



Hewlett Packard
Enterprise

HPE Smart Array SR Gen10 Configuration Guide

Abstract

This document identifies, and provides instructions for, the array configuration tools available for Hewlett Packard Enterprise controller and server products. This document is for the person who installs, administers, and troubleshoots servers and storage systems. Hewlett Packard Enterprise assumes you are qualified in the servicing of computer equipment and trained in recognizing hazards in products with hazardous energy levels.

Notices

© Copyright Hewlett Packard Enterprise Development LP

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

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

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

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

Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.



Contents

- Overview..... 6**
 - Array and controller configuration..... 6
 - Comparison of HPE SSA and UEFI System Utilities..... 6
 - UEFI System Utilities..... 7
 - HPE Smart Storage Administrator..... 8
 - Array configuration guidelines..... 16
- Smart Array configuration utility in UEFI System Utilities..... 17**
 - Viewing controller information..... 17
 - Configure controller settings..... 17
 - Modifying controller settings..... 17
 - Modifying advanced controller settings..... 19
 - Clearing the controller configuration..... 20
 - Viewing the backup power source status..... 21
 - Managing power settings..... 21
 - Configure arrays..... 22
 - Creating an array using UEFI System Utilities..... 22
 - Managing an array..... 23
 - Editing a logical drive..... 25
 - Deleting a logical drive..... 25
 - Disk Utilities..... 26
 - Viewing disk device information..... 26
 - Identifying a disk device..... 26
 - Setting bootable devices for Legacy Boot Mode..... 26
 - Setting the primary and secondary bootable devices (Legacy Boot Mode)..... 26
 - Setting the number of OS bootable drives (Legacy Boot Mode)..... 27
- HPE SSA graphical user interface..... 28**
 - Using the HPE SSA GUI..... 28
 - Icons..... 28
 - Keyboard shortcuts..... 29
 - Navigating the GUI..... 30
 - Configure the language of the GUI..... 34
 - Configuration tasks..... 34
 - Performing a Configuration task..... 38
 - Creating an array..... 38
 - Configuring a controller..... 39
 - Enabling or disabling HPE SSD Smart Path..... 40
 - Setting sanitize lock..... 41
 - Rapid parity initialization..... 41
 - Changing the Spare Activation Mode..... 43
 - Changing the Spare Management mode..... 44
 - Disabling Over Provisioning Optimization..... 45
 - Setting Auto RAID 0..... 45
 - Modifying parallel surface scan..... 46
 - Configuring controller cache..... 46



HPE Smart Array SR SmartCache.....	47
Working with mirrored arrays.....	50
Heal array.....	53
Change array drive type.....	54
Encryption Manager.....	54
Power modes.....	54
Viewing controller status.....	56
Options for erasing drives.....	56
Erasing a drive.....	57
Managing FLS.....	58
Diagnostics tasks.....	59
Performing a Diagnostics task using HPE SSA.....	60
Maintenance tasks.....	60
Checking online firmware activation readiness.....	61

HPE SSA CLI..... 62

Using the HPE SSA CLI.....	62
Opening the CLI in Console mode.....	62
Opening the CLI in Command mode.....	63
CLI syntax.....	63
Typical procedures.....	68
Encryption commands.....	88

HPE SSA Scripting..... 95

Using HPE SSA scripting.....	95
Capturing a configuration.....	95
Using an Input script.....	95
Creating an HPE SSA Scripting script file.....	95
Script file options.....	100
Basic scripting scenarios.....	120
Advanced commands.....	124
Advanced scripting scenarios.....	129
XML support.....	134
HPE SSA scripting warning messages.....	137
HPE SSA scripting error messages.....	138

Troubleshooting..... 146

HPE Smart Storage Administrator Diagnostics Utility CLI.....	146
About the utility.....	146
Reported information.....	146
Installing the utility.....	148
Launching the utility in CLI mode.....	148
Diagnostic report procedures.....	149
SmartSSD Wear Gauge report procedures.....	150
512e Physical drive support.....	150
Drive arrays and fault-tolerance methods.....	151
Drive arrays.....	151
Effects of a hard drive failure on logical drives.....	152
Fault-tolerance methods.....	153
Diagnosing array problems.....	160
Diagnostic tools.....	160
Troubleshooting resources.....	160



Optional components..... 161
HPE Smart Array SR Secure Encryption..... 161

HPE Smart Array Advanced Pack..... 162
About SAAP..... 162

Support and other resources..... 163
Accessing Hewlett Packard Enterprise Support..... 163
Accessing updates..... 163
Customer self repair..... 164
Remote support..... 164
Warranty information..... 164
Regulatory information..... 165
Documentation feedback..... 165
Websites..... 165



Overview

Array and controller configuration

You can configure arrays and controllers during the initial provisioning of the server or compute module and at any time after the initial configuration. Configuration tasks can be initiated using HPE Smart Storage Administrator (accessible through Intelligent Provisioning) or the configuration menus of the UEFI System Utilities.

During the initial provisioning of the server or compute module, an array is required to be configured before the operating system can be installed. You can configure the array using either of the options below:

- When you launch Intelligent Provisioning, you can specify options that enable Intelligent Provisioning to poll for any drives that are present and build an appropriate array for those drives. For example, if two drives are connected to the Smart Array card, the setup defaults to RAID 1. Hewlett Packard Enterprise recommends selecting this option when initially provisioning a server. For more information, see the Intelligent Provisioning documentation.
- You can use the UEFI System Utilities to create the primary array that is required.

After the initial provisioning of the server or compute module, you can use either HPE SSA or the UEFI System Utilities to configure the arrays and controllers.

Comparison of HPE SSA and UEFI System Utilities

HPE Smart Array SR Gen10 can be configured by using either HPE SSA or the Smart Array Configuration Utility within the UEFI System Utilities. Both HPE SSA and UEFI System Utilities can be used to configure the Smart Array S-class, E-class, and P-class.

HPE SSA provides a full set of array configuration features while the UEFI System Utilities provides a limited set of features. However, users may prefer using the UEFI System Utilities during the initial configuration of the server or compute module because the UEFI System Utilities loads faster than HPE SSA during that step.

To identify the standard configuration tasks that are supported within each interface, review the table.

Task	HPE SSA	UEFI System Utilities
Create or delete arrays and logical drives	+	+
Assign a RAID level to a logical drive	+	+
Identify devices by causing the LEDs to illuminate	+	+
Assign or delete a spare drive	+	+
Share a spare drive among several arrays	+	+
Assign multiple spare drives to an array	+	+
Set the spare activation mode	+	+
Specify the size of the logical drive	+	+
Create multiple logical drives per array	+	+
Set the stripe size	+	+
Migrate the RAID level or stripe size	+	

Table Continued



Task	HPE SSA	UEFI System Utilities
Expand an array	+	
Set the expand priority and migrate priority	+	
Set the cache ratio (accelerator) priority	+	+
Extend a logical drive	+	
Set the boot controller	+	
Configure HPE Smart Array SR SmartCache	+	

UEFI System Utilities

The UEFI System Utilities is embedded in the system ROM. The UEFI System Utilities enable you to perform a wide range of configuration activities, including:

- Configuring system devices and installed options
- Enabling and disabling system features
- Displaying system information
- Selecting the primary boot controller
- Configuring memory options
- Selecting a language
- Launching other pre-boot environments such as the Embedded UEFI Shell and Intelligent Provisioning

For more information on the UEFI System Utilities, see the *UEFI System Utilities User Guide for HPE ProLiant Gen10 Servers* on the **Hewlett Packard Enterprise website**.

Scan the QR code located at the bottom of the screen to access mobile-ready online help for the UEFI System Utilities and UEFI Shell. For on-screen help, press **F1**.

Launching the System Utilities

Procedure

1. Optional: If you access the server remotely, start an iLO remote console session.
 - a. Open a browser and enter `https://<iLO host name or IP address>` to log on to the iLO web interface.
 - b. On the login page, enter a directory or local user account name and password, and then click **Log In**.
 - c. Select **Remote Console & Media** in the iLO navigation tree.
The Launch tab is displayed.
 - d. Verify that your system meets the requirements for using the remote console application you want to use.
 - e. Click the launch button for your selected application.
You can also launch an iLO Remote Console session by selecting:

- The Integrated Remote Console link on the **Information - iLO Overview** page.
- The Console thumbnail in the low left corner of the iLO web interface, and then choosing the application type to launch.

2. Restart or power on the server.

The server restarts and the POST screen appears.

3. Press **F9**.

The **System Utilities** screen appears.

NOTE: Using System Utilities requires BIOS administrator authorization. If the BIOS administrator requires a password, the server prompts for the password to be entered prior to launching the System Utilities. For information on setting the administrator password, see **Server Security options**.

Using UEFI System Utilities

To use the System Utilities, use the following keys.

Action	Key
Access System Utilities	F9 during server POST
Navigate menus	Up and Down arrows
Select items	Enter
Save selections	F10
Access Help for a highlighted configuration option ¹	F1

¹ Scan the QR code on the screen to access online help for the UEFI System Utilities and UEFI Shell.

Default configuration settings are applied to the server at one of the following times:

- Upon the first system power-up
- After defaults have been restored

Default configuration settings are sufficient for typical server operations; however, you can modify configuration settings as needed. The system prompts you for access to the UEFI System Utilities each time the system is powered up.

HPE Smart Storage Administrator

HPE SSA is the main tool for configuring arrays on Smart Array controllers. It exists in three interface formats: the HPE SSA GUI, the HPE SSA CLI, and HPE SSA Scripting. All formats provide support for configuration tasks. Some of the advanced tasks are available in only one format.

The diagnostic features in HPE SSA are also available in the standalone software HPE Smart Storage Administrator Diagnostics Utility CLI.

HPE SSA is accessible both offline and online:

- **Accessing HPE SSA in the offline environment:** Using one of multiple methods, you can run HPE SSA before launching the host operating system. In offline mode, users can configure or maintain detected and supported HPE ProLiant devices, such as optional Smart Array controllers and integrated Smart Array controllers. Some HPE SSA

features are only available in the offline environment, such as setting the boot controller or performing split-mirror operations.

- **Accessing HPE SSA in the online environment:** This method requires an administrator to download the HPE SSA executables and install them. You can run HPE SSA online after launching the host operating system.

More information

[Accessing HPE SSA in the offline environment](#)

[Accessing HPE SSA in the online environment](#)

Benefits of using HPE SSA

HPE Smart Storage Administrator is an advanced utility that enables you to perform many complex configuration tasks.

HPE SSA includes a stand-alone **CLI diagnostics utility**.

Using HPE SSA over other configuration utilities provides the following benefits:

- GUI, CLI, and Scripting interfaces are available.
- English, French, German, Italian, Japanese, Portuguese, Russian, Simplified Chinese, and Spanish languages are supported.
- Applications can be executed using the following tools:
 - Any bootable media, such as a software CD
 - Most supported server platform host operating systems with a web browser (For a list of supported operating systems, see the **Hewlett Packard Enterprise website**.)
 - Intelligent Provisioning
- All formats can run in both online and offline environments.
- The utility can run on any machine that uses a supported browser.

