### CAUTION

**Adding drives to an existing RAID group may adversely impact performance. You cannot reverse this operation unless you delete the RAID group.**

### Change from one RAID level to a new RAID Level

Changing from one RAID level to another RAID level is called migration. The following migration levels are supported:

- JBOD to RAID Level 0
- JBOD to RAID Level 1
- RAID Level 0 to RAID Level 10
- RAID Level 1 to RAID Level 10

Select **Migrate RAID** in the **RAID Management** menu and follow the on-screen instructions.

### Delete a RAID group

To delete a RAID group using the ATTO ConfigTool, click on **Delete Group** in the **RAID Management** menu and follow the on-screen instructions.




### CAUTION

**Data can be compromised or lost when deleting storage or rearranging storage configurations.**

## Change RAID group properties

A RAID group has properties that are specific to the RAID group. The value of each property remains with the RAID group when it is moved from one computer to another. Some of the properties can only be specified during RAID group creation whereas others may be changed at any time during the life of the RAID group.

- 1 Select a RAID group in the **Groups** panel.
- 2 Click on **Properties** in the **RAID Management** menu.
- 3 View or change the current properties.



- **SpeedRead** specifies the cache policy to be used during read operations. Once a read command is given, the SAS/SATA RAID

storage controller retrieves the next set of sequential data from the RAID group and caches it in internal memory. If you select **Never**, read caching is never performed. If you select **Always**, read caching is always performed. If you select **Adaptive**, the default, SpeedRead is enabled or disabled depending on the sequential patterns detected in I/O requests.

- **Auto-Rebuild** controls the replacement of a faulted drive with any available unallocated drive. When you click on the **Auto-Rebuild** check box and the **Accept** button, Auto-Rebuild is enabled. If a drive becomes faulted, the SAS/SATA RAID storage controller replaces the drive with an unallocated drive.
- **Rebuild Priority** specifies the ratio of rebuild I/O activity to host I/O activity. A rebuild priority of **Same** (default value) indicates that rebuild I/O and host I/O are treated equally. A rebuild priority of low indicates that host I/O is given a higher priority than rebuild I/O. A rebuild priority of **High** indicates that rebuild I/O is given a higher priority than host I/O.
- **Prefetch** specifies the number of stripes that are read when SpeedRead is enabled or adaptive. The valid values for Prefetch are 0, 1, 2, 3, 4, 5 and 6, and the default value is 1. This property can only be changed after the RAID group is created. To access this property, select the RAID group and view its properties.

- 4 Click **Accept**.

## Replacing a faulted drive

If a drive in a RAID group fails, the RAID group's status becomes degraded. To return to optimal functionality, replace the faulted drive using one of the following mechanisms.



### Note

*All of these mechanisms start a RAID group rebuild after the drive is replaced. A RAID group rebuild may take several hours to complete. The RAID group is accessible during rebuild; however, performance may be impacted, depending on the Rebuild Priority assigned.*

## Create a Hot Spare Pool

A faulted drive is automatically replaced if a suitable disk is available in the Hot Spare Pool. You set up a Hot Spare Pool with drives reserved until a RAID group member fails; they are not available when creating a RAID group. Refer to [Creating a Hot Spare Pool](#) on page 15.

## Enable Auto-Rebuild

A faulted drive is automatically replaced if Auto-Rebuild is enabled and a suitable unallocated disk is available. Suitable unallocated drives are initialized, large enough to replace the degraded drive, and cannot contain any RAID group information. The unallocated drive may be a pre-existing drive or a newly-seated drive.

If a Hot Spare Pool exists, the SAS/SATA RAID storage controller chooses a suitable Hot Spare drive before selecting an unallocated drive.

Refer to [Change RAID group properties](#) on page 16.

## Manually replace a drive

You may replace a drive and rebuild a RAID group manually.

- 1 After following the [Preliminary Steps](#) on page 11, double-click on the degraded RAID group in the bottom panel.
- 2 A RAID group members tab displays in the bottom panel. Select an unallocated drive from the drive inventory and drag it over the degraded drive in the members tab.  
If the selected drive is appropriate, the faulted drive is replaced.

## Setup RAID Notification

The ATTO ConfigTool provides a mechanism to issue a notification when a RAID event occurs in ThunderStream SC storage controllers.

RAID events are divided into three categories:

- **Critical events** indicate a serious problem has occurred and the administrator of the RAID group should perform corrective action.
- **Warning events** are less serious but still warrant notification.
- **Information alerts** provide additional useful information about warnings or critical events.

The **Notifications** tab provides numerous drop-down boxes that allow for the selection of the types of RAID events to receive for a notification. There are four options available in these drop-down boxes:

- **Critical:** Only RAID events of a Critical level are reported for the notification.
- **Warning:** Only RAID events of Warning and Critical level are reported for the notification.
- **All:** All RAID events (Information, Warning, Critical) are reported for the notification.
- **None:** No RAID events are reported for the notification. This level is particularly useful if a notification feature is not desired.



The notifications are specified at the host system level and apply to all ATTO SAS/SATA RAID storage controllers installed in the host system.

## Basic alerts

The **Basic Alerts** section of the **Notifications** tab provides the controls to select the levels of RAID events for Audible and Visual notifications, as well as the level of RAID events to output to the system event log.

**Audible Alerts** utilize the system speaker and/or the hardware buzzer on a ThunderStream SC to produce the notification. When triggered, the alarm continuously sounds until it is stopped by the user. There is an icon  in the system tray (or system status area) that is used to turn off the alarm (see Exhibit 1.4-7 and 1.4-8). The user must right-click on the icon and then select **Mute Audible Alert** from the menu. This will stop the alarm. The system tray icon also provides a control that will allow the user to enable/disable the hardware buzzer, so that only the alert via the system speaker is used (see Exhibit 1.4-7 and 1.4-8). This allows the hardware buzzer to be turned off independently of the system speaker. It can also serve a remote purpose. If the daemon is on a server in a server room, the sysadmin can click the control to disable it. The hardware buzzer is enabled by default.



Note

*Audible alerts may not be available on your operating system.*


**Visual Alert** utilize pop-up message boxes to display the contents of a RAID event. The pop-up may be closed using the pop-ups button.

## Logging


The ATTO ConfigTool Service outputs RAID events and other useful informational messages to the system event log. The level of RAID events inserted into the system event log may be controlled by the System Log drop-down box in the Basic Alerts section of the Notifications tab. See [Exhibit 3.0-4](#)

The ConfigTool also generates a log on a per-adapter basis, as well as one for the SNMP functionality and the daemon itself. These are collected with the one-button diagnostics menu item. This can be useful when contacting ATTO Technical Support. Choose **Help, Run Diagnostics**, then choose a place to save the file output. It may take several minutes to gather the log. See [Exhibit 3.0-5](#)

*Exhibit 3.0-4 The **Basic Alerts** controls of the **Notifications** tab in the ATTO ConfigTool.*



*Exhibit 3.0-5 Run Diagnostics generates an extensive log to aid in troubleshooting.*



## Email

The ATTO ConfigTool Service periodically sends email notifications to the designated email address(es) that have been input under the Notification Addresses: field(s). There are three text fields under which email addresses can be specified, and more than one email address may be specified in each text field, as long as they are separated by a comma. Each email address text field can be configured to receive a specific level

of RAID event via the drop-down box displayed next to each email address text field.

The ATTO ConfigTool Service can be configured to use a TLS/SSL email server (eg. Gmail™, Yahoo®, etc.) as well as the port to use when connecting to the email server. The ATTO ConfigTool Service's connection to the email server may be configured using the following controls:

- **Server Address:** specifies the address of the SMTP email server to use when sending an email (eg. smtp.example.com)
- **Sender Address:** specifies the text that will show in the "From" field in the sent email. This may contain an actual email address, or any other text that may be useful in identifying the sent email (eg. example@example.domainexample, server123, etc.)
- **Username:** specifies the username or login ID that is required when logging into the email server to send an email
- **Password:** specifies the password for the account identified by the Username.
- **Enable SSL:** when checked, the ATTO ConfigTool Service will attempt to connect to the email server using the SSL protocol. When unchecked, the ATTO ConfigTool Service will attempt to connect to the email server using an unsecured channel.
- **Port:** specifies the port number to which the ATTO ConfigTool Service attempts to connect to the email server. Valid values are 1-65535.



### Note

*Contact your email provider for the appropriate Port number. Using any other number will result in an email notification failure.*

Email notifications are sent at 15-minute intervals. If a Critical RAID event is detected, an email notification will be sent after a 10-second delay. This delay allows the ATTO ConfigTool Service to gather supporting RAID events that may be useful in identifying why the Critical RAID event occurred.

Exhibit 3.0-6 The **Notifications** tab in the ATTO ConfigTool.





Exhibit 3.0-7 Mute Audible Alert (Mac OS X)



# Using SCSI Enclosure Services (SES)

---

SAS/SATA drive enclosures may provide a SCSI Enclosure Processor which indicates enclosure health status, drive identification and drive fault identification.

The ATTO ConfigTool recognizes drive enclosures that provide SCSI Enclosure Services (SES). You may use SES to identify individual drives, all the drives in the same enclosure, all the drives in a single RAID

group, or faulted drives. You may also select drives and monitor the status of the enclosure.

To use SES, open the ATTO ConfigTool and follow one of the procedures below.

## Identifying Drives

---

**Drive Identification** lights LEDs showing the drives you have selected using the ATTO ConfigTool. Most drive enclosures blink an LED next to the drive in the enclosure.

The ConfigTool provides a way to identify individually selected drives, all of the drives in the same enclosure, all drives in a RAID group, a selected mirror group in RAID 1 or RAID 10, and faulted drives.



### Note

*Drive Identification works for drives that are not controlled by a SCSI Enclosure Service. The ATTO ConfigTool provides functionality to locate all or specifically selected drives in a RAID group, when a SES processor is not present.*

- 1 Select one or more drives individually or in enclosures or drives in RAID groups:
  - Select individual drives in the **Attached Drives** panel.
  - Select one or more RAID groups in the **Groups** panel.
  - Select one or more drives in the **Attached Drives** panel, right click on one of the selected drives and select **Enclosure**.

- 2 Right-click on one of the selected drives or RAID groups.
- 3 Select **Locate**.



### Note

*If a RAID 1 or RAID 10 group was selected, the locate operation provides the capability to select a specific mirror of the group.*

- The status icon next to the selected drives blinks and the enclosure performs its specific identification method until you stop it.
- 4 To stop the drive identification, right-click on one of the selected drives or RAID groups.
  - 5 De-select **Locate**.

## Identify Faulted Drives

**Drive Fault Identification** is performed automatically by the ThunderStream SC storage controller when a member of a RAID group becomes degraded by exhibiting unrecoverable errors during I/O.

The ThunderStream SC storage controller reports the status of the drive and asks the SES device to perform fault identification. The SES device usually illuminates a blinking red LED as a fault identification.

The fault identification continues until the drive is replaced or the RAID group is deleted.

## Selecting All Drives in the Same Enclosure

---

**Drive Selection** selects all drives in the same enclosure using the ConfigTool. You may use drive selection if you need to select all the drives in one enclosure which are attached to a RAID storage controller without selecting other enclosures, or if you are selecting members for a RAID group during RAID group creation.

- 1 Right click on one drive in the **Attached Drives** panel.
- 2 Select **Enclosure**.
- 3 Select **Drives**. All drives in the same enclosure as the selected drive are selected.



## Monitoring S.M.A.R.T. Data

*Self-Monitoring, Analysis and Reporting Technology, or S.M.A.R.T., is a system built into SATA drives to detect and report on various indicators of drive health.*


The S.M.A.R.T. (Self-Monitoring, Analysis and Reporting Technology) monitoring feature monitors and reports the status of SATA drives using certain parameters recorded by the drives. Notifications are sent when the values exceed certain pre-determined values.

Use the ATTO ConfigTool to view the files that record changes to S.M.A.R.T. parameters. The files are permanent and can be viewed independently whether you have enabled monitoring or not.

## Enabling or Disabling S.M.A.R.T. Monitoring


You may enable or disable the monitoring feature at any time. Monitoring is disabled by default: if you want to use the feature, you must enable it.

- 1 Select the **ThunderStream SC** storage controller from the **Device Listing** panel.



- 2 Select the **RAID** tab in the right panel.
- 3 Select the **RAID Management** menu item at the top of the screen.

- 4 Select the **Monitor S.M.A.R.T.** menu item to change the monitoring status. The **Monitor S.M.A.R.T** menu item under the **RAID Management** menu has a check mark when monitoring is enabled and no mark when it is disabled.



## Checking S.M.A.R.T. Status

The ATTO ConfigTool interface displays the latest S.M.A.R.T. status record for a selected drive. All attributes reported by the drive are listed with each attribute's **Threshold**, **Worst**, **Current** and **Raw** value.

The threshold value is the value at which notification of a problem is generated by the software.

If there has been a change from a previous record of S.M.A.R.T. status, an arrow notes which way the change went, either higher or lower. For example, in

[Exhibit 3.0-9](#) on page 23, the temperature listed in this record is lower than the temperature listed in a previous record and the arrow next to that attribute points downward.

The S.M.A.R.T. status display also contains information such as the date and time the S.M.A.R.T. status was recorded, the total number of records for this drive, and the current monitoring status (enabled or disabled).