Support for configuration tasks

HPE SSA supports the configuration tasks listed below. Support for individual tasks varies according to your controller type. You should always have the latest firmware installed so that you have access to all of the features supported for your controller.

- Activate or delete license keys
- Assign a RAID level to a logical drive
- Assign spare drives to an array
- Configure multiple systems identically
- Copy the configuration of one system to multiple systems
- Create multiple logical drives per array
- Create or delete arrays and logical drives
- Enable or disable a physical drive write cache on physical drives that are configured as part of a logical volume
- Enable or disable physical drive write cache on unconfigured physical drives
- Enable solid state drives to be used as caching devices, using SmartCache (requires a registered SmartCache license key)

- Enable optimized data path to solid state drives using HPE SSD Smart Path
- Expand an array
- Extend a logical drive
- Heal an array
- Drive Erase (replace the content of a physical drive with zeros or random 0 and 1)
- Drive Erase (with sanitize options)
- Identify devices by causing their LEDs to flash
- Migrate the RAID level or stripe size
- Move an array (copy all array data to a new array and then delete the old array)
- Move and delete individual LUNs
- Optimize the controller performance for video
- Re-enable a failed logical drive
- Set the boot controller
- Set the expand priority, migrate priority, and accelerator ratio
- Set the spare activation mode
- Set the stripe size
- Set the surface scan delay
- Share a spare drive among several arrays
- Remove a drive from an array (restripe the data on an array to occupy fewer physical drives, then remove the excess drives from the array)
- Specify the size of the logical drive
- Split a RAID 1 array or recombine a split array (offline only)
- Split mirror backup and rollback of RAID 1, 1+0, 1 (ADM) and 10 (ADM) mirrors

Minimum requirements

For minimum operating system requirements to run any HPE SSA format, see the [Hewlett Packard Enterprise website](#).

Minimum video requirements to run the HPE SSA GUI include a minimum monitor resolution of 1024x768 and 16-bit color. The GUI supports the following browsers:

- Mozilla Firefox 9.0 or later
- Microsoft Internet Explorer 9.0 or later
- Google Chrome

For a list of supported controllers, see the [Hewlett Packard Enterprise website](#).

Native support for 64-bit and 32-bit operating systems

HPE SSA now offers a native 64-bit HPE SSA application for supported 64-bit operating systems, eliminating the need for compatibility libraries. A 32-bit HPE SSA application is also available. Users can choose to install the application that matches the OS installed on the server product.



The 64-bit HPE SSA application is not a direct upgrade from 32-bit HPE SSA. On 64-bit systems running 32-bit HPE SSA, you must uninstall the 32-bit application, and then install the 64-bit application.

Future versions of HPE SSA will be available as native 32-bit or native 64-bit applications, as long as 32-bit operating systems are supported.

Accessing HPE SSA

Accessing HPE SSA in the offline environment

To access and launch the HPE SSA GUI in an offline environment, use one of the following methods:

- **Launching HPE SSA with Intelligent Provisioning**
- **Launching HPE SSA during POST**
- **Launching HPE SSA from an ISO image**

To access the **HPE SSA CLI or HPE SSA Scripting in an offline environment**, you must launch HPE SSA from an ISO image.

When offline HPE SSA launches, an **Execution Mode** screen does not appear because HPE SSA does not support Remote Service Mode in an offline environment. For this functionality, use **HPE SSA in an online environment**.

Launching HPE SSA with Intelligent Provisioning

Procedure

1. Boot the server or compute module.
2. Press **F10** to launch Intelligent Provisioning.
3. At the main screen, select **Perform Maintenance**.
4. At the **Maintenance** screen, select **Intelligent Storage Configuration**.
5. At the **Controller Detail** screen, select **Configuration**.

The system launches the HPE SSA GUI.

Launching HPE SSA during POST

1. Boot the server or compute module.
During POST, the system recognizes devices.

2. Press **F10** to start Intelligent Provisioning.
A menu appears, listing the options for launching HPE SSA.

3. Select a method to launch HPE SSA.
The system launches the HPE SSA GUI, or if you are using Serial Console, the system launches the HPE SSA CLI.

Launching HPE SSA from an ISO image

To launch HPE SSA, you can also boot from an ISO image. To prepare the image, use one of the following methods:

- **Mounting the image through iLO**
- **Burning the image to a CD or DVD**

- **Flashing the image to a USB memory key or SD card on a HPE UEFI bootable server**
- **Installing the image on a PXE server**

Booting from the ISO image on a drive, on a key, or through iLO provides the same GUI interface. The user can select to run Offline HPE SSA GUI, HPE SSA CLI, or HPE SSA Scripting.

Mounting the image through iLO

This iLO functionality requires an iLO Advanced license.

To mount the image:

1. Download the ProLiant Offline HPE Smart Storage Administrator ISO image from the **Hewlett Packard Enterprise website**.
2. Browse to the server iLO page.
3. Start the remote console for the server.
4. On the remote console, use the iLO mount feature to browse to the location of the ISO image.
5. Select the ISO image to be mounted.
6. Reboot the server.
7. Press F11 to access the One Time Boot Menu.
8. Select the iLO Virtual USB CD/DVD ROM menu item.

Burning the image to a CD or DVD

1. Download the ProLiant Offline HPE Smart Storage Administrator ISO image from the **Hewlett Packard Enterprise website**.
2. Use third-party software to burn the ISO image to a CD or DVD.
3. Set the server to boot from the optical drive.
4. Insert the CD or DVD.
5. Reboot the server.

Flashing the image to a USB memory key or SD card on a UEFI bootable server

Booting from a USB key is supported only for UEFI mode.

Procedure

1. Download the ProLiant Offline HPE Smart Storage Administrator ISO image from the **Hewlett Packard Enterprise website**.



CAUTION: Before creating a bootable USB key from the Offline HPE SSA ISO image, back up any critical data stored on the key to a different location. The utility overwrites all data on the key.

2. Using ISO mounting software, mount the Offline HPE SSA ISO image to a Windows drive.

For this example, use "E:."

3. Insert a USB key into a USB connector on the Windows system.

For this example, use "F:."



4. Format the USB key, and copy the contents of the mounted ISO to the USB key.
5. Insert the USB key into the server.

The USB key can now be used to boot to the Offline HPE Smart Storage Administrator environment. When booting from the USB key, a menu is presented. Select "USB BOOT: Smart Storage Administrator (SSA)" to mount the correct device during the boot sequence.

Installing the image on a PXE server

To install the Offline HPE SSA ISO image on a PXE server and boot from the image over a network, use the following procedures:

1. Review the **prerequisites**.
2. Set up PXELinux.
3. **Configure PXELinux**.
4. Specify the ISO image path.

Depending on the network configuration, boot times might vary.

Prerequisites

Before proceeding with the configuration, you must have all the prerequisites:

- A good, working knowledge of PXE and TFTP
- A network with a DHCP server on it
- A TFTP server configured on the same network as the DHCP server
- A network file server that hosts the ISO images and is accessible by a PXE booted system
- **PXELinux**

These instructions presume that you are using a Linux TFTP server and the **TFTP package**. Other TFTP servers should work similarly.

Set up PXELinux

Before proceeding with the configuration, ensure that your TFTP server and PXELinux configuration is set up and configured properly.

To set up PXELinux:

Procedure

1. Download the ProLiant Offline HPE Smart Storage Administrator ISO image from the **Hewlett Packard Enterprise website**.

2. Copy the ISO image to the network file system, and note the location. NFS and Windows file shares are supported.

For this example, use the following NFS and path to the ISO image:

```
192.168.0.99:/path/to/ahpssacd/image/ssaooffline-2.60.x.x.iso
```

3. Before proceeding, test your network file system to ensure that it is accessible.
4. Access the `/system` directory of the CD in one of the following ways:



- Burn and mount the ISO image.
 - Extract the ISO image using a third-party tool.
5. Copy all the files from the `/system` directory of the CD to your TFTP server so that it is accessible by the TFTP software.

Configure PXELinux

1. Using the `isolinux.cfg` file from the `/system/` directory of the CD as a guide, copy the labeled targets to your PXELinux configuration file. You do not need to include the entire file:

```
label sos

MENU LABEL HP ProLiant Offline HP SSA Image

kernel hpboot_v.c32

append vmlinuz initrd=initrd.img media=net rw root=/dev/ram0
ramdisk_size=257144 init=/init loglevel=3 ide=nodma ide=noraid pnpbios=off
vga=791 splash=silent showopts TYPE=AUTOMATIC

label vsos

MENU LABEL HP ProLiant Offline HP SSA Image

kernel hpboot_v.c32

append vmlinuz initrd=initrd.img media=net rw root=/dev/ram0
ramdisk_size=257144 init=/init loglevel=3 ide=nodma ide=noraid pnpbios=off
vga=791 splash=silent showopts TYPE=MANUAL
```

2. Replace the lines `kernel hpboot_v.c32` with `kernel vmlinuz`.
3. Remove `vmlinuz` from the append line.

The paths to the files on the TFTP server are `vmlinuz` and `initrd.img`. You must modify them to include any directories or naming conventions you may have on your TFTP server.

Specify the ISO image path

For the PXE-booted server to find the ISO image, you must add the ISO image path to the append line in the PXELinux configuration file.

Add the following arguments:

```
iso1=nfs://192.168.0.99/path/to/hpssacd/image/ssaoffline-2.60.x.x.iso
isolmnt=/mnt/bootdevice
```

The `iso1` parameter helps the PXE-booted HPE SSA Offline CD locate the ISO image. The `iso1mnt` parameter tells the PXE-booted HPE SSA CD where the `iso1` image must be mounted.

Your final configuration must be similar to the following example:

```
label sos

MENU LABEL HP ProLiant Offline HPS SA Image

kernel vmlinuz

append initrd=initrd.img media=net rw root=/dev/ram0 ramdisk_size=257144
init=/init loglevel=3 ide=nodma ide=noraid pnpbios=off vga=791 splash=silent
showopts TYPE=AUTOMATIC iso1=nfs://192.168.0.99/path/to/hpssacd/image/
ssaoffline-2.60.x.x.iso isolmnt=/mnt/bootdevice
```

```
label vsos
MENU LABEL HP ProLiant Offline HP SSA Image
kernel vmlinuz
append initrd=initrd.img media=net rw root=/dev/ram0 ramdisk_size=257144
init=/init loglevel=3 ide=nodma ide=noraid pnpbios=off vga=791 splash=silent
showopts TYPE=MANUAL iso1=nfs://192.168.0.99/path/to/hpssacd/image/
ssaoffline-2.60.x.x.iso iso1mnt=/mnt/bootdevice
```

You can add additional ISO images by specifying the additional iso# and iso#mnt arguments, for example, iso2=/path/to/iso2.iso iso2mnt=/mnt/iso2.