You may move to previous or subsequent records, query the drive or refresh the view using controls on the interface.

- 1 Select a single drive in the **Attached Drives** panel.
- 2 Right-click on the selected drive.
- 3 Select **S.M.A.R.T. Status** in the sub-menu.
- 4 The **S.M.A.R.T. Status** box displays.
  - Use the left arrow or right arrow control to move between S.M.A.R.T. status records


- Use the **Refresh** button to query the drive for the latest values. If any values are different from the most recent record, a new record is created and displays.



**Note**

*If you click on the **Refresh** button when monitoring is disabled, a pop-up box displays. You can enable monitoring from the pop-up box to complete the refresh request.*

Exhibit 3.0-9 S.M.A.R.T. status for a selected drive.



## Filtering S.M.A.R.T. Attributes

Each of the S.M.A.R.T. status attributes is assigned one or more classification types:


- Performance
- Error Rate
- Event Count
- Critical

The **S.M.A.R.T. Status** dialog box can be filtered to display any combination of these types.

The default view is to display all types.

- 1 Open the **S.M.A.R.T. Status** box as described above.
- 2 Right click in the table area where the attribute values are displayed.
- 3 Each classification type that is visible has a check mark. (See [Exhibit 3.0-10](#)). Select any classification type to change the check mark.

Exhibit 3.0-10 S.M.A.R.T. Status box with classification types.



## S.M.A.R.T. Notifications

S.M.A.R.T. status is collected from each SATA drive at 60 minute intervals and, if the data is different than the previous status, a S.M.A.R.T. status record is added to the S.M.A.R.T. status file for that drive.

A notification of the S.M.A.R.T. status difference is generated based upon the current settings in the **Notifications** tab. Refer to [Setup RAID Notification](#) on page 17.

The notification level of S.M.A.R.T. status is determined as follows:

- **Critical:** One or more of the status values was below a threshold value and one was classified as a critical attribute.
- **Warning:** One or more of the status values was below a threshold value but none were classified as a critical attribute.
- **Information:** None of the status values were below the threshold value.

## Test Drive Performance

The performance testing feature in the ATTO ConfigTool records the performance level of a drive in a RAID group under an I/O (input/output) load generated by a host application. The information can be used to compare the performance of each drive relative to others.

The performance feature of the ATTO ConfigTool counts the Megabytes per second (MB/sec.) which elapses from the time a start button is clicked until a stop button is clicked.

Drive performance can be measured for any combination of single drives within a RAID group or across multiple RAID groups, or for all the drives in a selected RAID group.



### Note


*The performance test can only be performed on drives which belong to a RAID group.*

The **Performance** dialog box contains individual statistics for each drive selected. You should start I/O to the drives to be tested before beginning the test.

The collection of these statistics begins when the **Start** button is pressed and continues until the **Stop** button is pressed. You may press the **Reset** button and restart a test at any time.

The ATTO ConfigTool can be closed after starting the performance test and re-opened to show the statistics from the most recent test.

- 1 Start I/O from the host system.
- 2 Select the **ThunderStream SC** storage controller from the **Device Listing** panel.



- 3 Select the **RAID** tab in the right panel.
- 4 Select the drives to be tested:
  - Select individual drives from the **Attached Drives** panel.
  - Select a single RAID group in the **RAID Group** panel.
- 5 Right-click on the selected drive(s) or RAID group.
- 6 Select the **Performance** menu item.
- 7 The **Performance** dialog box displays. (See [Exhibit 3.0-11.](#))
- 8 Click the **Start** button.
- 9 Click the **Stop** button to halt monitoring. Click the **Reset** button to reset the performance values to zero at any time.

You can close the **Performance** dialog box while monitoring is in progress, and then open it later to see the performance results.
- 10 Stop I/O any time after you have stopped the performance test.

Exhibit 3.0-11 Detail of the **Attached Drives** window and the **Performance** dialog box with performance results for the selected drives before a test has been run.

Drives: 8    Faulted: 0    Raid Groups: 1    Hot Spares: 0

Attached Drives

Name	Vendor	Address	Usage
WD2002FYPS-02W3B	WDC	50:0 1:08:60:00:52:17:a8	r5x8 (1)
WD2002FYPS-01U1B	WDC	50:0 1:08:60:00:52:17:a9	r5x8 (8)
WD2002FYPS-02W3B	WDC	50:0 1:08:60:00:52:17:aa	r5x8 (4)
WD2002FYPS-01U1B	WDC	50:0 1:08:60:00:52:17:ab	r5x8 (3)
WD2002FYPS-01U1B	WDC	50:0 1:08:60:00:52:17:ac	r5x8 (5)
WD2002FYPS-01U1B	WDC	50:0 1:08:60:00:52:17:ad	r5x8 (6)
WD2002FYPS-01U1B	WDC	50:0 1:08:60:00:52:17:ae	r5x8 (7)
WD2002FYPS-01U1B	WDC	50:0 1:08:60:00:52:17:af	r5x8 (2)

Performance for r5x8

Member	Transferred	Max Time	Read Rate	Write Rate	Errors
1	0.00 bytes	0.00 ms	0.00 bytes/s	0.00 bytes/s	0
2	0.00 bytes	0.00 ms	0.00 bytes/s	0.00 bytes/s	0
3	0.00 bytes	0.00 ms	0.00 bytes/s	0.00 bytes/s	0
4	0.00 bytes	0.00 ms	0.00 bytes/s	0.00 bytes/s	0
5	0.00 bytes	0.00 ms	0.00 bytes/s	0.00 bytes/s	0
6	0.00 bytes	0.00 ms	0.00 bytes/s	0.00 bytes/s	0
7	0.00 bytes	0.00 ms	0.00 bytes/s	0.00 bytes/s	0
8	0.00 bytes	0.00 ms	0.00 bytes/s	0.00 bytes/s	0

Start    Stop    Reset

## RAID Media Scan Feature

---

The Media Scan feature scans disk drives for media errors and parity errors. All media errors are counted and fixed. All parity errors are reported in the event log. The two options are described below:

**Media Scan** - Media Scan works with parity RAID group members, Hot Spare drives and unallocated drives. Media Scan reads the selected drives and, if a Media Error is found, Media Scan re-writes the disk with the media error. The drive will relocate the bad sector to an alternate part of the drive. Media Scan re-writes the correct data for a disk that is a member of an on-line parity RAID group. Media Scan writes random data to Hot Spare drives and unallocated drives. The Media Scan feature records the number of media errors detected and corrected for each drive. These counts are stored persistently if the drive is a RAID group member or a Hot Spare drive. The counts are not persistently stored for any other drives.


**Media Scan with Parity Verify** - is a variation of Media Scan that is available for on-line parity RAID groups only. Parity Verification is performed on each RAID group stripe that has no detected media errors. The parity of the stripe is recalculated and compared to the original parity for the stripe. If there is a mismatch, an error is generated and recorded in the Event Log and System Log files.

### Modes of operation for:

- On-line RAID Group – Media Scan performs SCSI Read commands to each stripe group. Parity verification is performed on a stripe that has no media errors. Media Scan activity is scheduled in accordance with the Rebuild Priority Level. Media Scan is restarted after reboot if the Media Scan did not complete.
- Degraded or Offline RAID Group – Media Scan performs SCSI Read commands for each stripe group of the on-line drives. No parity verification occurs. The Media Scan is restarted after a re-boot if the Media Scan did not complete.
- Hot Spares & Unallocated Drives – Media Scan performs SCSI Read commands for each selected drive. Media Scan is not restarted after a re-boot if the Media Scan did not complete.


# Starting a Media Scan

- 1 Select the **ThunderStream SC** storage controller from the **Device Listing** panel.
- 2 Select the **RAID** tab in the right panel.
- 3 Select the RAID group, Hot Spare drive(s) or unallocated drive(s) which will have a Media Scan. One RAID group can be selected or multiple Hot Spare and unallocated drives can be selected.
- 4 Select the **RAID Management** menu item at the top of the screen.
- 5 Select the **Scanning** menu item.
- 6 Select the **Media Scan** or **Parity Verify** from the sub-menu.
- 7 The Media Scan starts as soon as the sub-menu item is selected.




# Media Scan Status

The Media Scan starts as soon as the menu item is selected. The **Attached Drives** panel displays a rotating icon next to each drive being scanned. The **RAID Group** panel displays the type of scan being performed and a status for percent complete.



## Viewing the Scan Report

- 1 Select the **ThunderStream SC** storage controller from the **Device Listing** panel.
- 2 Select the **RAID** tab in the right panel.
- 3 Select the RAID group, Hot Spare drive(s) or unallocated drive(s) whose scan report should be displayed. One RAID group can be selected or multiple Hot Spare and unallocated drives can be selected.
- 4 Select the **RAID Management** menu item at the top of the screen.
- 5 Select the  **Scan Report** from the menu.
- 6 A **Media Scan Error Report** message box appears with the scan results. If a scan is currently in process the **Scanned** column indicates the percent complete for the operation.



Device	Scanned	Errors	Total Errors	Scans
WD2002FYP...	4%	0	0	0
WD2002FYP...	4%	0	0	0
WD2002FYP...	4%	0	0	0
WD2002FYP...	4%	0	0	0
WD2002FYP...	4%	0	0	0
WD2002FYP...	4%	0	0	0
WD2002FYP...	4%	0	0	0
WD2002FYP...	4%	0	0	0

The Scan Report includes the counts for the most recent scan and the total results for all Media Scan operations. This information is persistent for RAID group drives and Hot Spare drives since it is stored in meta-data of the drive. The results may be erased from the drive when the meta-data is re-written. e.g. The information is lost from a Hot Spare when the Hot Spare is added to a RAID group. The results of a Media Scan for an unallocated drive are only available until a reboot occurs.

Media Scan provides the following metrics:

- 1 Media errors detected on most recent scan in the **Errors** column.
- 2 Media errors corrected on most recent scan.
- 3 Total stripe groups scanned on most recent scan. This data is not displayed in Configuration Tool.
- 4 Media errors detected on all scans in the **Total Errors** column.
- 5 Media errors corrected on all scans
- 6 Total number of all scans in the **Scans** column.


The **Media Scan Error Report** does not display the count of corrected errors. The corrected error count is displayed while the mouse pointer hovers over the **Errors** or **Total Errors** value for a specific drive. If the detected errors do not equal the corrected errors the warning icon  appears to the left of device name column.

The Scan Report is updated by the Configuration Tool every 10 seconds. The metrics are not updated when the Media Scan is paused.

## Pausing or Resuming Media Scan

Media Scan for a parity RAID group can be paused, resumed and cancelled. The Configuration Tool provides menu items to pause and resume and CLI provides a command to cancel the Media Scan. The results of the Media Scan up to the time of the pause or cancel are saved persistently.

Media Scan for Hot Spare and unallocated drives can be cancelled and cannot be paused. Media Scan of Hot Plug and unallocated drives are run immediately and they are not scheduled. The Configuration Tool provides menu items to start and cancel Media Scan on these drives.




## Scheduling Media Scan

Media Scan for a parity RAID group can be scheduled to occur at a specified time for a specified frequency. The scheduling mechanism provides the means to specify the following:

- Day of the week and the time of day to run a Media Scan once a week
- Day of the week and the time of day to run a media Scan once
- Time of day to run a media Scan every day of week
- A media Scan is scheduled using the following procedure.

- 1 Select the **ThunderStream SC** storage controller from the **Device Listing** panel.
- 2 Select the **RAID** tab in the right panel.
- 3 Select the parity RAID group that should have a scheduled Media Scan.
- 4 Select the **RAID Management** menu item at the top of the screen.
- 5 Select the  **Schedule...** from the menu.
- 6 Select **Scan Only** or **Parity Verify** and press the **Next** button.



**Schedule A Task**


Select a Type of Media Scan

Scan Only

Parity Verify

Back Next Finish Cancel

- 7 Select **Once**, **Daily** or **Weekly** as the frequency of the task.



**Schedule A Task**

Select a Frequency to Perform the Task

Time: 11:15 AM

Day: Thursday

Once

Daily


Weekly

Back Next Finish Cancel

- 8 Select the time of day to run the task.
- 9 Select the day of week to run the task. The day of week field is disabled when the frequency is daily.
- 10 Press the **Finish** button and a confirmation dialog box is displayed.
- 11 Press **Yes** to schedule the task. Media Scan for Hot Spare and unallocated drives cannot be scheduled.

## Viewing Scheduled Tasks

The Configuration Tool contains a Task tab for each RAID storage controller present in the system. The tab contains a table that displays information for all of the currently scheduled tasks for the associated storage controller. The table only contains an entry for items that have not occurred to date. It does not show any information for completed tasks. The table provides the capability to cancel a task or to change the schedule for the task. You cannot schedule a new task from this view.



## Automatic Cancellation of Media Scan

A Media Scan can be automatically cancelled by the FW for the following reasons.