Supported network file systems

The following network file systems are supported for use with PXE booting:

- NFS:

```
iso1=nfs://192.168.0.99/path/to/hpssacd/image/ssaoffline-2.60x.x.iso
iso1mnt=/mnt/bootdevice
```

NFS volumes are mounted with the following options:

- -o ro
- nolock

- Windows operating systems:

```
iso1=smbfs://192.168.0.99/share/path/to/hpssacd/image/
ssaoffline-2.60.x.x.iso iso1mnt=/mnt/bootdevice
```

- Windows operating systems with logon credentials:

```
iso1=smbfs://user:password@192.168.0.99/share/path/to/hpssacd/image/
ssaoffline-2.60.x.x.iso iso1mnt=/mnt/bootdevice
```

Accessing HPE SSA in the online environment

To access, install, and launch HPE SSA in the online environment, you must download the HPE SSA executables. All three formats have separate executables.

HPE SSA Scripting is a standalone application that is distributed with the HPE SSA CLI application.

Both HPE SSA and the HPE SSA CLI need sg driver (scsi generic) to access Smart Array controller in Linux. Red Hat Enterprise Linux 7.1 and later does not load sg driver automatically. You should load the driver and type 'modprobe sg' before running HPE SSA or HPE SSA-CLI.

To use HPE SSA in the online environment, obtain the executable files from the [Hewlett Packard Enterprise website](#).

Procedure

1. Follow the installation instructions provided with the executable.
2. After the executables are installed, launch each executable in the following manner:
 - GUI—Click **Start**, and then select **Programs>Windows System>Smart Storage Administrator**.
 - CLI—Click **Start**, and then select **Programs>Windows System>Smart Storage Administrator CLI**
 - Scripting—Run `ssascripting.exe` (Windows) or `ssascripting` (Linux).



Launching HPE SSA on a local server

Microsoft OS

Procedure

1. Click **Start**, and then select **Programs>Windows System>Smart Storage Administrator**.

HPE SSA launches in either a browser or application window (v1.50 and later). HPE SSA then scans the system and detects controllers. When controller detection is complete, the controllers are available on the **Controller/Device** menu.

2. Configure a controller. For more information, see "[Configuring a controller](#)."

When configuration is complete, continue with the next step.

3. (Optional) To make newly created logical drives available for data storage, use the operating system disk management tools to create partitions and format the drives.

Linux OS

1. From any command prompt, enter one of the following:

- For local mode, enter: `ssa -local`
- For remote mode, enter: `ssa -start`

HPE SSA launches in a browser window.

2. For a list of options, enter the following: `ssa -h`

Array configuration guidelines

Remember the following factors when you build an array:

- All drives grouped in a logical drive must be of the same type (for example, either all SAS or all SATA and either all hard drives or all solid state drives).
- For the most efficient use of drive space, all drives within an array should have approximately the same capacity. Each configuration utility treats every physical drive in an array as if it has the same capacity as the smallest drive in the array. Any excess capacity of a particular drive cannot be used in the array and is unavailable for data storage.
- The more physical drives configured in an array, the greater the probability that the array will experience a drive failure during any given period.
- To guard against the data loss that occurs when a drive fails, configure all logical drives in an array with a suitable fault-tolerance (RAID) method. For more information, see "[Drive arrays and fault-tolerance methods](#)."

Smart Array configuration utility in UEFI System Utilities

This section contains information about using the Smart Array Configuration utility within the UEFI System Utilities.

For more information about the options mentioned in this section, see the [Comparison of HPE SSA and UEFI System Utilities](#) section in this guide.

Viewing controller information

Procedure

1. From the System Utilities screen, select **System Configuration > controller > Controller Information**.
2. In the Controller Information screen, view the information.

Configure controller settings

Modifying controller settings

Procedure

1. From the System Utilities screen, select **System Configuration > controller > Configure Controller Settings > Modify Controller Settings**.
2. In the Modify Controller Settings screen, modify any of the following settings:

Setting	Description
Cache Ratio (Read)	Adjusts the amount of memory for read-ahead cache versus write cache. Range is from 0-100; value can be increased or decreased in increments of 5.
Configured Physical Drive Write Cache State	Enables or disables the setting for the write cache on all configured physical drives. Options are Enabled, Disabled, or Default.
Current parallel surface scan count	Controls how many controller surface scans can operate in parallel: <ul style="list-style-type: none">• 1: Disabled• 16: Maximum

Table Continued



Setting	Description
No Battery Write Cache	<p>Enables or disables write cache when the energy pack is not present or not charged.</p> <p>Options are Enabled or Disabled.</p>
Rebuild priority	<p>Determines the urgency with which the controller treats an internal command to rebuild a failed logical drive.</p> <ul style="list-style-type: none"> • Low: Normal system operations take priority over a rebuild. • Medium: Rebuilding occurs for half of the time, and normal system operations occur for the rest of the time. • Medium high: Rebuilding is given a higher priority over normal system operations. • High: The rebuild takes precedence over all other system operations.
Spare Activation Mode	<p>Predictive Spare Activation mode activates a spare drive any time a member drive within an array reports a predictive failure.</p> <p>Failure Spare Activation mode activates a spare drive when a member drive within an array fails using fault tolerance methods to regenerate the data.</p>
Surface Scan Analysis Priority	<p>Modifies the amount of delay or idle time of the controller before surface scan analysis is resumed.</p> <ul style="list-style-type: none"> • 0: Disabled • 1-30: Idle with delay • 31: High

Table Continued



Setting	Description
Transformation priority	(Does not apply to HPE Smart Array S100i SR Gen10 SW RAID.) Rate at which requests from the operating system are processed: <ul style="list-style-type: none"> High: Completes as fast as possible at the expense of normal I/O. Medium: Completes with some impact on normal I/O. Low: Performs when normal I/O is not occurring.
Unconfigured physical drive write cache state	Enables or disables the write cache on all unconfigured physical drives. Options are Enabled, Disabled, or Default.

3. Click **Submit changes**.

Modifying advanced controller settings

Procedure

1. From the System Utilities screen, select **System Configuration** > **controller** > **Configure Controller Settings** > **Advanced Controller Settings**.
2. In the Advanced Controller Settings screen, modify any of the following settings:

Setting	Description
Alternate Inconsistency Repair Policy	(Does not apply to HPE Smart Array S100i SR Gen10 SW RAID.) Controls the behavior of the controller inconsistency repair Policy. This option is used to tune the controller performance for Video applications and requires the installation of a valid license key. Options are Enable or Disable.
Degraded Mode Performance Optimization	Used to tune controller performance for video applications and requires the installation of a valid license key. Options are Enabled or Disabled.
HDD Flexible Latency Optimization	Reduces the maximum observed latency from a host request.

Table Continued



Setting	Description
Maximum drive Request Queue Depth	Controls the maximum number of physical drive requests that the firmware will submit to a drive at any given time. This option is used to tune controller performance for video applications. Options are 2, 4, 8, 16, 32, or Automatic.
Monitor and Performance Analysis Delay	Controls the behavior of the controller Monitor and Performance Analysis Delay and is expressed in values ranging from 0 to 60. This option is primarily used to tune controller performance for video applications and requires the installation of a valid license key.
Physical Drive Request Elevator Sort	Controls the behavior of the controller cache write Elevator sort algorithm. This option is used to tune controller performance for video applications and requires the installation of a valid license key. Options are Enabled or Disabled.
RAID 6/60 Alternate Inconsistency Repair Policy	Sets the Inconsistency Repair Policy for the controller. Options are Enabled and Disabled.

3. Click **Submit changes**.

Clearing the controller configuration

Clearing the controller configuration destroys the controller metadata, including array configurations and partition information.



CAUTION: When you clear the controller configuration, all data on the attached media is no longer accessible and cannot be recovered.

Procedure

1. From the System Utilities screen, select **System Configuration** > **controller** > **Configure Controller Settings** > **Clear Configuration**.
2. In the Clear Configuration screen, select one or both of the following:
 - **Delete All Array Configurations**—deletes all the arrays in the controller. All the data in the arrays is also deleted.
 - **Delete Configuration Metadata on All Physical Drives**—deletes RAID metadata on the drives that are not part of the array.



Viewing the backup power source status

Procedure

1. From the System Utilities screen, select **System Configuration > controller > Configure Controller Settings**.
2. In the Backup Power Source screen, view the status of the backup power.

Status options are:

- Failed
- Not present
- Charging
- Charged

Managing power settings

The power management features do not apply to HPE Smart Array S100i SR Gen10 SW RAID.

Procedure

1. From the System Utilities screen, select **System Configuration > controller > Configure Controller Settings > Manage Power Settings**.
2. In the Manage Power Settings screen, update any of the following settings.

Setting	Description
Power Mode	<p>Options are:</p> <ul style="list-style-type: none">• Maximum performance (default)—Power savings options that affect performance are disabled.• Balanced—Use this setting to save power with minimal effects on performance.• Minimum power—When settings are selected without regard to system performance, maximum power savings is achieved. <p>NOTE: Hewlett Packard Enterprise recommends the minimum power setting for specific applications, but it is not appropriate for most customers. Most applications will suffer significant performance reduction.</p>
Survival Mode	<p>Allows controller to throttle back dynamic power settings to minimum when temperature exceeds the threshold. This minimum setting allows the server to run in most situations, but performance might decrease.</p>

3. Click **Submit changes**.



Configure arrays

Creating an array using UEFI System Utilities

When you create an array, you can select drives, specify RAID level, and configure array settings, including strip size and logical drive size.

Procedure

1. From the UEFI System Utilities screen, select **System Configuration >controller > Array Configuration > Create Array**.
2. In the Create Array screen, select each drive that you want to include in the array and click **Proceed to next Form**.

NOTE: SAS drives are not supported on the S100i SW RAID.

3. In the Set RAID Level screen, select the RAID Level from the drop-down menu and click **Proceed to next Form**.
4. In the Set Logical Drive Configuration screen, specify the configuration settings or use the default selection.

Setting	Description
Logical Drive Label	Use the default selection for the drive label or enter a new label. The characters in the label can be alphanumeric or spaces.
Strip Size/Full Stripe Size	<p>Strip size is the amount of data that is stored on each physical drive in the array. The full stripe size is the amount of data that the controller can read or write simultaneously on all the drives in the array. For RAID levels that support fault tolerance through parity, the parity information is calculated one full strip size at a time.</p> <p>For hardware RAID, you can specify from 16KiB up to 1024KiB, depending on the number of disks and RAID level. The default value is all available space.</p> <p>For Smart Array S100i SW RAID, the minimum size is 16KiB and the maximum size is up to 256KiB, depending on RAID level and device type.</p>
Size	Values in decimal; minimum RAID size is 16 MiB.
Unit Size	Logical drive unit size (MiB/GiB/TiB).
Acceleration Method	Logical drive acceleration method (controller cache or none).