- A parity RAID group is being scanned and it requires a rebuild operation
- A parity RAID group is being scanned and the last drive is marked 'Faulted'
- A Hot Spare is being scanned and it is needed for a RAID group rebuild
- An unallocated drive is being scanned and it is needed for a RAID group rebuild

## Data Recovery for OFFLINE RAID Groups

---

ATTO SAS/SATA RAID storage controllers use DriveAssure™ to maintain error-free operation of RAID Group member drives. This includes features like Drive Initialization, Media Scan, Parity Verify, Performance Monitoring, and algorithms for latency and response management. Occasionally, a drive will suffer a mechanical or fatal media error that makes it unusable. A connection error may cause multiple drives to be inaccessible, resulting in an array that cannot be accessed normally. This document describes procedures to follow that may allow data to be partially recovered when drives fail or become inaccessible and cause the ATTO RAID Group to go OFFLINE.



### Note

Anytime a RAID Group goes OFFLINE, data integrity may be compromised (unless the conversion was halted by the user). It is imperative that data consistency checks and/or file system repair tools are used to validate recovered data.

## Definitions

---

Exhibit 3.0-12 RAID Group Status Definitions

RAID Group Status	Definition
ONLINE	The RAID Group state is ONLINE for RAID Groups with all members available and fully operational.
DEGRADED	RAID Group does not have full data protection. This is the case when a non critical number of RAID Group Members are UNAVAILABLE, DEGRADED or FAULTED (i.e., 1 member in a RAID5 Group, 2 in a RAID6 Group, and so forth). On a READ command to inaccessible drive(s), the data is derived using redundancy or parity. A DEGRADED RAID Group may be initializing or rebuilding.
OFFLINE	RAID Group data cannot be accessed because a critical number of drives have FAULTED or been removed. For RAID 5, for example, two or more drives have failed.
RECOVERY BASIC	RAID Group is in BASIC data recovery mode. I/O may be limited to reads.
RECOVERY EXTREME	RAID Group is in EXTREME data recovery mode, which returns data at all costs. I/O may be limited to reads.
WAITING	The transient state of a group <b>while</b> it is being discovered.
DELETED	The transient state of a group <b>while</b> it is being deleted.
NEW	The state of a group that has been created but not yet committed.

Exhibit 3.0-13 RAID Member Status Values

<b>Rebuild Status Values</b>	
Ok	Data is up to date
Faulted	The operation was stopped due to an error
Halted	The operation was stopped by the user
Interrupted	The operation was stopped due to a disk removal or firmware restart
<b>Rebuild Types</b>	
Unknown	A pseudo value used if the type of operation is not recognized
Rebuilding	Represents the operation where a degraded group is repaired
Erasing	Represents the operation where the data contained in a group is purged by writing to every sector of the RAID group
Writing	Represents the operation where a specific data pattern is written to every sector of the RAID group
Converting	Represents the operation where data is shuffled around within a group either because the group is being converted from one RAID level to another, or the capacity of a group is being expanded
Initializing or Advanced Initializing	Represents the operation where new disks added to the group are undergoing full initialization. This can happen either when the group is created, migrated or expanded
Initializing or Express Initializing	Represents the operation where the group undergoes a parity rebuild (if applicable) for the new members added. This can happen either when the group is created, migrated or expanded
Media Scanning	Represents the operation where I/O is performed to each block on a device to force the reallocation of bad sectors. All error correcting is handled automatically by the device, and any errors detected will increment the appropriate counters
Parity Verifying	Represents the sector scan operation that also verifies the parity blocks on a RAID group are correct. If errors are found, they are left uncorrected, and the appropriate error counters are incremented
Parity Repairing	Represents the parity scan operation that also attempts to fix any parity errors that are encountered. Any errors detected will increment the appropriate error counters
Recovery Rebuilding	Represents the operation where an offline group is forced online prior to rebuilding it

## Failure Scenario Table

RAID Groups cannot be accessed normally when their member disks fail, and the RAID Group is marked OFFLINE. RAID Groups of different RAID levels are marked OFFLINE for different reasons, as follows:

Exhibit 3.0-14 Failure Scenarios and Recovery Methods

RAID Level	Reason(s) for being marked OFFLINE	Recovery Method
JBOD and RAID 1	Any drive failure.	See <a href="#">Faults on Critical Number of Drives</a> on page 42
RAID 1 and RAID 10	Error during rebuild	See <a href="#">Recovery from Failed Rebuild</a> on page 40
	Mistaken replacement of a good drive when its mirror has failed	See <a href="#">Recovery from Replacement of Wrong Drive</a> on page 39
RAID 4 and RAID 5	Errors on two or more drives	See <a href="#">Faults on Critical Number of Drives</a> on page 42
	Error during rebuild	See <a href="#">Recovery from Failed Rebuild</a> on page 40
	Mistaken replacement of a good drive when another member of the RAID Group has failed	See <a href="#">Recovery from Replacement of Wrong Drive</a> on page 39
RAID 6	Errors on three or more drives	See <a href="#">Faults on Critical Number of Drives</a> on page 42
	Error during rebuild	See <a href="#">Recovery from Failed Rebuild</a> on page 40
	Mistaken replacement of good drive(s) when other members of the RAID Group have failed	See <a href="#">Recovery from Replacement of Wrong Drive</a> on page 39

## Drive Replacement on a Failure Condition

### Replacing RAID Group Member Drives as Soon as They Fail

With parity and redundancy RAID levels, the RAID Group can withstand the loss of one member, and the data is still valid and accessible. In this case, the RAID Group goes into DEGRADED mode and uses parity or redundancy to generate the data. Although the RAID Group is fully operational, the RAID Group is at risk because if any other drive fails, data integrity is called into question.

### A Warning about Drive Replacement

A very common reason that an array goes from DEGRADED mode to OFFLINE mode is when the wrong drive is replaced. By pulling out a perfectly good drive, a double-drive fault occurs and there are insufficient drives to generate data. The following procedure is very important when you are considering removing a failed drive, to ensure the correct drive is pulled.

## Identifying Failed Drives

If a failed drive is in an enclosure that supports SES (SCSI Enclosure Services), the drive's Fault LED should be blinking. In that case, it is clear which drive should be replaced. If multiple drives are blinking, power cycling the array and/or the drives, and reseating the drives can sometimes correct intermittent conditions.

ATTO supplies other facilities to aid in failed drive identification. When there is some question about which drive has failed, use either the CLI command-line interface or the ATTO ConfigTool procedures described below to be very sure as to which drive failed, prior to replacing the drive.

## Using ATTO ConfigTool to Identify a Failed Drive

The ATTO ConfigTool provides a graphical display of RAID Groups and attached drives. [Exhibit 3.0-15](#), below, shows a RAID Group with a status of DEGRADED. The upper window shows attached disks and their association with RAID Groups.

Exhibit 3.0-15 RAID Display of Degraded Group

The screenshot shows the ATTO ConfigTool interface. On the left is a 'Device Listing' tree showing the hardware configuration. The main window has tabs for 'Basic Info', 'Flash', 'RAID', 'RAID CLI', and 'Tasks'. The 'RAID' tab is active, showing summary statistics: 'Drives: 7', 'Faulted: 0', 'Raid Groups: 1', and 'Hot Spares: 0'. Below this is a table of 'Attached Drives' with columns for Name, Vendor, Address, Usage, and Capacity. The bottom section shows a 'Groups' tab for 'R5x3' with a table of RAID members and their status.

Name	Vendor	Address	Usage	Capacity
ST336754SS	SEAGATE	50:00:c5:00:00:54:7...		34.18 GB
ST336754SS	SEAGATE	50:00:c5:00:00:54:2...		34.18 GB
ST336754SS	SEAGATE	50:00:c5:00:00:53:b...		34.18 GB
ST336754SS	SEAGATE	50:00:c5:00:00:53:b...		34.18 GB
ST336754SS	SEAGATE	50:00:c5:00:00:54:2... R5x3 (1)		34.18 GB
ST336754SS	SEAGATE	50:00:c5:00:00:52:4... R5x3 (2)		34.18 GB
ST336754SS	SEAGATE	50:00:c5:00:00:54:7... R5x3 (3)		34.18 GB


  

Member	Status	Rebuild Status
1	Degraded	Not applicable
2	Online	Not applicable
3	Online	Not applicable

By double-clicking the RAID Group line, each RAID Member and its status is displayed. In [Exhibit 3.0-16](#), below, RAID Member 1 indicates DEGRADED.

If the drive is in an enclosure that supports SES (SCSI Enclosure Services), the drive's Fault LED should be blinking.


Exhibit 3.0-16 Member List for Degraded RAID Group



The member drive can be physically identified, in its enclosure, by right-clicking in the Attached Drives list and selecting Locate (See [Exhibit 3.0-17](#) below). The drive's Identify LED (or activity LED, in a non-SES enclosure) should blink.

Sometimes the activity LED won't blink because of a failure in the drive or enclosure electronics. One possible action is to blink all the other LEDs, and find the failed drive by process of elimination.

Exhibit 3.0-17 Identify a Drive Using Locate



## Recovering from Failed Drives

### Recovery from Replacement of Wrong Drive

Exhibit 3.0-18 RAID 5 Group with 4 Operational Drives




Exhibit 3.0-19 A Drive Fails; RAID Group goes DEGRADED





Exhibit 3.0-20 Drive 3 is replaced instead of drive 4! RAID Group goes OFFLINE



When a drive fails, but the wrong drive is replaced, a rebuild will not initiate and the RAID Group will be taken off line. The user should follow the procedure described earlier in this section to identify the failed drive, and ensure the correct drive is replaced.

The situation described above can be corrected. Replace the erroneously replaced drive (3A) with the original drive 3 in the above example; then identify and replace the failed drive (4); and then start a rebuild.

## Recovery Mode

Sometimes, despite careful operation and maintenance, drives will coincidentally fail in such a way that the RAID Group integrity is compromised. After a RAID Group has been marked OFFLINE because of problems with member drives, there is a way to possibly recover some of the data. The following guidelines and commands can help recover data from an OFFLINE RAID Group. The following descriptions refer to RAID 5 specifically, but the principles extend to other RAID types.

## Recovery from Failed Rebuild

### Fault During Rebuild - Recover by Forcing Rebuild to Continue



#### Note

The diagrams which follow illustrate a scenario where a failed drive is removed and replaced. The error recovery procedure is also valid if there is a Hot Spare, or if there is an unallocated drive and Auto Rebuild is enabled.

Exhibit 3.0-21 RAID 5 Group with 4 Operational Drives




Exhibit 3.0-22 A Drive Fails; RAID Group goes DEGRADED




Exhibit 3.0-23 Drive is Removed and Replaced






Exhibit 3.0-24 Media Error During Rebuild; Rebuild Halts, RAID Group goes OFFLINE



Use the ATTO ConfigTool to force the RAID Group to continue the rebuild. First, display the RAID Group information as shown in [Exhibit 3.0-15](#). Then right-click on the Offline RAID Group that has the Rebuild status **Rebuilding (Faulted)**. Select the option Start Recovery Rebuild as shown below in [Exhibit 3.0-25](#).


Exhibit 3.0-25 Enable Rebuild Recovery



The RAID Group status will change from **Offline** to **Degraded**. The Rebuild status will change from **Rebuilding (Faulted)** to **Recovery Rebuilding (n%)**, where n% is the completion percentage of the rebuild process.

At the end of the Rebuild Recovery, barring major problems, the RAID Group will be put back online. However, at the location corresponding to the Media Error, the data may be invalid! Anytime a RAID Group goes OFFLINE, data integrity **has been compromised**. It is imperative that data consistency checks and file system repair tools are used to validate the data before trying to use it.

Exhibit 3.0-26 Recover Rebuild; Rebuild Completes; RAID Group goes ONLINE



**Note**

*If the rebuild cannot continue to the end because of unrecoverable errors on multiple drives, use the Extreme recovery mode, described below, to try to recover data.*

## Serious Disk Failure During Rebuild - Using "Extreme" Mode

Exhibit 3.0-27 RAID 5 Group with 4 Operational Drives




Exhibit 3.0-28 A drive Fails; RAID Group goes DEGRADED




Exhibit 3.0-29 Drive is Removed and Replaced, and Rebuild Starts




Exhibit 3.0-30 Second Unrecoverable Failure; Rebuild Halts, RAID Group goes OFFLINE





Exhibit 3.0-31 Recover Extreme: All Drives Are Used to Try to Recover Data



When a serious disk failure occurs during a rebuild, the rebuild cannot be forced to continue, but the Extreme recovery mode may be viable. If there is an available slot, reinstall the original failed drive (the drive provides another possible source of data even if the data is stale). This original failed drive will only be accessed if there is an error on one of the other RAID Group members. Note that in this failure scenario, data may be extensively compromised.

Extreme Recovery Mode is only accessible from the Command Line Interface. Use the CLI command `RGRecover Extreme` to enter the "extreme, read-at-all-costs" recovery mode. Use a host application to read the data (whatever you can) off of the drives and back it up to another location. The RAID Group will stay in recovery mode. If you use the CLI command `RGRecover Disable` to disable this mode, the RAID

Group will return to the OFFLINE state. There is no mechanism to put this RAID Group back online.

Once you have the data backed up, use only good drives to create a new RAID Group, and then restore the data to the new RAID Group.

## Faults on Critical Number of Drives

Exhibit 3.0-32 RAID 5 Group with 4 Operational Drives




Exhibit 3.0-33 A drive Fails; RAID Group goes DEGRADED





Exhibit 3.0-34 A Media Error occurs; RAID Group goes OFFLINE




In this scenario, there is no Hot Spare drive, and either Auto Rebuild is not enabled or there is no spare drive for Auto Rebuild to use. When the first drive fails, the array will continue to operate in DEGRADED mode and return data when requested. However, once another error occurs, the array is OFFLINE and data has been compromised. Basic recovery mode should be used.

## Basic Recovery Mode

CLI commands or the GUI can be used to put the drives into basic recovery mode. This allows you to read data to a backup location, replace all failed or marginal drives, and then restore the data to the new drives. However, at the location corresponding to the Media Error, the data may be invalid! Anytime a RAID Group goes OFFLINE, data integrity **has been compromised**. It is imperative that data consistency checks and file system repair tools are used to validate the data.

To enter recovery mode from the ATTO ConfigTool, select the Offline RAID Group and right-click to get the menu options. Select the option Basic Rebuild as shown below.

Exhibit 3.0-35 Enable Basic Recovery Mode



The RAID Group status will change from **Offline** to **Recovery (Basic)**. Once you enter Recovery Mode, use a host application to read the data (whatever you can) off of the drives and back it up to another location. The RAID Group will stay in recovery mode. If you use the CLI command `RGRecover Disable` to disable this mode, the RAID Group will return to the OFFLINE state. There is no mechanism to put this RAID Group back online.

Once you have the data backed up, use only good drives to create a new RAID Group, and then restore the data to the new RAID Group.

### Extreme Recovery Mode

When a critical number of drives fail, or if a Rebuild Recovery operation fails, the Extreme Recovery Mode may be an option. This will only work if one or both of the failed drives are still operational at some level, there are spare slots, and the data is not too far out of date. When this is the case, re-insert one or both drives into the array and enter Extreme Recovery mode. The original failed drives will only be accessed if there is an error on the other (non-failed) RAID Group drives. Follow the procedure outlined in section, [Serious Disk Failure During Rebuild - Using "Extreme" Mode](#) on page 42.

# File System Repair Tools

---

When a RAID array has been made accessible by setting the drive states appropriately, the host's file system may still have integrity problems. Most operating systems have utilities that try to detect and repair file access issues.

## Windows: CHKDSK

Chkdsk (Chkdsk.exe) is a Windows command-line tool that checks formatted disks for problems. Chkdsk then tries to repair any problems that it finds. For example, Chkdsk can repair problems related to bad sectors, lost clusters, cross-linked files, and directory errors. To use Chkdsk, you must log on as an administrator.

Click **Repair Disk** to verify and repair any problems on the selected volume(s). The results of the repair will be posted in the window. When no problems are found, it reports: "The volume "diskname" appears to be OK." Occasionally, at the end of a scan, a message similar to the following will appear: "Scan complete. Problems were found, but Disk First Aid cannot repair them."

## Linux: FSCK

The system utility fsck ("file system check") performs a file system consistency check and interactive repair. It has options to do a quick check, preen files, preen and clean files, and others.

## Mac OS X: First Aid (part of Disk Utilities)

Launch Disk Utility from OS X or the OS X Installer In the column on the left, select the volume you want to repair, and then on the right, click the **First Aid** tab.



### Note

*If there are problems found, you should click **Repair Disk** to scan again. If the same problem(s) are found during the second scan, it means First Aid cannot repair the disk. This indicates that some of your files may be more severely damaged. To repair this damage, you can either back up your disk and reformat, or obtain a more comprehensive disk repair program.*

# Command Line Interface

---

## Using CLI to Identify a Failed Drive

### Blockdevscan

The `blockdevscan` CLI command shows a list of all attached drives and their usage. Use this command to update the list of drive IDs, which are needed by other CLI commands.

```
blockdevscan

6
;ID Vendor      Product                SerialNumber Capacity InUse      Port
;-----
0          ST3250310NS          9SF0800F 232.88GB r5
1          ST3250310NS          9SF076NJ 232.88GB r5
2          ST3160812AS-1       5LS25P51 149.10GB r5
3          ST3250620AS          6QF0LJ17 232.88GB r5

Ready.
```

For members of a RAID Group, the InUse column shows the RAID Group name.

## RGdisplay

This is the RAID Group display command. Use this command to display all RAID Groups and their statuses.

```
rgdisplay

3
;GroupName          Type   Interleave Capacity Partitions Members Status
;-----
r5                   RAID5  64      KB   447.0GB           1         4 DEGRADED

Ready.
```



### Note

If `rgdisplay` shows a RAID Group is OFFLINE, then data recovery methods should be initiated; skip this section and go directly to the Section, [Recovering from Failed Drives](#) on page 39.

If `rgdisplay` shows a RAID Group is DEGRADED, then it is important to replace the failed drive as soon as possible. `Rmstatus` shows the status of each of the member drives. If a member is FAULTED or UNAVAILABLE, it should be replaced. To match the displayed member to the physical drive, you can use `sesidentify`, which is the preferred method, or `blockdevidentify`, if there is no SES capability.

## Rmstatus

Next, use `rmstatus`, the RAID Member status CLI command, to show information for each member of the RAID Group in question.

```
rmstatus r5

6
;ID Status      RebuildStatus   Vendor      Product              Rev  SN
;-----
0 ONLINE      OK
1 ONLINE      OK
2 FAULTED     OK
3 ONLINE      OK
   ST3250310NS    SDN1 9SF0800F
   ST3160812AS-1 H     5LS25P51
   ST3250310NS    SDN1 9SF076NJ
   ST3250620AS    K     6QF0LJ17

Ready.
```

## SESIidentify DRIVE (only valid if enclosure supports SES)

The drive FAULT LED should already be blinking. Use this CLI command to also blink the Identify LED for the drive. The ID from the `blockdevscan` command must be used, and not the member ID shown in the `rmstatus` command. Match the S/N of the FAULTED drive (9SF076NJ) from `rmstatus` to the `blockdevscan` drive IDs, and use that number to select the correct drive to blink.

```
SesIdentify DRIVE 1
```

```
Ready.
```

```
SesIdentifyStop All
```

```
Ready.
```

## BlockDevIdentify

Use this command to blink the activity LED for the drive (if the drive is working enough so the activity LED can operate). The ID from the blockdevscan command must be used, and not the member ID shown in the rmstatus command. Match the S/N of the FAULTED drive (9SF076NJ) from rmstatus to the blockdevscan drive IDs, and use that number to select the correct drive to blink.

Sometimes, when `blockdevidentify` must be used, the activity LED won't blink because the drive has failed so badly. One possible action is to blink all the other LEDs, and find the failed drive by process of elimination.

```
BlockDevIdentify 1
```

```
Ready.
```

```
BlockDevIdStop
```

```
Ready.
```



### Note

If `sesidentify` is used, follow it up with the CLI command `sesidentifystop All` to turn the drive LEDs off. When `blockdevidentify` is used, follow it up with the CLI command `blockdevidstop` (no parameters) to turn the drive LEDs off.

## Using CLI to Enable / Disable Recovery Mode

Display the RAID Group and RAID Group member drives:

```
rgdisplay
```

```
3
;GroupName      Type      Interleave Capacity Partitions Members Status
;-----
r5              RAID5    64         KB  447.0GB          1         4 OFFLINE
```

```
Ready.
```

```
rmstatus r5
```

```
6
;ID Status      RebuildStatus      Vendor      Product              Rev  SN
;-----
0 ONLINE      OK                  ST3250310NS  SDN1 9SF0800F
1 FAULTED    OK                  ST3160812AS-1 H      5LS25P51
2 FAULTED    OK                  ST3250310NS  SDN1 9SF076NJ
3 ONLINE      OK                  ST3250620AS  K      6QF0LJ17
```

```
Ready.
```

Put the RAID Group into Basic or Extreme recovery mode with read and write access:

```
Set RGRrecover r5 Basic (or Extreme or Rebuild)
```

```
Ready.
```

Display the RAID Group to see the results:

```
rgdisplay

3
;GroupName      Type   Interleave Capacity Partitions Members Status
;-----
r5              RAID5  64      KB    447.0GB      1      4 RECOVER

Ready.
```

Display the RAID Group members to see the results:

```
rmstatus r5

6
;ID Status      RebuildStatus Vendor      Product              Rev  SN
;-----
0 ONLINE      OK              ST3250310NS SDN1 9SF0800F
1 ONLINE      OK              ST3160812AS-1 H 5LS25P51
2 ONLINE      OK              ST3250310NS SDN1 9SF076NJ
3 ONLINE      OK              ST3250620AS K 6QF0LJ17

Ready.
```

Once as much data as possible is retrieved from the drives, the Recovery Mode should be disabled and the FAULTED drives removed from the system.

Turn off Recovery Mode:

```
Set RGRRecover r5 Disable

Ready.
```

## Using CLI to Enable Recovery with Writes

On occasion, the data that has been disrupted is part of the File System on the disk drives. When this is the case, it is useful to run the File System Repair Tools for your operating system. This requires write access to the drives.



### CAUTION

**Use the `rgrecoverwithwrites` command only when necessary and use with caution!**

Put the RAID Group into *Basic* or *Extreme* recovery mode with read and write access:

```
Set RGRRecoverWithWrites r5 Basic (or Extreme)

Ready.
```

Run the appropriate file system repair utility prior to recovering data from the drives.

## 4 SNMP Configuration

SNMP is a standard network protocol that provides the ability to monitor SNMP enabled systems from anywhere on the network. Once configured, 3rd party MIB browsers on a Management Station can be used to retrieve the current configuration and receive TRAP messages when significant events occur. Via the Configuration.

The ATTO SNMP agent has two modes, depending on your configuration. If the only SNMP agent on your server is the ATTO agent, the mode should be **Enabled**. If you have other SNMP agents running and you want them to integrate their data into a single tree of information or if you want to use advanced SNMP features, choose **SubAgent** mode (See [SubAgent Mode](#) on page 49).

### Definitions

- **SNMP** - Simple Network Management Protocol
- **Network Node** - An addressable device attached to a computer network.
- **Management Station** - The host system that monitors network nodes.
- **Trap** - An alert that is sent to a management station by agents.
- **Agent** - A software process on the system being monitored that responds to queries using SNMP to provide status and statistics about a network node.
- **SubAgent** - A software process on the system being monitored that responds to SNMP queries from another agent, instead of a MIB browser. The request from the other agent is generated when it is contacted by a MIB browser.


### Details

The ATTO SNMP agent supports protocol versions 1 and 2c. Communication with the agent requires UDP over IPv4. The ATTO specific information can be found under "attotech" in the "enterprises" node (1.3.6.1.4.1.4547). Read-only support is provided.

#### Enabled Mode

This mode of operation provides Management Stations with basic information for the standard OID tree, as well as the ATTO specific information. While this mode will also work if a system service is installed and running, the ATTO agent must be configured to

listen on a different port than the system service to ensure both function properly. This mode is required to support SNMP if the system SNMP service (such as, Windows SNMP service or snmpd for OS X and Linux) is not already running.



## SubAgent Mode

This mode of operation relies on the operating system's SNMP service being installed as a Master Agent and running. All Management Station requests must be directed to this Master Agent, **not** the ATTO agent. It is expected that only experienced SNMP administrators will use this mode, as installing and configuring the operating system services are beyond the scope of this document. Examples are Windows SNMP services, or the open-source Net-SNMP package.

In this mode, the system service is automatically reconfigured so that when the service receives requests for the ATTO branch of the OID tree, the request is delegated to the ATTO agent through the loopback interface on a different port. While the primary benefit of subagent mode is to allow the Master Agent to handle most requests directly (the operating system services implement many more MIBs than the ATTO agent), and to allow multiple subagents to be merged into a single OID tree, this mode can be used to enable features that are not supported by the ATTO agent.

For example, if the system service supports version 3 of the protocol, all communication over the network between the Management Station and the service can be encrypted using the desired privacy options of version 3. The system service will then communicate to the ATTO agent using one of the supported protocols.




### Note

*For security purposes, the ATTO agent will only respond to requests through the loopback interface when in subagent mode. If the system service is stopped, the ATTO agent will be inaccessible from any other host on the network.*



### Note

*Other features that are not directly supported by the ATTO agent, but can be enabled with this configuration include source address filtering and stronger authentication than simple community strings provide. All of these features depend on properly configuring the system service, and are beyond the control of the ATTO agent. Refer to the documentation for the system service to determine how to enable these features.*



## Disabled Mode

This mode disables the operation of the ATTO SNMP agent. Any network ports it has open to support SNMP are closed, making it inaccessible to any Management Station. The system service will not forward requests to it, and no TRAPs are sent out when events occur.

## Basic Setup

The ATTO SNMP agent can be configured through the SNMP tab on the localhost node. When you first enter

the SNMP tab, and until you commit a new mode, the SNMP mode will be **Disabled**. The various settings can be manually edited, but the simplest way to enable SNMP support in this case is to click the **Default** button and then click **Commit**. This will load the proper settings based on your current system configuration, and then reconfigure the agent to use them. The operational mode (enabled or subagent) and port will be selected based upon whether or not an operating system SNMP service is running.



Note

*The **System Service** and **TRAP Service** states are not updated dynamically as they change in the system. They can be manually updated at any time by clicking **Refresh**.*

This basic setup will allow you to browse the information that is made available by the agent with a 3rd party application, but will not send out TRAPs. TRAP destinations are network specific, and cannot be determined programmatically.