5. Click **Submit Changes**.
6. Return to the main menu.
7. Click **OK** when prompted to save your changes.
8. Reboot the server.



Managing an array

Viewing logical drive properties

Procedure

1. From the System Utilities screen, select **System Configuration > controller > Array Configuration > Manage Arrays > array > List Logical Drives > logical drive > Logical Drive Details**.
2. In the Logical Drive Details screen, view the details.

Creating a logical drive

Procedure

1. From the System Utilities screen, select **System Configuration > controller > Array Configuration > Manage Arrays > array > Create Logical Drive**.
2. In the Create Logical Drive screen, select the RAID level, and then click **Proceed to next Form**.
3. In the Set Logical Drive Configuration screen, use the default values for the configuration or specify different values.

Setting	Description
Logical Drive Label	Use the default selection for the drive label or enter a new label. The characters in the label can be alphanumeric or spaces.
Strip Size/Full Stripe Size	<p>Strip size is the amount of data that is stored on each physical drive in the array. The full stripe size is the amount of data that the controller can read or write simultaneously on all the drives in the array. For RAID levels that support fault tolerance through parity, the parity information is calculated one full strip size at a time.</p> <p>You can specify from 8KiB to 1024KiB, depending on the number of disks and RAID level. The default value is all available space.</p> <p>If you use Smart Array S100i SW RAID, the minimum size is 16KiB, and the maximum size is 256KiB.</p>
Size	Values in decimal; minimum RAID size is 16 MiB.
Unit Size	Logical drive unit size (MiB/GiB/TiB).
Acceleration Method	Logical drive acceleration method (controller cache or none).

4. Click **Submit Changes**.

Assigning spare drives

A spare is a drive that automatically replaces a failed drive in a logical drive.



Prerequisites

A spare drive must meet the following criteria.

- It must be an unassigned drive or a spare drive for another array.
- It must be the same type as existing drives in the array (for example, SATA or SAS).
- The drive capacity must be greater than or equal to the smallest drive in the array.

Procedure

1. From the System Utilities screen, select **System Configuration > controller > Array Configuration > Manage Arrays > array > Manage Spare Drives**.
2. In the Manage Spare Drives screen, select the spare activation type:
 - **Assign Dedicated Spare**
 - **Assign Auto Replace Spare**
3. Select the drive that you want to assign as a spare.

NOTE: Only drives that meet the criteria listed in the prerequisites are displayed.

Deleting a spare drive

Procedure

1. From the System Utilities screen, select **System Configuration > controller > Array Configuration > Manage Arrays > array > Manage Spare Drives > Delete Spare Drives**.
2. From the Delete Spare Drives screen, select the spare that you want to delete, and click **Delete Spare Drives**.

Identifying a device

Use the UEFI System Utilities to identify a drive by turning on its device identification LED.

Procedure

1. From the System Utilities screen, select **System Configuration > controller > Array Configuration > Manage Arrays > array > Identify Device**.
2. In the Identify Device screen, specify the duration (in seconds) that you want the LED to be on, select the drive configuration type, and click **Start**.

To turn off the LED, click **Stop**.

Deleting an array

This procedure deletes:

- All the logical drives on the array.
- All data on the logical drives that are part of the array.



If the deleted array is the only one on the controller, the controller settings are erased, and the default configuration is restored.

To delete an individual logical drive, see "Deleting a logical drive."

Procedure

1. From the System Utilities screen, select **System Configuration > controller > Array Configuration > Manage Arrays > array > Delete Array**.
2. In the Delete Array screen, click **Submit Changes**.

More information

[Deleting a logical drive](#)

Editing a logical drive

Procedure


1. From the System Utilities screen, select **System Configuration > controller > Array Configuration > Manage Arrays > array > List Logical Drives > logical drive > Edit Logical Drive**.
2. In the Edit Logical Drive screen, edit any of the following settings.

Setting	Description
Acceleration method	Acceleration method can increase database performance by writing data to the cache memory instead of directly to the logical drives. Options are: <ul style="list-style-type: none">• Controller cache--writes data to the cache memory.• None--disables caching to reserve the cache module for other logical drives on the array.
Logical drive label	This label value appears in the Logical Drive Details screen. The label can contain alphanumeric characters and spaces only.

3. Click **Submit Changes**.

Deleting a logical drive

Use this procedure to delete an individual logical drive. To delete all logical drives in an array, see "Deleting an array."

 **IMPORTANT:** If you delete the logical drive, any data on the logical drive is deleted as well. If the logical drive that you are deleting is the only logical drive in the array, the array is also deleted.

Procedure

1. From the System Utilities screen, select **System Configuration > controller > Array Configuration > Manage Arrays > array > List Logical Drives > logical drive > Delete Logical Drive**.
2. In the Delete Logical Drive screen, click **Submit Changes**.



More information

[Deleting an array](#)

Disk Utilities

Viewing disk device information

Procedure

1. From the System Utilities screen, select **System Configuration** > **controller** > **Disk Utilities** > **disk** > **Device Information**.
2. In the Device Information screen, view the information.

Identifying a disk device

Procedure

1. From the System Utilities screen, select **System Configuration** > **controller** > **Disk Utilities** > **disk** > **Identify Device**.
2. In the Identify Device screen, specify the duration (in seconds) that you want the LED to be on, select the drive configuration type, and then click **Start**.

To stop blinking the LED, click **Stop**.

Setting bootable devices for Legacy Boot Mode

NOTE: HPE Smart Array S100i SR Gen10 SW RAID does not support legacy boot mode.

Setting the primary and secondary bootable devices (Legacy Boot Mode)

Procedure

1. From the System Utilities screen, select **System Configuration** > **controller** > **Set Bootable Device(s) for Legacy Boot Mode** > **Select Bootable Logical Drive** > **logical drive**.
2. In the *Logical Drive* screen, select either of the following:
 - **Set as Primary Bootable Device**
 - **Set as Secondary Bootable Device**



Setting the number of OS bootable drives (Legacy Boot Mode)

Procedure

1. From the System Utilities screen, select **System Configuration** > **controller** > **Set Bootable Device(s) for Legacy Boot Mode** > **Number of OS bootable drives**.
2. In the Number of OS bootable drives screen, specify the number of OS bootable drives.
3. Click **Submit Changes**.



HPE SSA graphical user interface

Using the HPE SSA GUI

Access HPE SSA with one of the many methods available:

- **Accessing HPE SSA in the offline environment**
- **Accessing HPE SSA in the online environment**

When you launch the HPE SSA GUI, the application opens and HPE SSA scans the system and detects controllers. This process can last up to 2 minutes. When controller detection is complete, the controllers are available on the **Devices/Tools** menu.

After the GUI is open, tasks are distributed among categories. For more information, see "**Navigating the GUI.**"

Icons

The HPE SSA GUI includes many icons (also defined in the Help file) to help with identification and troubleshooting.














Image	Description
	Critical
	Warning
	Informational
	Active Task(s)
	Paused/Offline Drive
	HPE ProLiant Server
	Array Controller
	Array Controller (Embedded)
	Array/Logical Devices or Solid State Devices
	Logical Drive
	Assigned Physical Drive
	Unassigned Physical Drive
	Unassigned Drives

Table Continued



Image	Description
	Spare Drive
	Transient Drive
	Storage Enclosure
	Port or Unassigned Physical Drive Connected to Boot Port
	Tape Drive
	Locked
	License Manager/Encryption Manager
	Cache Manager
	Array Diagnostic Report
	Smart SSD Wear Gauge Report
	None

Keyboard shortcuts

Keyboard functions and shortcuts can be used for navigating or performing actions in the GUI.

Key	Description
Tab	Cycle through selectable items on a page.
Shift + Tab	Cycle backwards through selectable items on a page.
F5	Rescan system (equivalent to clicking the Rescan System button).
B	Browse main menu.
H	Open HPE SSA Help.
X	Exit HPE SSA.
Enter	Perform the action of the currently selected link or button.
Escape	Close nonaction popups. ¹
R	Refresh selected controller. ¹
Spacebar (toggle)	Select or clear a check box selection. ¹



¹ Local keyboard shortcuts are available only when the action that the key activates is accessible.

Navigating the GUI

When you open HPE SSA, the **Welcome** screen appears.



The following elements are visible:

- The **HPE Smart Storage Administrator** quick navigation menu is in the top, left-hand corner of the screen. Clicking the down arrow displays the available devices, and clicking one of the available devices displays additional information and options for the device. You can also return to a server **Home** screen, or you can choose **Configuration** or **Diagnostics** for a device listed. For more information, see "[Configure screen](#)" or "[Diagnostics screen](#)."



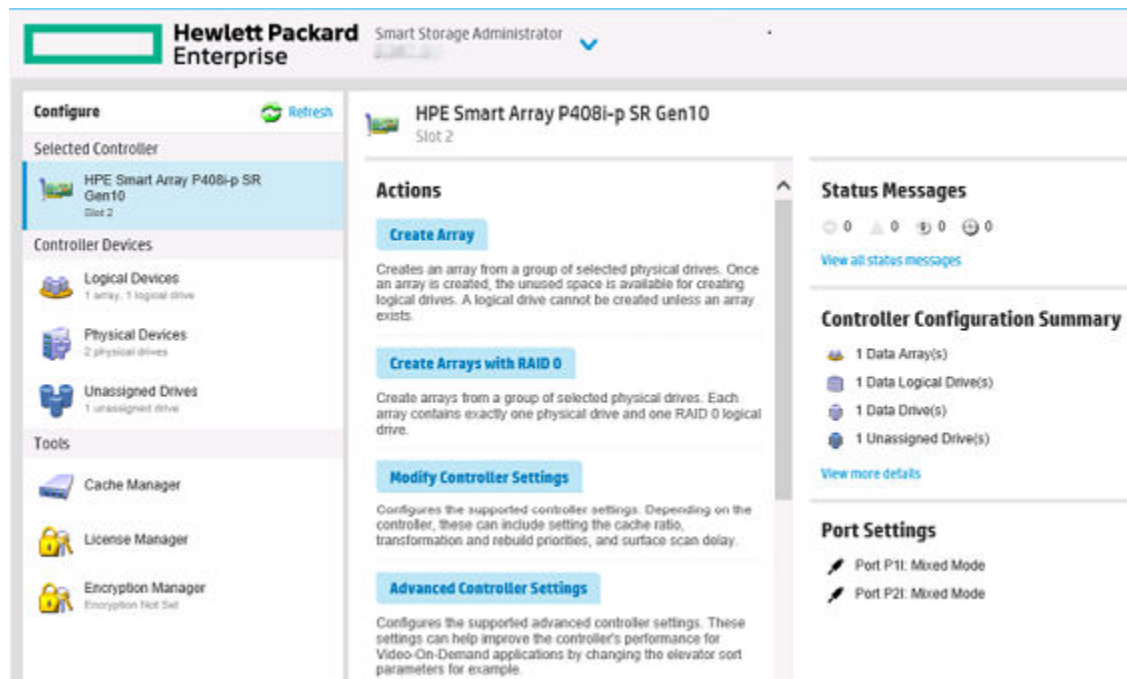
- **Available device(s)** are listed on the left-hand side of the screen. Clicking on a server or array controller displays the available actions, alerts, and summary for that device. You can point to the status alerts to see details on an alert.
- The **Rescan System and Reset Application** button is near the top right of the screen.
After adding or removing devices, click **Rescan System and Reset Application** to update the list of available devices.
- The **Help** button is near the top right of the screen.
To access help topics, press the **H** key or click **Help**. For more information, see "[HPE SSA Help](#)."
- The **Exit** button is near the top right of the screen.