Note

*In order to properly secure access to the agent from the local machine, it is suggested to change the default string in the **Communities** list before committing the default settings.*

## Configuration Options

---

### Agent Port

This value specifies the UDP port that the ATTO SNMP agent listens on for incoming Management Station requests. The port cannot be used by any other process on the system, or the behavior of both the agent and the other process is undefined. The standard port value for SNMP is 161, but that is not the default value in subagent mode, because it is assumed that the system service is using that port.

enabled, one will be sent to each configured destination. The list can be manipulated through the **Add**, **Edit** and **Remove** buttons below it. An existing community must be selected to edit or remove it. A valid community string has a length between 1 and 128 (inclusive) and can include any keyboard character. (See Exhibit 1-12)



Note

*A binding error may not occur if the port is already in use. If you are not sure if a port is in use, use the `netstat` command.*



Note

*In subagent mode, only a single community can be specified since that is the community used by the Master Agent when talking to the ATTO SNMP subagent. The communities that Management Stations must use are configured through the system service.*

In enabled mode, this is the port the Management Station uses to communicate with the agent. In subagent mode however, the agent will **not** respond on this port to Management Station requests from a different machine. The port the Management Station needs to use is the system service's port.



Note

*For added security in subagent mode, you should change the default community string. The agent cannot tell the difference between a local Management Station request and a request from the system service, which means a local user can bypass the authentication checks done by the system service if they know this community string and the agent's port.*

### Communities

This is a list of community strings accepted by the agent when it receives an incoming request. If a Management Station makes a request and provides a community string that is not in this list, the request is dropped by the agent. If authentication traps are

## Send Authentication TRAP

When checked, the agent will send a TRAP to the configured TRAP destinations indicating that a Management Station attempted to contact the agent and used a community string that is not in the community list.



### Note

*In subagent mode, this may only be useful for debugging. The system service will perform Management Station authentication based on its configuration, and only forward the request to the ATTO agent once the Management Station is authenticated. Since the system service is reconfigured as needed to use the community string on this panel when forwarding requests, there should not be authentication errors when the agent processes requests from the system service. The system service usually has the same option that can be enabled to see these authentication failures.*

## Enable TRAPs

When checked, the agent will send traps to the configured destinations, if any, if the agent itself is not disabled. (See [Exhibit 1-11](#))

## Trap Destination Table

This table lists the host address and community string used for each destination when a trap is triggered. The host column consists of the hostname or IP address of the Management Station to which the traps are sent,

and the UDP port on which the trap receiver is listening. The community column displays the community string sent with the trap to that host. If the community string is not in the list of communities accepted by the receiver, the receiver will ignore the trap message. (See [Exhibit 1-13](#))

The table can be manipulated through the Add, Edit and Remove buttons below it. An existing destination must be selected to edit or remove it. A valid community string has a length between 1 and 128.



### Note

*The trap configuration is completely independent of the enabled or subagent modes, but no traps are sent if the agent is disabled.*

Many operating systems have a trap receiver service as part of their SNMP package, which can be configured to do various things when it receives a trap, such as writing to a log file or forwarding to another destination. In a similar manner to the way the system service and subagent mode can enable SNMPv3 support, the operating system trap service can be configured to translate traps it receives into SNMPv3 traps, with optional privacy, authentication and INFORM support. The ATTO agent then only needs to specify a single trap destination to this service in order to enable SNMPv3 support over the network. The manner in which the system service is configured is system specific, and beyond the scope of this document.

## Control Buttons

### Commit

This button will reconfigure the ATTO agent to use the settings as they are currently displayed on the screen, after validating them to the extent possible. The settings are persistently stored, so that they are available after a reboot. If necessary, the system service will be reconfigured to support the agent's current operational mode.

### Restore

This button will update the display to reflect the settings the agent is currently using, rolling back any changes made in the panel that have not been committed. Since these are the current settings, they do not need to be committed.

### Default

This button will update the display with the simplest configuration, based on the current system configuration. For example, if the operating system service is detected as running, the default mode is subagent. Otherwise, the mode is enabled and the registered SNMP port is displayed.



### WARNING

**Since it isn't possible to determine trap destinations and community strings programmatically, all trap destinations will be cleared.**

## Test

This button will validate the settings displayed on screen, and then send a warm Start trap to each of the displayed trap destinations as if the displayed settings had been committed. Selecting disabled mode, unchecking the enable traps option or removing all trap destinations will prevent the test trap from being sent. Once an acceptable configuration is entered, a commit is needed for it to take effect.



### Note

*The test is not performed by changing the settings used by the agent at the time of the test, so testing will have no effect on any concurrent Management Station requests.*

## Save MIBs

This button will save the non standard MIBs implemented by the agent to the user specified directory as individual text files, therefore they can be loaded into a MIB browser. If one of the files already exists in that directory, the user is prompted to overwrite or skip the file, or cancel the operation.

## Troubleshooting

---

### Requests from a MIB browser time out.

- Ensure the agent is not disabled
- Ensure the Management Station is using the port for the system service if the agent is in subagent mode or the port displayed on the SNMP panel if the agent is enabled.
- Ensure the community string used by the Management Station is accepted by the system service when in subagent mode or is displayed in the communities list on the SNMP panel if the agent is enabled. Be sure the Send Authentication TRAP option is selected, committed, and then try again. If the community is a problem, each of the configured trap destinations will be notified.
- Ensure the Management Station is using SNMPv1 or SNMPv2c, unless the system service is being used.
- If in subagent mode, the system service may be misconfigured. Try enabling the agent and using the Management Station to talk directly to it. You will have to change the port value and likely the community string.

### Traps are not received when testing

- Ensure the agent is not disabled
- Ensure the Enable TRAPs option is checked.
- Ensure a trap recipient is properly configured and running on each machine in the destination table, at the specified port.
- Ensure the community associated with the destination is in the trap recipient's community list.

### The RAID OIDs are skipped when walking or the tables are empty

- Ensure the driver for the RAID adapter is installed. This can be seen on the Basic Info panel for the adapter.



### Note

*The architecture of Net-SNMP installed has to match the O/S; i.e. if you are using a 64bit O/S the 64bit version of Net-SNMP must be installed.*

Exhibit 4-1 Panel Image




Exhibit 4-2 Community Image





Exhibit 4-3 Edit Trap Image



## 5 Troubleshoot the ATTO ConfigTool

You may see an error message informing you about an unexpected event or incorrect information discovered by the application. Using the help text presented with the error message, correct the issue before proceeding.

Warnings and error messages are posted in the **Status** area of the configuration utility.

### Messages from NVRAM tab actions

#### An error occurred loading NVRAM data.

The first time a channel is highlighted, the application tries to read NVRAM from the card. This message usually indicates the application could not communicate with the driver, probably because the application does not support the driver version in use.

#### Warning: NVRAM could not be read, defaults returned.

NVRAM is corrupt and the driver returns to the default configuration. The defaults are presented via the graphical user interface. These defaults may be modified but the defaults or modifications must be committed in order to correct NVRAM.

#### An error occurred updating the NVRAM.

The driver cannot put the new settings on the card; no changes are made to the card.

### Feature bounds checking

When the **Commit** button is clicked, each NVRAM feature is validated before being sent to the card. If any one of these features is deemed inappropriate based on the implemented checks, further NVRAM feature validation checks are stopped and the message is displayed.

#### Execution Throttle is greater than the maximum allowable value of 255. No NVRAM configuration changes have been made to your card.

The exact message varies based on the first field with an out-of-range value.

### Messages from Flash tab actions

#### This is not a flash file, or it is corrupt.

The ATTO-created flash file is corrupt or the **ConfigTool** does not recognize the file as a flash file. Only ATTO-created flash files may be selected using the flash file dialog box.

#### This HBA is not compatible with the selected flash file.

ATTO flash files are created based on the type of card flashed. ATTO flash files are only compatible with certain ATTO cards. When a flash file is selected, the flash file is inspected to determine if it is compatible.

#### A valid file was not selected.

You clicked the **Cancel** button on the flash file selection dialog.

#### An error occurred reading from the flash file, the file may be corrupt.

You selected a compatible flash file but the contents are corrupt.

#### An error occurred updating the flash.

You tried to flash a card when the firmware was not able to accept a flash.

#### The card has been prepared for firmware updating, but the machine must be rebooted for the changes to take effect. You need to repeat this process after rebooting to actually update the firmware.

Some firmware upgrades need to prepare the existing firmware in order to successfully update the adapter. Rebooting allows the changes made during the preparation process to take effect, and the same file should be flashed again.

# Appendix A CLI provides an ASCII-based interface

The **RAID CLI** tab in the ATTO ConfigTool allows experienced users to enter RAID Command Line Interface (CLI) commands to the ThunderStream SC. The CLI uses ASCII commands typed while in CLI mode.



## CAUTION

Do not use the CLI unless you are directed to by an ATTO technician.

Changing parameters may cause loss of data and/or disruption to performance and reliability of the ThunderStream SC storage controller.

The ATTO ConfigTool interface is the preferred method to operate and manage the ThunderStream SC storage controller. Refer to [RAID Configuration](#) on page 10 for details.

The command line interface (CLI) is a set of ASCII-based commands which perform configuration and diagnostic tasks. Refer to [RAID Configuration](#) on page 10.

- CLI commands are context sensitive and generally follow a standard format

```
[Get|Set] Command [Parameter1|Parameter2]
```

followed by the **return** or **enter** key.

- CLI commands are case insensitive: you may type all upper or all lower case or a mixture. Upper and lower case in this manual and the **help** page are for clarification only.

- Commands generally have three types of operation: get, set and immediate.
- The get form returns the value of a parameter or setting and is an informational command.
- Responses to get commands are followed by **Ready**.
- The set form is an action which changes the value of a parameter or configuration setting. It may require a **SaveConfiguration** command and a restart of the system before it is implemented. The restart can be accomplished using a separate **FirmwareRestart** command. A number of set commands may be issued before the **SaveConfiguration** command.
- Responses to **set** commands are either an error message or **Ready**. \*. The asterisk indicates you must use a **SaveConfiguration** command to finalize the **set** command.
- Set commands which do not require a **SaveConfiguration** command, defined as immediate commands, are immediately executed.



## Note

*Using certain CLI commands during normal operation can cause a performance drop. Once command actions are complete, performance should return to normal levels.*

*Exhibit A-1 Symbols, typefaces and abbreviations used to indicate functions and elements of the command line interface used in this manual.*

Symbol	Indicates
[ ]	Required entry
< >	Optional entry
	Pick one of
n-n	A range (6 – 9 = 6, 7, 8, 9)
BlockDevID	Index designation of a block device not assigned to any other RAID group; the index of a block device provided by the <a href="#">BlockDevScan (Immediate, Disabled on Error)</a> command. 0<=n<=63
GroupName	The name of the RAID group to which the block device is assigned, or blank if the block device is available
MemberIndex	Index designation of a RAID group member as found in the <a href="#">RMStatus (Immediate)</a> command
PartIndex	Index designation of a partition as found in the <a href="#">PartitionDisplay (Immediate)</a> command
SASIndex	Index designation of SAS drives as found in <a href="#">SASTargets (Immediate, Disabled on Error)</a>
TID	Target ID: 0<=n<=255

## CLI error messages

---

The following error messages may be returned by the Command line Interface

```

ERROR Invalid Command. Type 'Help' for command list.
ERROR Command Not Processed
ERROR Wrong/Missing Parameters
ERROR Invalid Hot Spare Serial Number
ERROR Invalid RAID GroupName
ERROR Invalid RAID Group State
ERROR Insufficient number of RAID Group members
ERROR RAID Group does not exist
ERROR No RAID Groups found
ERROR Invalid RAID Type
ERROR RAID Group is already unmapped
ERROR Invalid Block Device Index
ERROR Invalid RAID MemberIndex
ERROR Invalid RAID Member State
ERROR Missing RAID Member
ERROR Invalid RAID Member Capacity
ERROR Invalid Partition Index
ERROR Maximum number of RAID Groups exceeded
ERROR Maximum number of Partitions exceeded
ERROR Invalid number of Partitions
ERROR Maximum number of RAID Members exceeded
ERROR Maximum stripe width

```

ERROR Invalid number of Partitions specified  
 ERROR Invalid Span Depth specified  
 ERROR Cannot perform operation on mapped Partition  
 ERROR Cannot perform operation. RAID Group has mapped Partitions  
 ERROR Cannot perform operation. RAID Group has Outstanding Commands  
 ERROR Block Device at specified index no longer available  
 ERROR Insufficient RAID Group members for RAID type  
 ERROR Incorrect number of RAID Group members for QuickVideo configuration  
 ERROR Invalid Virtual Drive ID  
 ERROR Specified capacity is invalid  
 ERROR Too many Indices specified.  
 ERROR Only one add storage operation is permitted at any given time.  
 ERROR No free block devices  
 ERROR Cannot benchmark a drive that is being initialized  
 ERROR Specified drive is not being monitored

## CLI summary

The following chart summarizes the Command Line Interface commands, their defaults, and an example of how to enter the commands. Commands which have no default values have a blank entry in that column of the table.