Configure screen

To access this screen, click either a device under **Configuration** in the quick navigation menu, or select an available device from the **Home** screen, and then click **Configure** under the available options.

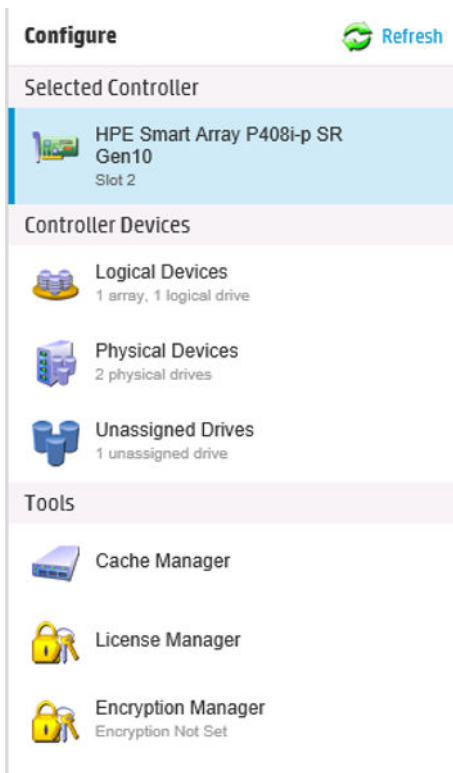
The **Configure** screen displays the GUI elements from the **Welcome** screen and lists available actions, status messages, more detailed information, and a controller configuration summary for a selected controller.



When a controller is selected, the following elements appear:

- **Controller Devices and Tools**—This panel, at left, displays systems, controllers, arrays, physical devices, unassigned drives, cache managers, and license managers. The panel also displays encryption managers.





- **Actions**—This panel, in the middle, provides the following information and functionality:
 - Tasks that are available for the selected device based on its current status and configuration
 - Options and information pertinent to the task, after a task is selected
- **Status Messages**—This panel provides the following information and functionality:
 - Status icons (critical, warning, and informational) with the number of individual alerts for each category
 - A **view all status messages** link that displays device-specific alerts in a pop-up window
- **Controller Configuration Summary**—This panel provides a summary of the following elements for the selected controller:
 - Data arrays
 - Data logical drives
 - Data drives
 - Unassigned drives
 - Spare drives
 - A **View more details** link that displays more information in a pop-up window
- **Port settings**—This panel provides more information about the drives attached to the controller. The values vary according to the controller. For Gen10 controllers, mixed mode combines RAID and HBA modes, occurs by default, and cannot be disabled.

For a list of possible tasks that are available on the **Configure** screen, see "[Configuration tasks](#)."



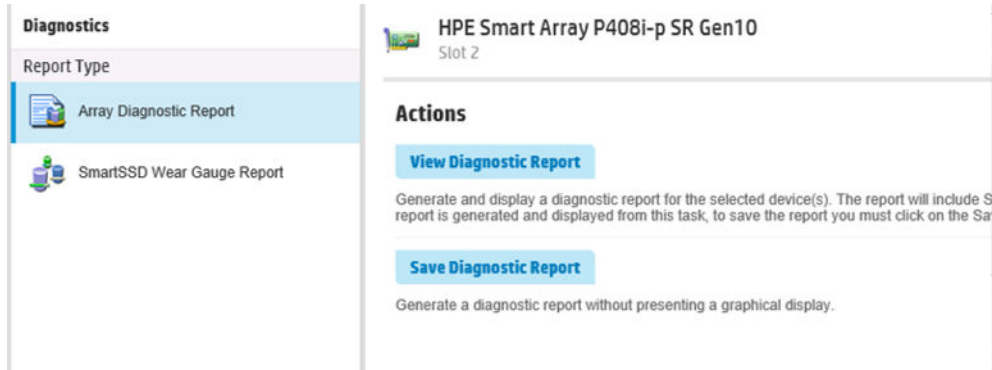
Diagnostics screen

To access this screen, either click a device under **Diagnostics** in the quick navigation menu, or select an available device from the **Home** screen, and then click **Diagnose** under the available options.

From the **Diagnostics** screen, you can run one of the following reports:

- Diagnostics Report
- SmartSSD Wear Gauge Report

When selecting either report, the available actions on the **Actions** panel include viewing the report or saving the report.



For a list of possible tasks that are available on the **Diagnostics** screen, see "[Diagnostic tasks](#)."

SmartSSD Wear Gauge Report

The SmartSSD Wear Gauge Summary provides a brief overview of the following:

- Total Solid State Drives with Wearout Status
- Total Smart Array Solid State Drives
- Total Non Smart Array Solid State Drives
- Total Solid State SAS Drives
- Total Solid State NVMe Drives
- Total Solid State Drives

When running a report, you can either view a graphic representation of the report with SSD usage and estimated lifetime information, or generate a report without a graphical display, with the option of saving the report.

HPE SSA help

The Help button, at upper right, opens the embedded HPE SSA help file. In addition to providing information about the main screens and tabs, Help also provides several useful topics for new users, including the following:

- **Image Legend**—A visual reference list defining the icons and graphical buttons used in HPE SSA
- **Keyboard Shortcuts**—A list of keys and operations they perform within the GUI

To view these help topics and others, press the **H** key or click **Help**. When the Help window opens, expand the topic "Getting Started with Smart Storage Administrator Overview."

The glossary in HPE SSA help defines industry standard and Hewlett Packard Enterprise terms as they relate to the HPE SSA application.



Configure the language of the GUI

Configuring the language of the GUI (Windows)

Procedure

1. On the local server (host), click **Start**, and then select **Programs>Windows System>Smart Storage Administrator Preferences**.

The **Smart Storage Administrator Preferences** screen appears.

2. Under Preferred Language, select a language or leave as **System Default**.
3. Click **Exit**.

Configuring the language of the GUI (UNIX)

Procedure

1. Identify the two-digit code for the language that you want by running the command ssa-h:

```
# ssa -h
Smart Storage Administrator 3.30.13.0 2018-04-12
Command Line Functions:
-local : Run ssa application
ssa -local
Run the application in English.
ssa -local [ -lang languageCode ]
Run the application in the language specified by languageCode.
Supported languageCode/languages are:
en - English (default)
ja - Japanese
de - German
es - Spanish
fr - French
it - Italian
pt - Portuguese
ru - Russian
zh - Simplified Chinese
```

2. Specify the language using the appropriate two-digit language code:
`ssa -local -lang <languageCode>`


Configuration tasks

From the **Configure** screen, you can perform tasks related to controllers, arrays, physical drives, and logical drives.



When a controller or device is selected, the tasks that appear are a subset of the total number of possible tasks for the selected item. HPE SSA lists or omits tasks based on the controller model and configuration. For example, if the selected controller has no unassigned physical drives, Create Array is not an available task.

The following table lists all the possible tasks for every type of item.

 **IMPORTANT: HPE Special Reminder:** Before enabling encryption on the Smart Array controller module on this system, you must ensure that your intended use of the encryption complies with relevant local laws, regulations and policies, and approvals or licenses must be obtained if applicable.

For any compliance issues arising from your operation/usage of encryption within the Smart Array controller module which violates the above mentioned requirement, you shall bear all the liabilities wholly and solely. HPE will not be responsible for any related liabilities.



Item	Tasks
Controller	Accelerated I/O Path
	Advanced Controller Settings ¹
	Array Accelerator Settings
	Check Online Firmware Activation Readiness
	Clear Configuration
	Controller Settings
	Create Array
	Disable Standby Controller
	Enable HBA/RAID/Smart Array Mode operations ¹
	Enable SmartCache
	Manage Encryption
	Manage License Keys ¹
	Modify Power Modes ¹
	More Information
	Parallel Surface Scan
	Physical Drive Write Cache Settings
	Redundancy Settings ¹
	Set Sanitize Lock
	View Status Alerts
Array	Bypass RAID components using HPE SSA Smart Path
	Change Array Drive Type
	Create Array
	Create Logical Drive
	Create Split Mirror Backup
	Convert Plaintext Data to Encrypted Data
	Delete
	Expand Array
	Heal Array
	Manage Split Mirror Backup
	More Information

Table Continued

Item	Tasks
	Move Drives Re-Mirror Array Shrink Array Spare Management Split Mirrored Array View Status Alerts Volume Key Rekey
Logical drive	Create Logical Drive Create SmartCache for Logical Drive Convert Plaintext Data to Encrypted Data Delete Extend Logical Drive Instant Secure Erase Migrate RAID/Stripe Size Modifying Cache Write Policy Move Logical Drive ¹ More Information Re-enable Failed Logical Drive View Status Alerts Volume Key Rekey
Unused space	Create Logical Drive More Information
Physical drive	Erase Drive View Status Alerts
Unassigned drive	Create Array Erase Drive More Information

¹ This task is not available on all controller models.



Performing a Configuration task

Procedure

1. Open HPE SSA.

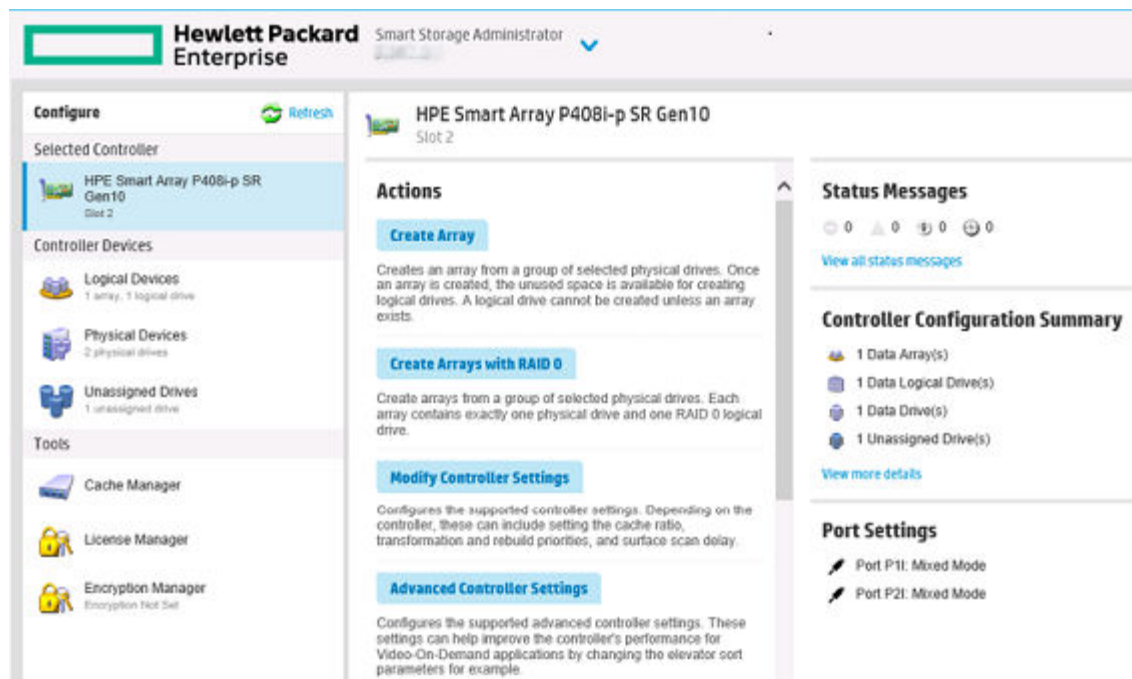
For more information, see "[Using the HPE SSA GUI](#)."