### CAUTION

**Do not use CLI unless you are directed to by an ATTO technician.**  
**Changing parameters may cause loss of data and/or disruption to performance and reliability of the ThunderStream SC.**

Command	Default	Example
AutoMap		automap
AutoResume	rebuild = enabled erase = disabled initialize = enabled	set autoresume erase enabled raidgroup1
BlockDevClean		blockdevclean 30
BlockDevIdentify		blockdevidentify 30
BlockDevIDStop		blockdevidstop
BlockDevInfo		blockdevinfo
BlockDevScan		blockdevscan
ClearEventLog		cleareventlog
Date		
DeleteScheduledTasks		
DisplayScheduledTasks		
DriveAssureReport		driveassurereport groupname1
DriveHealth	disabled	set drivehealth enabled
DriveHealthDisplay		drivehealthdisplay all
DriveHealthStatus		drivehealthstatus
DriveTest		drivetest begin
DriveTestClearList		drivetestclearlist all
DriveTestConfig	not initiated	set drivetestconfig read
DriveTestList		get drivetestlist all

Command	Default	Example
DriveTestStatus		get driveteststatus
DumpConfiguration		dumpconfiguration
DumpEventLog		dumpeventlog
FUAWrites	enabled	set fuawrites disabled
Help		help eventlog
HSAdd		hsadd 3
HSDisplay		hsdisplay
HSRemove		hsremove 3
IdentifyBridge	Disabled	set IdentifyBridge enabled
Info		info
IsReserved		isreserved
MediaScanErrorReport		
Metrics		metrics display all
MirrorIdentify		set mirroridentify rg1
MirrorIdentifyStop		mirroridentifystop rg1
OEMConfigFile	ATTO	get oemconfigfile
Partition		partition alpha1 6 4 GB
PartitionDisplay		partitiondisplay alpha1
PartitionMerge		partitionmerge all
PartitionSplit		partitionsplit alpha1 22 2
PartitionWriteCache		set partitionwritecache enabled
PassThroughMode		set passthroughmode all
PassThroughPersistent		passthroughpersistent
PassThroughRediscover		passthroughrediscover
RAIDRebuildPriority	same	set raidrebuildpriority low
RAIDSpeedWriteLimit	8	set raidspeedwritelimit 16
RebuildContinueOnError	Disabled	set rebuildcontinueonerror all enabled
Reserve		reserve
RestoreConfiguration		restoreconfiguration default
RGAddStorage		rgaddstorage groupname1 span commit
RGAutoRebuild	disabled	set rgautorebuild all enabled
RGCancelAddStorage		rgcanceladdstorage groupname1
RGCancelMediaScan		
RGCommit		rgcommit all
RGCreate		rgcreate groupname1 raid0
RGDiskWriteCache	enabled	set rgdiskwritecache all disabled
RGDisplay		rgdisplay all
RGErase		rgerase groupname1
RGHaltConversion		rghaltconversion groupname1
RGHaltErase		rghalterase groupname1
RGHaltInitialization		rghaltinitialization groupname1
RGHaltMediaScan		
RGHaltRebuild		rghaltrebuild groupname1
RGMediaScan		
RGMediaScanErrorReport		set rghdparameter groupname1 30

Command	Default	Example
RGMemberAdd		rgmemberadd groupname1 30
RGMemberRemove		rgmemberremove groupname1 30
RGPrefetch		
RGReadDirect	enabled	set rgreaddirect groupname 1 disabled
RGRebuild		rgrebuild groupname1
RGRecover		rgrecover g1 rebuild
RGRecoverWithWrites		rgrecoverwithwrites g1 basic
RGResumeConversion		rgresumeconversion groupname1
RGResumeErase		rgresumeerase groupname1
RGResumeInitialization		regresumeinitialization groupname1
RGResumeMediaScan		
RGResumeRebuild		rgresumerebuild groupname1
RGSectorSize	512	set rgsectorsize groupname1 4096
RGSpanDepth	1	set rgspandepth groupname1 8
RGSpeedRead	all disabled	set rgspeedread groupname1 enabled
RGUnmap		rgunmap groupname1
RGWaitTimeout	3	rgwaittimeout 30
RMStatus		rmstatus groupname1
Route		route host 1 raid groupname1 6
RouteDisplay		routedisplay host 0
SasPortList		
SASTargets		sastargets
SaveConfiguration		saveconfiguration
SerialNumber		get serialnumber
SES	enabled	
SESArmTest		
SESDiskFailureAlarm	disabled	
SESEnclosures		
SESIIdentify	off	set sesidentify all
SESIIdentifyStop		sesidentifystop all
SESMute		
SESPoll	30	set sespoll 0
SESStartingSlot	1	
SESStatus		
Temperature		get temperature
Time		set time 03:32:30
TimeZone	EST	set timezone pst
VerboseMode	enabled	set verbosemode disabled
VirtualDriveInfo		virtualdriveinfo

# CLI command explanations

---

Command Line Interface commands are listed alphabetically with explanations of what they are used for, their defaults and syntax.



## CAUTION

**Using CLI without contacting an ATTO technician is not recommended because changing parameters may cause loss of data and/or disruption to performance and reliability of the ThunderStream SC.**

### AutoMap (Immediate, Disabled on Error)

Automap automatically maps each RAID Partition to a Target ID on the host system (maximum 256 maps allowed). If the optional 'passthrough' parameter is entered then AutoMap stores maps for currently connected SAS/SATA Pass Through Devices. Existing maps for available targets are preserved.

Usage: AutoMap <passthrough>

### AutoResume

AutoResume sets or gets the AutoResume features for interrupted rebuild, write pattern, erase, initialization, media scan, and conversion operations at startup. Optional parameter GroupName specifies the RAID group to operate on. If no GroupName is specified, the command operates on all existing RAID groups.

Usage: set AutoResume [Rebuild|WritePattern  
|Erase|Initialization|MediaScan|Conversion|all]  
[enabled|disabled] <GroupName>  
get AutoResume  
[Rebuild|WritePattern|Erase|Initialization|MediaScan|  
Conversion|all]

### BlockDevClean (Immediate, Disabled on Error)

BlockDevClean removes any RAID configuration data from the block device with the specified BlockDevID. BlockDevID is the index of a block device provided by the BlockDevScan CLI command. Caution: All RAID Group setup information is lost and you lose all RAID Group data.

Usage: BlockDevClean [BlockDevID]

### BlockDevIdentify (Immediate, Disabled on Error)

Lights the LED of a disk drive. Use either RAID Group name and member index, or BlockDevID. BlockDevID is the index of a block device provided by the BlockDevScan CLI command.

Usage: BlockDevIdentify [[Groupname MemberIndex] |  
BlockDevID]

### BlockDevIdStop (Immediate, Disabled on Error)

Turns off the IO LED of a previously identified disk drive. Specify RAID group name and member index, BlockDevID, or no parameters (which will turn off all previously identified drives). BlockDevID is the index of the block device provided by the BlockDevScan CLI command.

Usage: BlockDevIdStop <[GroupName MemberIndex] |  
BlockDevID>

### BlockDevInfo (Immediate, Disabled on Error)

BlockDevInfo gives detailed information about connected physical block devices along with any potential RAID Group association.

Usage: BlockDevInfo [BlockDevID | all]

### BlockDevScan (Immediate, Disabled on Error)

BlockDevScan lists all currently connected physical block devices along with any potential RAID Group association. Each block device listed is assigned a unique index at the time of the scan. This index is used to identify block devices for other CLI operations.

Usage: BlockDevScan

### ClearEventLog (Immediate)

ClearEventLog clears the contents of the event log.

Usage: ClearEventLog

### Date

Sets/displays the current date. The date range is 01/01/2000 to 12/31/2099.

Usage: set Date [MM/DD/YYYY]  
get Date

### DeleteScheduledTasks (Immediate, Disabled on Error)

Deletes a scheduled task with the Id returned by DisplayScheduledTasks.

Usage: DeleteScheduledTasks [Id]

### **DisplayScheduledTasks (Immediate, Disabled on Error)**

Immediate command that displays all outstanding scheduled tasks.

Usage: DisplayScheduledTasks

### **DriveAssureReport (Immediate, Disabled on Error)**

Displays drive timeout error statistics for all member drives in the specified RAID group.

Usage: DriveAssureReport [GroupName]

### **DriveHealth (Disabled on Error)**

Changes the system's ability to acquire drive health data from connected drives. Issuing this command during I/O operations may adversely affect performance.

Usage: set DriveHealth [enabled | disabled]  
get DriveHealth

### **DriveHealthDisplay (Immediate, Disabled on Error)**

Retrieves and displays S.M.A.R.T. data from SATA disk drives, and MEDIUM DEFECT and INFORMATION EXCEPTIONS counts from other drives. Issuing this command during I/O operations may adversely affect performance. S.M.A.R.T data may be optionally filtered by attribute type: Performance, Error Rate, Event Count, Critical, or All. If filtering is used, only data for SATA drives is shown.

Usage: DriveHealthDisplay [BlockDevID | all] <perf | error | evtcnt | crit | all>

### **DriveHealthStatus (Immediate, Disabled on Error)**

Displays the current S.M.A.R.T. support of specified SATA disk drives and MEDIUM DEFECT and INFORMATION EXCEPTION support in other disk drives.

Usage: DriveHealthStatus [BlockDevID | all]

### **DriveTest (Immediate, Disabled on Error)**

Immediate command which starts or stops a drive test with the previously specified configuration and drive list. Drives which are in-use by the test are not available for RAID configuration or RAID operations. Only one test can be run at a time.

Usage: DriveTest [Begin | Cancel]

### **DriveTestClearList (Immediate, Disabled on Error)**

Specifies drives to be removed from the drive test list. The 'drive BlockDevID' parameter will remove the specified drive from the list. The 'all' parameter automatically removes all drives from the list.

Usage: DriveTestClearList [drive [BlockDevID] | all]

### **DriveTestConfig (Disabled on Error)**

Configures the drive test to perform one of the following operations: initialize (destructive write-only), mediascan (destructive for sectors with medium errors), read (non-destructive read-only), verify (destructive verify), or init-verify (destructive write-read-verify). The test is not started until the DriveTest Begin command is given. A new configuration may not be set while a drive test is being performed.

Usage: set DriveTestConfig [init | read | verify | mediascan | init-verify]  
get DriveTestConfig

### **DriveTestList (Disabled on Error)**

Specifies drives to be run in the next drive test. DriveTestConfig should be setup prior to adding any drives into the test list. This command can be called with different eligible BlockDev IDs and each one will be added to the list. Drives which are part of a RAID Group are only eligible for read drive tests. Additionally, Hot Spare drives are only eligible for mediascan and read drive tests. The 'all' parameter automatically chooses eligible drives. The test is not started until the DriveTest Begin command is given.

Usage: set DriveTestList [drive [BlockDevID] | all]  
get DriveTestList

### **DriveTestStatus**

Displays the status of the currently running drive test. This command does not display performance metrics. If a block device ID is not running or cannot be found, its state will be 'idle' and percent complete will be 0.

Usage: get DriveTestStatus <drive [BlockDevID]>

### **DumpConfiguration (Immediate)**

Dumps system's configuration

Usage: DumpConfiguration

### **DumpEventLog (Immediate)**

DumpEventLog can be used to dump the entire contents of the event log.

Usage: DumpEventLog

### **FUAWrites**

FUAWrites is used to enable or disable whether or not FUA writes will be able to be sent to disk drives.

Usage: set FUAWrites [enabled | disabled]  
get FUAWrites

### **Help (Immediate)**

The Help command issued with no parameters displays a list of available CLI commands. When a CLI Command name is specified, a command usage string and command description is presented on the CLI.

Usage: Help <command>

### **HSAAdd (Immediate)**

Assigns a Block Device to the Hot Spare pool.

Usage: HSAAdd [BlockDevID]

### **HSDisplay (Immediate)**

HSDisplay outputs a list of all of the devices in the Hot Spare pool.

Usage: HSDisplay

### **HSRemove (Immediate)**

Removes a Block Device from the Hot Spare pool.

Usage: HSRemove [BlockDevID | all]

### **IdentifyBridge**

IdentifyBridge causes the 'Alert' LED to blink to enable identification of this system. Disable this option to cancel the blinking.

Usage: set IdentifyBridge [enabled | disabled]

### **Info (Immediate)**

Info displays version numbers and other product information for key components. Use the optional 'brief' parameter to show a more concise subset of system information.

Usage: Info <brief>

### **IsReserved (Immediate)**

IsReserved displays the reservation status of the current services session/interface.

Usage: IsReserved

### **MediaScanErrorReport (Immediate, Disabled on Error)**

Displays media scan error statistics for either a single block device or all block devices in the system.

Usage: MediaScanErrorReport [BlockDevID | all]

### **Metrics (Immediate)**

The Metrics CLI command offers control over the collection of standard data metrics within a product via the command's 'Start', 'Stop', and 'Display' parameters.

Usage: Metrics [Start|Stop|Display|Clear] [[drive [BlockDevID]]|all|running]

### **MirrorIdentify(Immediate)**

For RAID 1 and RAID 10, MirrorIdentify blinks the Identify LEDs for a set of RAID Group Drives that contain a complete copy of the data. Use MirrorNumber 0 to identify the original set.

Usage: set MirrorIdentify [GroupName <MirrorNumber>]  
get MirrorIdentify [GroupName <MirrorNumber>]

### **MirrorIdentifyStop(Immediate)**

MirrorIdentifyStop stops identifying the specified set of RAID Group drives. GroupName with no mirror number stops identifying all drives in the RAID Group.