2. Open the **Configure** panel by doing one of the following:

- Choose a device and click **Configure** in the quick navigation menu.
- Select an available device from the **Home** screen, and then click **Configure** under the available options.

3. Select a device from the **Devices** menu.

The Actions, Status Messages, and Controller Configuration Summary panels appear. The listed tasks are available for this device in its current configuration. For more information, see "[Configuration tasks](#)."



4. Click a task button.

A list of all possible options for that task appears on the right side of the screen, replacing the task list.

5. Select the settings or configuration options for the device.
6. Use the Next and Back buttons to navigate multiple screens of options.
7. Click **Save** or **OK**.

Creating an array

Use this procedure to configure RAID and other settings for the array/logical drive.

Procedure

1. Open HPE SSA.



For more information, see "[Using the HPE SSA GUI](#)."

2. Open the **Configure** panel by doing one of the following:
 - Choose a device and click **Configure** in the quick navigation menu.
 - Select an available device from the **Home** screen, and then click **Configure** under the available options.
3. Select a controller from the **Devices** menu.

The **Actions** panel appears.
4. Click **Create Array**.
5. Select the physical drives for the new array.
6. Click **Create Array**.
7. Make selections for **RAID Level**, **Strip Size/Full Stripe Size**, **Sectors/Track**, and **Size**.
8. Click **Create Logical Drive**.
9. Click **Finish**.

By default, HPE SSD Smart Path is enabled on SSD drives.


Configuring a controller

Procedure

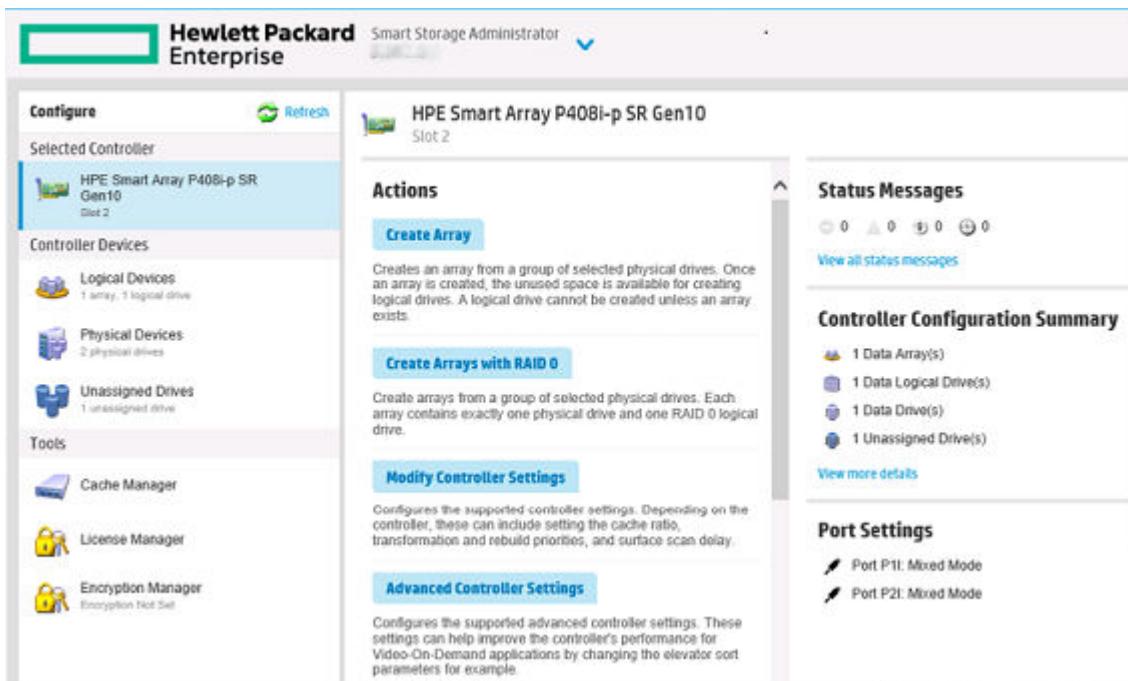
1. Open HPE SSA.

For more information, see "[Using the HPE SSA GUI](#)."
2. Open the **Configure** panel by doing one of the following:
 - Choose a device and click **Configure** in the quick navigation menu.
 - Select an available device from the **Home** screen, and then click **Configure** under the available options.

The **Configure** panel appears.

 **IMPORTANT:** Screens may have different options available, depending on the server configuration.





3. Configure the controller. See "**Performing a Configuration task.**"
4. When prompted, save the configuration.
5. Do one of the following:
 - Configure an additional controller. Repeat steps 3 through 5.
 - Click **Exit HPE SSA**.

Enabling or disabling HPE SSD Smart Path

HPE SSD Smart Path is enabled by default when you create an array. For more information, see "**Creating an array.**"

Procedure

1. Open HPE SSA.
For more information, see "**Using the HPE SSA GUI.**"
2. Open the Configure panel by doing one of the following:
 - Choose a device and click **Configure** in the quick navigation menu.
 - Select an available device from the Home screen, and then click **Configure** under the available options.
3. Select a controller from the Devices menu.
The **Actions** panel appears.
4. Under **Configure**>**Logical Devices**, choose the array with logical drives.
5. Click **Disable HPE SSD Smart Path**.
6. Click **Save**.



Setting sanitize lock

The sanitize lock feature enables or disables sanitize operations for physical SATA drives. This setting is enabled at the controller level.

NOTE: This feature is applicable to SATA drives only and is not applicable to the Smart Array S100i.

Procedure

1. Open HPE SSA.

For more information, see "[Using the HPE SSA GUI](#)."

2. Open the **Configure** panel by doing one of the following:

- Choose a device and click **Configure** in the quick navigation menu.
- Select an available device from the **Home** screen, and then click **Configure** under the available options.

3. Select a controller from the **Devices** menu.

The **Actions** panel appears.

4. In the **Actions** panel, click **Set Sanitize Lock**.

5. In the Set Sanitize Lock panel, select one of the following options:

- **None:** This state is the normal state of the physical disk. No freeze or antifreeze commands are sent to any drives.
- **Freeze:** This setting prevents a drive sanitize operation.
- **Anti-freeze:** This setting prevents physical disks from being frozen. The setting enables drive sanitize operations.

6. If you are modifying the settings, reboot the server.

7. If you are enabling the freeze or antifreeze setting for the first time or modifying the settings, you might need to power cycle or hot-plug the drives.

Rapid parity initialization

RAID levels that use parity (RAID 5, RAID 6, RAID 50, and RAID 60) require that the parity blocks be initialized to valid values. Valid parity data is required to enable enhanced data protection through background controller surface scan analysis and higher write performance (backed out write). After parity initialization is complete, writes to a RAID 5 or RAID 6 logical drive are typically faster because the controller does not read the entire stripe (regenerative write) to update the parity data.

The rapid parity initialization method works by overwriting both the data and parity blocks in the foreground. The logical drive remains invisible and unavailable to the operating system until the parity initialization process completes. Keeping the logical volume offline eliminates the possibility of I/O activity, thus speeding the initialization process, and enabling other high-performance initialization techniques that wouldn't be possible if the volume was available for I/O. Once the parity is complete, the volume is brought online and becomes available to the operating system

This method has the following benefits:

- It speeds up the parity initialization process.
- It ensures that parity volumes use backed-out writes for optimized random write performance.



Initializing rapid parity

Procedure

1. Open HPE SSA.

For more information, see "[Using the HPE SSA GUI](#)."

2. Open the **Configure** panel by doing one of the following:

- Choose a device and click **Configure** in the quick navigation menu.
- Select an available device from the **Home** screen, and then click **Configure** under the available options.

3. Select **Logical Devices** from the **Devices** menu.

A list of arrays appears.

4. Select an array, and select **Create Logical Drive** from the **Actions** menu.

5. Make selections for **RAID Level**, **Number of Parity Groups**, **Strip Size/Full Strip Size**, **Sectors/Track**, **Size**, **Parity Initialization Method**, and **Caching**

Depending on the controller used and drives selected, your options may appear different from the image below.



RAID Level (What's this...?)

☐ RAID 0
☐ RAID 1+0
☐ RAID 5
☒ RAID 50

Number of Parity Groups (NPG) (What's this...?)

2 ▾

Strip Size / Full Stripe Size (What's this...?)

☐ 8 KiB / 16 KiB
☐ 16 KiB / 32 KiB
☐ 32 KiB / 64 KiB
☐ 64 KiB / 128 KiB
☐ 128 KiB / 256 KiB
☒ 256 KiB / 512 KiB
☐ 512 KiB / 1024 KiB
☐ 1024 KiB / 2 MiB

Sectors/Track (What's this...?)

☐ 63
☒ 32

Size (What's this...?)

☒ Maximum Size: 381416 MiB (372.4 GiB)
☐ Custom Size

Parity Initialization Method (What's this...?)

☒ Default: Online, parity block initialization
☐ Rapid: Offline, full zero-overwrite of all data and parity blocks

Caching (What's this...?)

6. Click **Create Logical Drive** to continue.

A summary page appears.

7. Click **Finish**.

Changing the Spare Activation Mode

The spare activation mode feature enables the controller firmware to activate a spare drive under the following conditions:

- When a data drive reports a predictive failure status
- When a data drive fails; this mode is the default.

In normal operations, and for older controllers, the firmware starts rebuilding a spare drive only when a data drive fails. With the predictive failure activation mode, rebuilding can begin before the drive fails, reducing the likelihood of data loss that could occur if an additional drive fails.



Procedure

1. Open HPE SSA.

For more information, see "[Using the HPE SSA GUI](#)."

2. Open the **Configure** panel by doing one of the following:

- Choose a device and click **Configure** in the quick navigation menu.
- Select an available device from the **Home** screen, and then click **Configure** under the available options.

3. Select a controller from the **Devices** menu.

The **Actions** panel appears.

4. In the **Actions** panel, click **Modify Spare Activation Mode**.

5. From the menu, select one of the following modes:

- Failure Spare Activation
- Predictive Spare Activation

6. Click **Save**.

Changing the Spare Management mode

The Spare Management feature provides multiple methods for handling spare behavior. You can choose from the following options:

- **Dedicated**—When the failed data drive is replaced, it must be rebuilt from the data on the spare drive. In Dedicated mode, one spare can be dedicated to multiple arrays.
- **Auto-Replace Drives**—The spare for the failed data drive automatically becomes the replacement data drive. When the spare is replaced, the data drive does not need to be rebuilt. In Auto-replace mode, spare drives cannot be shared between arrays.

If assigning **Auto-Replace Drives** mode to an array with a RAID 0 drive, **Spare Activation Mode** must be set to **Predictive Spare Activation** mode.

Procedure

1. Open HPE SSA.

For more information, see "[Using the HPE SSA GUI](#)."

2. Open the **Configure** panel by doing one of the following:

- Choose a device and click **Configure** in the quick navigation menu.
- Select an available device from the **Home** screen, and then click **Configure** under the available options.

3. Select a controller from the **Devices** menu.

4. Select **Create Array** from the **Actions** panel.

The Array details screen appears.

5. Select a controller, drive type, and physical drives, and click **Create Array**.



6. Select settings for **Create Plaintext Volume**, **RAID Level**, **Strip Size/Full Strip Size**, **Sectors/Track**, **Size**, and **Caching**. Click **Create Logical Drive** when complete.
7. Click **Manage Spare Drives**.
8. From the menu, select one of the following spare drive types:
 - Dedicated Spare Drives
 - Auto-Replace Drives
9. Select which drives will operate as spare drives in the array.
10. Click **Save**.

A confirmation screen appears.
11. Click **Yes** to continue.
12. Click **Manage Spare Drives** to make additional selections, or click **Finish**.

Disabling Over Provisioning Optimization

Solid state devices can be optimized by deallocating all used blocks before any data is written to the drive. The optimization process is performed when the first logical drive in an array is created and when a physical drive is used to replace a failed drive. Not all controllers support this option.

The SSD Over Provisioning Optimization feature allows the user to disable Over Provisioning Optimization in the GUI.

Procedure

1. Open HPE SSA.

For more information, see "[Using the HPE SSA GUI](#)."
2. Select **Unassigned Drives** located under **Controller Devices**.
3. To create an array, select from the list of available drives listed. When finished, click **Create Array**.
4. The **Create Logical Drive** window appears. Under the option **SSD Over Provisioning Optimization**, select **Do not perform SSD Over Provisioning Optimization on the Array**.
5. Click **Create Logical Drive**.

Setting Auto RAID 0

CAUTION: If you select this option for any logical drives, you will experience data loss for that logical drive if one physical drive fails. Assign RAID 0 to drives that require large capacity and high speed, but pose no data safety risk.

Auto Array RAID 0 creates a single RAID 0 volume on each physical drive specified, enabling the user to select multiple drives and configure as RAID 0 simultaneously. Each array contains one physical drive and one RAID 0 logical drive.

For more information about RAID 0, see "[RAID 0—No fault tolerance](#)."

Procedure

1. Open HPE SSA.



For more information, see "[Using the HPE SSA GUI](#)."

2. Select the controller.
3. Click **Create Arrays with RAID 0**.
4. A new window appears, confirming each array will contain a single RAID 0 logical drive. Click **Yes** to continue.
5. A new window appears, confirming RAID 0 configuration. Click **Finish** to complete.

Modifying parallel surface scan

Procedure

1. Open HPE SSA. For more information, see "[Using the HPE SSA GUI](#)."
2. Select a controller.
3. Click **Modify Controller Settings**.
4. Under **Current Parallel Surface Scan Count**, use the slider to select the parallel surface scan count.
5. Click **Save Settings**.

A summary page appears.

6. Click **Finish** to exit.

Configuring controller cache

Caching increases database performance by writing data to the cache memory, instead of directly to the logical drives. Caching can be disabled to reserve the cache module for other logical drives on the array.

Procedure

1. Open HPE SSA.
For more information, see "[Using the HPE SSA GUI](#)."
2. Open the **Configure** panel by doing one of the following:
 - Choose a device and click **Configure** in the quick navigation menu.
 - Select an available device from the **Home** screen, and then click **Configure** under the available options.
3. Select **Cache Manager** from the **Tools** menu.
4. Click **Controller Cache** in the **Cache Manager** menu.
5. Click **Modify Caching Settings**.
6. Select one or more logical drives to be cached.
7. Verify caching settings.
8. Click **OK**.



HPE Smart Array SR SmartCache

SmartCache enables solid-state drives to be used as caching devices for hard drive media. SmartCache provides the following features:

- Accelerates application performance.
- Provides lower latency for transactions in applications.
- Supports all operating systems without the need for changes to OS, driver, or applications.

SmartCache will be fully enabled after the first SmartCache is created on the controller.

The following features are not available unless SmartCache is disabled:

- Expand Array
- Advanced Capacity Expansion
- Move Logical Drive
- Change Array Drive Type
- Mirror Splitting and Recombining (offline only)
- Split Mirror Backup and Rollback (online and offline)
- Heal Array
- Extend Logical Drive
- Migrate RAID/Strip Size
- Change Cache Ratio
- Align Logical Drives

SmartCache requires an energy pack and a SmartCache license. If you purchase the HPE Smart Array P816i-a SR Gen10, a license is included.

For more information, see the [Hewlett Packard Enterprise website](#).

Enable HPE Smart Array SR SmartCache

Enabling SmartCache with an array will accelerate data input/output for the assigned logical drives. At least one logical drive must be created on the controller prior to enabling SmartCache.

Procedure

1. Open HPE SSA.
For more information, see "[Using the HPE SSA GUI](#)."
2. Open the **Configure** panel by doing one of the following:
 - Choose a device and click **Configure** in the quick navigation menu.
 - Select an available device from the **Home** screen, and then click **Configure** under the available options.
3. Select **Cache Manager** from the **Tools** menu.
4. Click **Enable HPE SmartCache** in the **Actions** menu.



5. Select one or more physical drives from the list of available drives.
6. Click **OK**.
7. A pop-up window appears, indicating certain features will not be available with SmartCache enabled. If you want to continue, click **Yes**.
8. The dialog **Create SmartCache for Logical Drive** appears.
9. Select the following:
 - Logical drive to cache
 - Cache write policy and RAID type
 - Size of the cache. Hewlett Packard Enterprise recommends 10% of the drive size, and it must be a minimum of 16 GiB.
10. Click **Create SmartCache for Logical Drive**.
11. **SmartCache Details**, **SmartCache Statistics**, and **Device Path** appear. Click **Finish** to proceed.

The SmartCache is created for the Logical Drive.

Installing a license key with HPE SmartCache

You can use HPE SSA to install the license key and activate SmartCache functionality.

To install a license key:

Procedure

1. Open HPE SSA.

For more information, see "[Using the HPE SSA GUI](#)."
2. Open the **Configure** panel by doing one of the following:
 - Choose a device and click **Configure** in the quick navigation menu.
 - Select an available device from the **Home** screen, and then click **Configure** under the available options.
3. In the **Tools** menu, click **License Manager**.
4. Under actions, click **Add License Key**.
5. Enter the license key number.
6. Click **Save**.

Caching write policies

SmartCache Write-back

HPE SSA contains two different policies for data writes when caching: write-back and write-through.

Write-back is a caching method where data is not copied to the data volume until absolutely necessary. Write-back might accelerate performance in comparison to the write-through policy by reducing the number of write operations to data volumes. The performance improvement can involve the risk of lost data if the cache volume fails.





Write-through is a caching method where data is written to the cache and the data volumes simultaneously. Write-through is the preferred write policy in applications where data loss cannot be tolerated, but has lower performance compared to the write-back policy.

A write-back SmartCache cannot be deleted until it is converted to a write-through SmartCache using the **Modify Caching Write Policy** button. This conversion forces user data to be flushed from the SmartCache volume to the primary hard drive volume, to avoid data loss when the SmartCache is deleted. The time to flush data from write-back SmartCache to hard drive volume varies, depending on several variables including how much dirty data is held in the write-back SmartCache, host workload, and number of hard drives in the primary volume.

SmartCache volumes must be deleted from newest to oldest, in reverse order from how they were created.

Some controllers might not support this option or might require a license key to enable this feature. The target can be any valid SSD drive and existing non-cached logical drive for the data.

 **CAUTION:** Specifying the write-back cache write policy could result in data loss in the event of failure of the cache volume. When using a RAID 0 cache volume, a single SSD failure might result in data loss.

 **IMPORTANT:** If a demo license key expires, all SmartCache volumes configured with write-back cache write policy are converted to write-through. When this happens, the logical drive details show different values for Cache Write Policy and Cache Write Policy Requested. If the license is reinstalled, the SmartCache volumes are restored to their original write-back cache write policy.

Modifying the physical drive write cache policy setting

The steps that you follow vary according to whether your controller supports modifying the policy on configured drives or on configured *and* unconfigured drives.

If your controller supports drive write caching of configured drives only

The Physical Drive Write Cache State setting appears in the Modify Controller Settings panel if your controller supports drive write caching of configured drives only.

Procedure

1. Open HPE SSA.

For more information, see "[Using the HPE SSA GUI](#)."

2. Select the controller.
3. Click **Modify Controller Settings**.
4. Under **Physical Drive Write Cache State**, select one of the following:
 - Enabled
 - Disabled
5. Click **Save Settings**.
6. A summary page appears. Click **Finish** to exit.

If your controller supports drive write caching of configured and unconfigured drives

Prerequisites

The Manage Drive Write Cache Policy option appears on the controller Actions menu if your controller supports drive write caching of configured and unconfigured drives.



Procedure

1. Open HPE SSA
2. Select your controller from the Configure menu
3. In the Actions menu, select **Manage Drive Write Cache Policy**.
4. In the Manage Drive Write Cache Policy dialog, select one of the following for configured drives or unconfigured drives:
 - Default -- selecting this option for configured drives allows the controller to optimize the drive write cache policy of the drives; selecting this option for unconfigured drives uses the drive's existing write cache policy.
 - Enable -- selecting this option can increase write performance but risks losing the data in the cache on sudden power loss.
 - Disable
5. Click **OK**.

Modifying the write cache bypass threshold

All writes larger than the specified value will bypass the write cache and be written directly to the disk for nonparity RAID volumes. A smaller value allows the controller to reserve write caching to I/Os smaller than the threshold.