Usage: MirrorIdentifyStop [GroupName <MirrorNumber>]

### **OEMConfigFile (Disabled on Error)**

This command returns the "name" (i.e., the contents of the first record) of the OEM configuration file stored in persistent memory.

Usage: get OEMConfigFile

### **Partition (Immediate)**

Partition sets the specified partition to the specified capacity in gigabytes (GB), megabytes (MB), or blocks. The specified capacity must be smaller than the specified partition's current capacity. A new partition is created to acquire the remainder of the original partition's space.

Usage: Partition [GroupName] [PartIdx] [capacity] [GB | MB | blocks]

### **PartitionDisplay (Immediate)**

PartitionDisplay outputs a list of all of the partitions available in the specified RAID Group. The partitions are listed in order of contiguousness (as opposed to index order). GroupName is the ASCII name of the RAID Group for which partitions will be displayed.

Usage: PartitionDisplay <GroupName>

### **PartitionMerge (Immediate)**

PartitionMerge merges the specified contiguous partitions into one partition. GroupName is the ASCII name of the RAID Group containing the partitions to merge. PartIdx is the index of a partition to merge, along with a number of contiguous partitions to merge to that index. 'All' indicates that all partitions in the RAID Group will be merged into a single Virtual Disk. The RAID Group must not be in a NEW state. None of the partitions to merge may be mapped.

Usage: PartitionMerge [GroupName] [[[PartIdx] [2-128]] | all]

### **PartitionSplit (Immediate)**

PartitionSplit splits the specified partition into one or more partitions whose capacities are evenly distributed among the capacity of the original partition. GroupName is the ASCII name of the RAID Group containing the partition to split. PartIdx is the index of the partition to split. The partition to split cannot be mapped. The RAID Group must not be in a NEW state.

Usage: PartitionSplit [GroupName] [PartIdx] [2-128]

### **PartitionWriteCache**

Enable RAID internal Write Cache for higher write performance with a small risk of data loss after a system failure. Disable RAID internal Write Cache for a higher level of data integrity with lower write performance.

Usage: set PartitionWriteCache [GroupName] [PartIdx]  
[enabled | disabled]  
get PartitionWriteCache [GroupName] [PartIdx] | all]

### **PassThroughMode (Disabled on Error)**

PassThroughMode specifies the non-disk device types which will be automatically mapped at boot time. The "all" option allows all non-disk devices to be mapped. The "SES" option allows dedicated SES processor LUNs to be mapped. The "non-SES" option allows all non-SES devices to be mapped. The "disabled" option disables the pass-through mode.

Usage: set PassThroughMode [all | SES | non-SES | disabled]  
get PassThroughMode

### **PassThroughPersistent (Immediate, Disabled on Error)**

PassThroughPersistent stores maps for currently-attached passthrough devices to persistent memory.

Usage: PassThroughPersistent

### **PassThroughRediscover (Immediate, Disabled on Error)**

PassThroughRediscover will make any previously deleted pass through target devices visible to the host.

Usage: PassThroughRediscover

### **RAIDRebuildPriority**

Set the RAID rebuild priority. A RAID rebuild priority that is set to high, will give higher priority to RAID rebuilds and lower priority to the processing of simultaneous I/O transactions. A RAID rebuild priority that is set to low, will give lower priority to RAID rebuilds and higher priority to the processing of simultaneous I/O transactions. A RAID rebuild priority that is set to same, will give equal priority to RAID rebuilds and the processing of simultaneous I/O transactions. If all or no groups are specified, the system default and all of the individual RAID Groups are set. If a group name is specified, only the group specified is set.

Usage: set RAIDRebuildPriority <GroupName | all> [high | low | same]  
get RAIDRebuildPriority <GroupName | all>

### **RAIDSpeedWriteLimit (Disabled on Error)**

Set or get the limit on the coalescing factor. Warning: changing the default setting (8) may result in poor performance or timeouts. A lower setting is recommended when using multiple initiators. A higher setting may improve performance with multiple streams of sequential write I/O, but too high a setting will cause timeouts.

Usage: set RAIDSpeedWriteLimit [0 - 32]  
get RAIDSpeedWriteLimit

### **RebuildContinueOnError (Disabled on Error)**

Allows a rebuild to continue to completion even if media read errors are encountered during the rebuild operation.

Usage: set RebuildContinueOnError [GroupName | all]  
[enabled | disabled]  
get RebuildContinueOnError [GroupName | all]

### **Reserve (Immediate)**

Reserve prevents other CLI sessions from modifying the system. When the management interface is reserved, set commands are unavailable to other sessions, but get commands are available.

Usage: Reserve

### **RestoreConfiguration (Immediate, Disabled on Error)**

RestoreConfiguration issued with the 'default' option will force the NVRAM settings to their original defaults. The 'saved' option will undo any changes made to this session since the last save.

Usage: RestoreConfiguration [default | saved]

### **RGAddStorage (Immediate)**

RGAddStorage adds additional storage to an existing RAID Group. GroupName is an ASCII name for the RAID Group. MIRROR|STRIPE|SPAN specifies the method used to expand the storage. Optional parameter list BlockDeviceID specifies up to 10 indices of available block devices, provided by the BlockDevScan CLI command, to be added to the RAID Group. If this list is omitted, the CLI command RGMemberAdd must be used. Optional parameter commit runs the RGCommit command automatically and all user data will be erased from each new member drive. If the parameter is omitted, the CLI command RGCommit must be entered. Any time before RGCommit is entered, the command RGCcancelAddStorage can be used to cancel the process. NOTE: MIRRORS cannot be added to a RAID 5, RAID 4, or DVRAID Group.

Usage: RGAddStorage [GroupName]  
[MIRROR|STRIPE|SPAN] <BlockDeviceID  
<commit> >

### **RGAutoRebuild**

RGAutoRebuild enables and disables Auto-Rebuild functionality for one or more RAID Groups. Auto-Rebuild uses drives assigned as Hot Spares, followed by available drives, as automatic replacements for any member that fails. Auto-Rebuild is disabled by default.

Usage: set RGAutoRebuild [GroupName | all] [enabled | disabled]  
get RGAutoRebuild [GroupName | all]

### **RGCancelAddStorage (Immediate, Disabled on Error)**

RGCancelAddStorage cancels the RGAddStorage command.

Usage: RGCancelAddStorage [GroupName]

### **RGCancelMediaScan (Immediate, Disabled on Error)**

RGCancelMediaScan cancels a media scan that is running on the specified existing RAID Group.

Usage: RGCancelMediaScan [GroupName]

### **RGCommit (Immediate, Disabled on Error)**

RGCommit stamps a NEW RAID Group's configuration to its member drives. Advanced Initialization is highly recommended for new drives; this erases and verifies the drive media. The RAID Group is unavailable until the operation completes. Express Initialization performs a background initialization and the RAID Group is immediately available for use. When RGCommit is issued after adding storage, it stamps an EXISTING RAID Group's configuration to the added drives, and initializes the drives if Advanced or Express are specified. GroupName is the ASCII name of the NEW RAID Group to commit.

Usage: RGCommit < GroupName <Advanced | Express> | all <Advanced | Express> >

### **RGCreate (Immediate)**

RGCreate creates a NEW empty RAID Group. GroupName is an ASCII name for the RAID Group (14 chars max, no spaces). The optional value after the RAID Group type parameter represents the desired interleave for the RAID Group. KB denotes interleave in kilobytes; without the KB suffix, interleave is set in 512 byte blocks. If interleave is not provided, the system-default interleave is used.

Usage: RGCreate [GroupName] [RAID[0|1|10|4|5|6]|JBOD] <64KB|128KB|256KB|512KB|1024KB |128|256|512|1024|2048>

### **RGDiskWriteCache**

Enable RAID member disk Write Cache for higher write performance with a small risk of data loss after a system failure. Disable RAID member disk Write Cache to force the drives to update the storage media at the expense of some write performance.

Usage: set RGDiskWriteCache [GroupName | all] [enabled | disabled]  
get RGDiskWriteCache [GroupName | all]

### **RGDisplay (Immediate)**

RGDisplay displays RAID Group status information. GroupName is an ASCII name for an existing RAID Group.

Usage: RGDisplay <GroupName | all>

### **RGErase (Immediate)**

RGErase erases the data from the specified existing RAID Group. WARNING: All data will be lost!

Usage: RGErase [GroupName]

### **RGHaltConversion (Immediate)**

RGHaltConversion halts the conversion on the specified existing RAID Group.

Usage: RGHaltConversion [GroupName]

### **RGHaltErase (Immediate)**

RGHaltErase halts the erase on the specified existing RAID Group.

Usage: RGHaltErase [GroupName]

### **RGHaltInitialization (Immediate)**

RGHaltInitialization halts the initialization of the specified RAID Group.

Usage: RGHaltInitialization [GroupName]

### **RGHaltMediaScan (Immediate, Disabled on Error)**

RGHaltMediaScan halts a media scan on the specified existing RAID Group.

Usage: RGHaltMediaScan [GroupName]

### **RGHaltRebuild (Immediate)**

RGHaltRebuild halts the rebuild(s) on the specified existing RAID Group. Optional parameter MemberIndex specifies the RAID Member whose rebuild will be halted. For RAID6 Groups, if a MemberIndex is specified, all rebuilding RAID Members on the span with that MemberIndex will halt as well. If no MemberIndex is specified, all rebuilds on that RAID Group will be halted.

Usage: RGHaltRebuild [GroupName] <MemberIndex>

### **RGMediaScan (Immediate, Disabled on Error)**

RGMediaScan initiates a Media Scan. A Media Scan reads all member drives and corrects Media Errors by calculating the expected data and rewriting it, so the drive can relocate it to a good sector. The 'verify' option adds a data integrity check by verifying that the data and parity match. 'Verify' plus 'fix' causes the parity to be re-written when a verify mismatch occurs. Enter time (must be HH:MM in 24-hour time format) and day of week without the 'daily/weekly' option to schedule a one-time scan for a later date. Enter time, day of week and 'daily' or 'weekly' to schedule a scan on a recurring basis.

Usage: RGMediaScan [GroupName] <verify | verify fix> <[HH:MM] <day of week> <daily | weekly>>

### **RGMediaScanErrorReport (Immediate, Disabled on Error)**

Displays error statistics for either a single RAID Group or all RAID Groups in the system.

Usage: RGMediaScanErrorReport [GroupName | all]

### **RGMemberAdd (Immediate)**

RGMemberAdd adds available block devices to a NEW RAID Group or as part of an RGAddStorage operation. GroupName is the ASCII name of the RAID Group to receive the RAID Member. BlockDevID is the index of an available block device provided by the BlockDevScan CLI command. Up to 10 BlockDevIDs may be specified. If all is specified, then all available unused BlockDevIDs will be added to the RAID Group until the maximum number of RAID Group members has been met. This command also resets the number of RAID Group partitions to 1.

Usage: RGMemberAdd [GroupName] [BlockDevID] all.

### **RGMemberRemove (Immediate)**

RGMemberRemove removes a RAID Member from a NEW RAID Group. GroupName is the ASCII name of the NEW RAID Group from which to remove the RAID Member. MemberIndex is the index of the RAID Member to remove. This also resets the number of partitions to 1.

Usage: RGMemberRemove [GroupName] [MemberIndex]

### **RGPrefetch**

Set or Get the prefetch for all or for the specified RAID Group. This command will fail if the RAID Group does not exist. GroupName is the ASCII name of the RAID Group for which the parameter will apply.

Usage: set RGPrefetch [GroupName | all] [Value 0 to 6]  
get RGPrefetch [GroupName | all]

### **RGReadDirect**

Directly read data from SAS hardware into host memory. This feature eliminates store-and-forward delays that are inherent with reading data into RAID internal cache first and then transferring it to host memory.

Usage: set RGReadDirect [GroupName | all] [enabled | disabled]  
get RGReadDirect [GroupName | all]

### **RGRebuild (Immediate)**

RGRebuild starts rebuilding the specified existing RAID Group. Optional parameters MemberN specify the members to rebuild. If no member is specified, all degraded members will be rebuilt. Optional parameters BlockDevIDN allows an available block device to be substituted for the RAID Member currently assigned to the Member Index. RAID 6 groups can rebuild two members using the optional 'and'.

Usage: RGRebuild [GroupName] <Member1>  
<BlockDevID1> <and> <Member2> <BlockDevID2>

### **RGRecover (Immediate)**

RGRecover provides mechanisms for assisting in the read-only access of data in OFFLINE RAID Groups. Specify "Rebuild" if the RAID Group was rebuilding and the rebuild faulted. "Basic" forces a RAID Group ONLINE with only up-to-date members and returns CHECK CONDITION on READ errors. "Extreme" forces a RAID Group ONLINE with older members and replaces READ error data with zeros. "Disabled" turns off RGRecover for "Basic" and "Extreme". NOTE: Before running RGRecover, power off and power back on all drives in the affected RAID Group drives to ensure the drives are at a known state and ready for data recovery.

Usage: RGRecover [GroupName] [Rebuild | Basic | Extreme | Disabled]

### **RGRecoverWithWrites (Immediate)**

RGRecoverWithWrites performs the same functions as RGRecover except the affected RAID Group allows WRITES to the RAID Group to occur normally. NOTE: Care must be exercised to minimize WRITE activity. Proceed with WRITES at your own risk.

Usage: RGRecoverWithWrites [GroupName] [Basic | Extreme]

### **RGResumeConversion (Immediate)**

RGResumeConversion resumes the halted conversion on the specified existing RAID Group.

Usage: RGResumeConversion [GroupName]

### **RGResumeErase (Immediate)**

RGResumeErase resumes the erase on the specified existing RAID Group.

Usage: RGResumeErase [GroupName]

### **RGResumeInitialization (Immediate)**

RGResumeInitialization resumes the initialization of the specified RAID Group.

Usage: RGResumeInitialization [GroupName]

### **RGResumeMediaScan (Immediate, Disabled on Error)**

RGResumeMediaScan resumes a media scan on the specified existing RAID Group.

Usage: RGResumeMediaScan [GroupName]

### **RGResumeRebuild (Immediate)**

RGResumeRebuild resumes the rebuild(s) on the specified existing RAID Group. Optional parameter MemberIndex specifies the RAID Member whose halted rebuild will be resumed. For RAID6 Groups, if a MemberIndex is specified, all halted RAID Members on the span with that MemberIndex will resume as well. If no MemberIndex is specified, all halted rebuilds on that RAID Group will be resumed.

Usage: RGResumeRebuild [GroupName] <MemberIndex>

### **RGSectorSize**

Set or get the sector size of the specified RAID Group. The desired RAID Group sector size must be evenly divisible by the sector size of any member disk. 512 bytes is the default size for most operating systems. Use 4 KB sectors to enable large volume support (greater than 2 TB) in Windows XP (32-bit).

Usage: set RGSectorSize [GroupName] [512-8192]  
get RGSectorSize [GroupName | all]

### **RGSpanDepth**

Set or get the span depth on the specified existing NEW RAID Group. All RAID Group types are supported except JBOD, which implicitly supports spanning as members are added.

Usage: set RGSpanDepth [GroupName] [SpanDepth [1-16]]  
get RGSpanDepth [GroupName | all]

### **RGSpeedRead**

Perform look-ahead during reads from RAID Group member disks for all or the specified RAID Group. GroupName is the ASCII name of the RAID Group for which look-ahead reads will be performed. Auto will choose the algorithm based on each I/O command.

Usage: set RGSpeedRead [GroupName | all] [enabled | disabled | auto]  
get RGSpeedRead [GroupName | all]

### **RGUnmap (Immediate)**

RGUnmap removes all of the mapped partitions of the specified RAID Group or the specified RAID Group's partition or contiguous partitions from the routing table. The partitions themselves will be unaffected, though they will now be inaccessible by any initiators.

Usage: RGUnmap [[[GroupName] <PartIdx> <2-128>>] | all]

### **RGWaitTimeout (Disabled on Error)**

The RGWaitTimeout field specifies the number of seconds which the driver waits for devices to appear. The time out is used during system boot and when the BlockDevScan command is issued. You should adjust the time out setting when you have more than one RAID group in a chassis so that the chassis has enough time to bring all the drives involved in RAID groups online. If not, there is a chance that the RAID adapter may not discover all the drives for a RAID group and will either fault or degrade the RAID group.

Usage: set RGWaitTimeout [1-300]  
get RGWaitTimeout

### **RMStatus (Immediate)**

RMStatus displays the status of all RAID Members within the specified RAID Group or a specific RAID member (if specified) within the specified RAID Group. This command will fail if the specified RAID Group does not exist or a specified member index within the RAID Group does not exist.

GroupName is the ASCII name of the RAID Group for which status will be displayed.

Usage: RMStatus <[GroupName] <MemberIndex>>

### **Route (Immediate, Disabled on Error)**

Route is used to map a RAID Partition or SAS/SATA Pass Through device to a Target ID on the host system. If a map with the specified Target Id already exists, then it will be overwritten. Use 'Delete' for a Target ID to remove the map.

Usage: Route host [tid] [ [RAID [GroupName] [PartIdx]] | [SAS [SasIdx]] | Delete]

### **RouteDisplay (Immediate)**

RouteDisplay will display a list of host protocol address to target destination device mappings. The optional 'tid' parameter will limit the list to the maps which satisfy a search for the given Target ID. If the "Passthrough" parameter is entered then all SATA/SATA Pass Through devices are displayed. Otherwise all mapped RAID Partitions are displayed. If the "Persistent" parameter is entered then only persistent maps will be displayed. Otherwise both persistent and non-persistent maps will be displayed.

Usage: RouteDisplay host < <tid> | <<passthrough>> <<persistent>> >

### **SasPortList (Immediate)**

SasPortList lists the status of all available SAS ports.

Usage: SasPortList

### **SASTargets (Immediate, Disabled on Error)**

This command lists the physical devices that are connected to all SAS ports.

Usage: SASTargets

### **SaveConfiguration (Immediate, Disabled on Error)**

SaveConfiguration will save configuration changes. Please note that certain modifications require a system restart.

Usage: SaveConfiguration

### **SerialNumber**

SerialNumber displays the serial number. The serial number is a 13 character field. The first seven alphanumeric characters are an abbreviation representing the product name. The remaining six digits are the individual system's number.

Usage: get SerialNumber

### **SES (Disabled on Error)**

SES enables support for SES enclosures that have been discovered by the system.

Usage: set SES  
get SES [enabled | disabled]

### **SESAAlarmTest (Immediate, Disabled on Error)**

SESAAlarmTest commands the specified enclosure's audible alarm to be turned on at the specified warning level. "Reset" turns off the alarm that has been set at any warning level. Note that SESEnclosures must be executed prior to executing SESAAlarmTest.

Usage: SESAAlarmTest [EnclIdx] [SET | RESET] [INFO | NON-CRIT | CRIT | UNRECOV]

### **SESDiskFailureAlarm (Disabled on Error)**

SESDiskFailureAlarm when enabled, activates an audible alarm when the system determines that a RAID member disk drive has failed. The enclosure which contains the failed disk drive will be sounded, other enclosures will be unaffected.

Usage: set SESDiskFailureAlarm [enabled | disabled]  
get SESDiskFailureAlarm

### **SESEnclosures (Immediate, Disabled on Error)**

SESEnclosures displays a list of SES-enabled enclosures which have been discovered by the system.

Usage: SESEnclosures

### **SESIIdentify (Disabled on Error)**

SESIIdentify commands the appropriate SES enclosure to identify the specified element(s). "ALL" identifies all disks. "RAID" and RAID Group name identifies all disks in a RAID Group. If the MemberIndex is also specified, only that disk is identified. "ENC" and enclosure index identifies all slots in the specified enclosure. "DRIVE" and BlockDevID identifies the specified disk. Note that SESEnclosures must be executed prior to executing SESIdentify with the qualifier ENC and BlockDevScan must be executed prior to executing SESIdentify with the qualifier ALL.

Usage: set SESIdentify [ALL | RAID GrpName  
<MemberIndex> | ENC EnclIdx | DRIVE BlockDevID]  
get SESIdentify [ALL | RAID GrpName  
<MemberIndex> | ENC EnclIdx | DRIVE BlockDevID]

### **SESIIdentifyStop (Immediate, Disabled on Error)**

SESIIdentifyStop commands the appropriate SES enclosure to stop identifying the specified element(s). "ALL" stops identifying all enclosures' drive slots. "RAID" and RAID Group name stops identifying disks in a RAID Group. "ENC" and enclosure index stops identifying all slots in the specified enclosure. "DRIVE" and BlockDevID stops identifying the specified drive. Note that SESEnclosures must be executed prior to executing SESIdentify with the parameters ALL or ENC.

Usage: SESIdentifyStop [ALL | RAID GrpName  
<MemberIndex> | ENC EnclIdx | DRIVE BlockDevID]

### **SESMute (Immediate, Disabled on Error)**

SESMute causes all known enclosures' audible alarms to be set to either the "mute" or "remind" state. The default action is "mute". The enclosure index qualifier is optional. The optional parameter "REMIND" may be specified to set the "remind" state, which causes an occasional audible reminder of the alarm condition (if supported). Note that SESEnclosures must be executed prior to executing SESMute.

Usage: SESMute <EnclIdx> <REMIND>

### **SESPoll (Disabled on Error)**

SESPoll specifies the SES enclosure polling interval, in seconds. At the specified interval, all known SES enclosures are polled for their current status. A setting of 0 disables SES enclosure polling.

Usage: set SESPoll [0 | 30 - 3600]  
get SESPoll

### **SESStartingSlot (Disabled on Error)**

SESStartingSlot establishes the starting slot/ID number for all attached SES enclosures.

Usage: set SESStartingSlot [0 | 1]  
get SESStartingSlot

### **SESStatus (Immediate, Disabled on Error)**

SESStatus displays the last polled status of the specified element type in the specified enclosure. SupportLevel indicates the SES features supported by the specified enclosure: Fan,Power,Temp,Alarm, DriveLEDs. If no element type is specified, all status is displayed. Note that SESEnclosures must be executed prior to executing SESStatus.

Usage: SESStatus [EnclIdx | all] <ENC | DRIVE | FAN |  
POWER | TEMP | ALARM | SUPPORTLEVEL>

### **Temperature**

Temperature returns the current internal operating temperature in degrees Celsius. The value is read-only.

Usage: get Temperature

### **Time (Disabled on Error)**

Time sets/displays the current time in 24 hour format.

Usage: set Time [HH:MM:SS]  
get Time

### **TimeZone (Disabled on Error)**

Timezone sets/displays the time zone or an offset from GMT. GMT offset must be in the format +/-HH:MM

Usage: set TimeZone [[EST | CST | MST | PST] | [[+|-]  
[HH]:[MM]]]  
get TimeZone

**VerboseMode**

VerboseMode controls the level of detail in CLI 'Help' output and command response output for the current CLI session.

Usage: set VerboseMode [enabled | disabled]  
get VerboseMode

**VirtualDriveInfo (Immediate)**

VirtualDriveInfo displays characteristics and statistics for all the available virtual drives or any available virtual drive identified by its virtual drive ID.

Usage: VirtualDriveInfo <VirtualDrive ID>

## Appendix B Glossary

---

Some terms used in the Fibre Channel industry are defined below. More information is available through the ATTO Technology website ([www.attotech.com](http://www.attotech.com)), the Fibre Channel Industry Association, Cern, the Storage Area Networking Industry Association ([www.snia.org](http://www.snia.org)), and the Fibre Channel Consortium.

Term	Definition
ANSI	American National Standards Institute.
BER	Bit Error Rate: a measure of transmission accuracy; the ratio of bits received in error to bits sent.
bit	The smallest unit of data a computer can process: a single binary digit, either 0 or 1.
bus	A collection of unbroken signal lines used to transmit information from one part of a computer system to another. Taps on the lines connect devices to the bus.
Byte	An ordered set of 8 bits.
channel	A Point-to-Point link which transports data from one point to another.
CPU	Central Processing Unit: the portion of the computer which performs computations.
CRC	Cyclic Redundancy Check: an error-correcting code which calculates a numeric value for received and transmitted data. If no error has occurred during transmission, the CRC for both received and transmitted data should be the same.
destination address	A value in the frame header of each frame which identifies the port in the node where the frame is being sent.
device driver	A program which allows a microprocessor to direct the operation of a peripheral device.
DMA	Direct Memory Access: a way to move data from a storage device directly to RAM without using the CPU's resources.
DMA bus master	Allows a peripheral to control the flow of data to and from system memory by block as opposed to allowing the processor to control the data by bytes (PIO or programmed I/O).
full duplex	A communication protocol which allows transmission in both directions at the same time.
half duplex	A communication protocol which allows transmission in both directions, but only one direction at a time.
host	A processor, usually a CPU and memory, which communicates with devices over an interface.
initiator device	A component which originates a command.
LED	Light-Emitting Diode: a type of diode which emits light when current passes through it. Visible LEDs are used as indicator lights on all sorts of electronic devices.
LUN	Logical Unit Number: an identifier for a logical unit (0-7).
originator	An initiating device; a component which originates a command.
parity checking	A method which verifies the accuracy of data transmitted over the SCSI bus by adding one bit in the transfer to make the sum of all the bits either odd or even (for odd or even parity). An error message occurs if the sum is not correct.
PCI	Peripheral Component Interconnect. allows peripherals to be connected directly to computer memory, bypassing the slower ISA and EISA busses.
Point-to-Point	A topology where two ports communicate.
port address	The address, assigned by the PCI bus, through which commands are sent to the storage controller. Also known as "port number".
read direct	A RAID Group feature in which data is passed directly to the OS without first being read into the storage controller cache. This can result in read performance improvements, especially for large, sequential transfers. This feature is only available for RAID 5.

<b>Term</b>	<b>Definition</b>
receiver	The ultimate destination of data transmission; a terminal device.
SCSI	Small Computer Systems Interface: a processor-independent standard for system-level interface between a computer and intelligent devices including hard disks, floppy disks, CD-ROM, printers, scanners, etc.
SES	SCSI Enclosure Services: a processor which identifies individual drives, all the drives in the same enclosure, all the drives in a single RAID group and faulted drives; selects drives, and monitors the status of the enclosure.
topology	The logical layout of the parts of a computer system or network and their interconnections.
transceiver	A transmitter/receiver module.
transfer rate	The rate at which bytes or bits are transferred, as in megabytes or gigabits per second.