Procedure

1. Open HPE SSA.
For more information, see "[Using the HPE SSA GUI](#)."
2. Open the **Configure** panel by doing one of the following:
 - Choose a device and click **Configure** in the quick navigation menu.
 - Select an available device from the **Home** screen, and then click **Configure** under the available options.
3. Select **Cache Manager** from the **Tools** menu.
4. Click **Controller Cache** in the **Cache Manager** menu.
5. Click **Modify Caching Settings**.
6. Select one or more drives.
7. In the **Write Cache Bypass Threshold** section, use the slider bar to select the value, or enter the value into the field.
A valid value is from 16 KiB to 1024 KiB. If the value that you enter is not a multiple of 16, the user interface will update the value to the nearest valid value.
8. Click **OK**.

Working with mirrored arrays

Among the advanced tasks possible with the HPE SSA GUI, you can split a mirrored array and then recombine it. This process entails splitting a RAID 1 or RAID 1+0 mirror into two identical new arrays consisting of RAID 0 logical drives.

Support for these procedures requires the following:



- The HPE SSA GUI must be run in **offline mode**.
- Mirrored arrays being split can have RAID 1, RAID 1+0, RAID 1 (ADM), or RAID 10 (ADM) configurations. Arrays with other RAID configurations cannot be split.

Several reasons exist for splitting and recombining a mirrored array. For more information, see the "*RAID 1(+0): breaking mirrors and rebuilding drives*" how-to white paper on the [Hewlett Packard Enterprise website](#).

Splitting a mirrored array

Procedure

1. Run the HPE SSA GUI in offline mode. See "[Accessing HPE SSA in the offline environment](#)."
2. Open the **Configure** panel by doing one of the following:
 - Choose a device and click **Configure** in the quick navigation menu.
 - Select an available device from the **Home** screen, and then click **Configure** under the available options.
3. Select **Arrays** from the **Devices** menu.
4. From the **Arrays** menu, select the appropriate array.
5. Select **Manage Data Drives** from the **Actions** panel.
6. Under **Available Array Action(s)**, select **Mirror Array**.
7. Select a physical drive.
8. Click **OK**.
The mirrored array details are displayed.
9. Click **Finish**.
10. When HPE SSA finishes splitting the array, two logical drives appear in the **Arrays** menu:
 - When a RAID 1 or RAID 1+0 array splits, two RAID 0 logical drives are created.
 - When an array that contains a RAID 1 (ADM) logical drive splits, a RAID 1 logical drive and a RAID 0 logical drive are created.
 - When an array that contains a RAID 10 (ADM) logical drive splits, a RAID 1+0 logical drive and a RAID 0 logical drive are created.
11. Shut down the OS.
12. Power down the server.
13. With power off, remove the physical drives that constitute one of the new arrays.
If you do not remove the physical drives for one of the arrays, the OS will be unable to distinguish between the two arrays when the server is restarted because the arrays are identical.
14. Power up the server.
15. Restart the OS.



Recombining a split mirrored array

Procedure

1. Run the HPE SSA GUI in offline mode. See "[Accessing HPE SSA in the offline environment](#)."
2. Open the **Configure** panel by doing one of the following:
 - Choose a device and click **Configure** in the quick navigation menu.
 - Select an available device from the **Home** screen, and then click **Configure** under the available options.
3. Select **Arrays** from the **Devices** menu.
4. Select the appropriate array from the **Arrays** menu.
5. Select **Manage Data Drives** from the **Actions** panel.
6. In the **Available Tasks** panel, click **Manage Split Mirror Backup**.
7. Select the array to be mirrored to the source array.

This array is usually the array that was split from the original mirrored array. However, it can be any other array of the correct size.
8. Click **OK**.
9. When HPE SSA finishes re-mirroring the array, restart the OS.

The controller uses the rebuild process to synchronize the mirrored drives. The drive online LED flashes during the rebuild process. Depending on the hard drive size and the server load, this process can take up to 2 hours. You can boot the OS during this time, but the logical drive is not fault-tolerant until the rebuild is complete.

Creating a split mirror backup

This task splits an array that consists of one or more RAID 1, RAID 1+0, RAID 1 (ADM), or RAID 10 (ADM) logical drives, and then creates two arrays: a primary array and a backup array.

Procedure

1. Run the HPE SSA GUI in offline mode. See "[Accessing HPE SSA in the offline environment](#)."
2. Open the **Configure** panel by doing one of the following:
 - Choose a device and click **Configure** in the quick navigation menu.
 - Select an available device from the **Home** screen, and then click **Configure** under the available options.
3. Select **Arrays** from the **Devices** menu.
4. Select the appropriate array from the **Arrays** menu.
5. In the **Actions** panel, click **Create Split Mirror Backup**.

A verification and message dialog box appears.
6. Click **OK**.
7. A details window appears. Click **Finish**.

HPE SSA creates the array according to the following rules:



- If the original array contained RAID 1 or RAID 1+0 drives, then the primary array will contain RAID 0 drives.
 - If the original array contained RAID 1 (ADM) drives, the primary array will contain RAID 1 drives.
 - If the original array contained RAID 10 (ADM) drives, the primary array will contain RAID 1+0 drives.
 - The backup array always contains RAID 0 logical drives.
 - The primary array continues to be fully accessible to the operating system while the backup array is hidden from the operating system.
8. When HPE SSA finishes creating the split mirror backup, the new backup array appears in the **Devices** menu:
The array includes the designation "Backup" at the beginning of the array name.

Re-mirroring, rolling back, or re-activating a split mirror backup

Procedure

1. Run the HPE SSA GUI in offline mode. See "[Accessing HPE SSA in the offline environment](#)."
2. Open the **Configure** panel by doing one of the following:
 - Choose a device and click **Configure** in the quick navigation menu.
 - Select an available device from the **Home** screen, and then click **Configure** under the available options.
3. Select **Arrays** from the **Devices** menu.
4. Select the appropriate array from the **Arrays** menu.
5. In the **Available Tasks** panel, click **Manage Split Mirror Backup**.
6. Select one of the following actions:
 - Re-mirror the array and preserve the existing data. Discard the backup array.
This option re-creates the original mirrored array with the current contents of the primary array.
 - Re-mirror the array and roll back to the contents of the backup array. Discard existing data.
This option re-creates the mirrored array but restores its original contents, which are in the backup array. Hewlett Packard Enterprise recommends that you do not perform this option under the following circumstances:
 - In an online environment
 - If the logical drive to be rolled back is mounted
 - If the logical drive to be rolled back is in use by the operating system
 - Activate the backup array.
This option makes the backup array fully accessible to the operating system. HPE SSA removes the designation "backup" from the name of the array.

Heal array

The Heal Array operation allows you to replace failed physical drives in the array with healthy physical drives. The original array and logical drive numbering is unaffected after the replacement. Note the following conditions and restrictions for the Heal Array operation:



- The replacement physical drives and the original drives must be the same interface type (such as SAS or SATA) as the original drives.
- The operation is available only if enough unassigned physical drives of the correct size are available.
- The array has at least one failed drive.
- The array is not transforming (for example, rebuilding to a spare).
- The array has a working cache, making it capable of transformation.

Change array drive type

Some controllers may not support this option or might require a license key to enable the feature.

HPE SSA enables you to transfer the contents of an array to an existing empty array or a new array. During this operation, all logical drives transfer from the original array to the destination array. The original array is deleted, and the drives that were being used are freed and listed as unassigned drives.

Changing an array drive type is a time-consuming process for two reasons: all data in each logical drive is copied to the destination array, and the controller performs all data transformations while servicing IO requests to other logical drives.

To perform the operation, you must meet the following conditions:

- The destination array must have the same number of physical drives as the source or original array.
- Both the source and destination arrays must be in OK state.
- All existing logical drives in the source array must be in OK state.
- The destination array must have sufficient capacity to hold all logical drives present in the source array.

Encryption Manager

! **IMPORTANT: HPE Special Reminder:** Before enabling encryption on the Smart Array controller module on this system, you must ensure that your intended use of the encryption complies with relevant local laws, regulations and policies, and approvals or licenses must be obtained if applicable.

For any compliance issues arising from your operation/usage of encryption within the Smart Array controller module which violates the above mentioned requirement, you shall bear all the liabilities wholly and solely. HPE will not be responsible for any related liabilities.

Secure Encryption is a controller-based, enterprise-class data encryption solution that protects data at rest on bulk storage hard drives and SSDs attached to a compatible Smart Array Controller. The solution is compatible with the Enterprise Secure Key Manager and can operate with or without the presence of a key manager in the environment, depending on individual customer settings.

Secure Encryption is supported on Smart Array P-class and E-class Gen10 controllers.

For more information about performing a full setup and configuring Secure Encryption, see *HPE Smart Array SR Secure Encryption Installation and User Guide*.

Power modes

There are three available power modes:



- Maximum performance
- Minimum power
- Balanced

Maximum performance (default)

This is the default setting. All settings are selected based on maximum performance. Power savings options that affect performance are disabled.

Balanced

You can use this setting to save power with minimal effects on performance. For large queue depths, this setting affects throughput by 10% or less.

At lower queue depths or infrequent I/O, impacts on performance may be greater. This command is typically useful in environments using only hard drives, and is not recommended when using SSDs.

Settings are based on the user configuration, such as the number or types of drives, the RAID level, and storage topology. Significant changes to the configuration may require a reboot for optimal setting selection. If a reboot is required to change settings, HPE SSA generates a warning.

Minimum power

When settings are selected without regard to system performance, maximum power savings is achieved. Hewlett Packard Enterprise recommends this setting for specific applications, but it is not appropriate for most customers. Most applications will suffer significant performance reduction.

❗ **IMPORTANT:** A reboot may be required after switching power modes to optimize savings and performance.

❗ **IMPORTANT:** When the power mode is set to Balanced, future controller configuration changes may require a reboot for optimal performance.

Modifying power modes

Procedure

1. Open HPE SSA. For more information, see "[Using the HPE SSA GUI](#)."
2. Select the controller.
3. Click on **Configure**.
4. Click **Modify Power Mode**.
5. Select a power mode:
 - Min Power
 - Balanced
 - Max Performance
6. Click **OK**.
7. A summary page appears. Click **Finish** to exit.



Viewing controller status

Use the HPE SSA GUI to view the controller status, including the status of the cache and an energy pack.

Procedure

1. Open HPE SSA. For more information, see ["Using the HPE SSA GUI."](#)
2. Select a controller.
3. Click **Configure**.
4. Under **Controller Configuration Summary**, click **View more details**.

A new window appears.

5. To view the status of the controller, cache, and energy pack, scroll down to **Controller Status**.

Options for erasing drives

When you erase a drive, you remove all sensitive information from a physical drive. Effective methods of erasing sensitive data involve replacing the data with patterns of data and changing the internal encryption keys.

HPE Smart Storage Administrator provides multiple drive-erase options, but not all drives support each option. The product provides the following erase options:

- Sanitize overwrite – (HDD only) Sanitize overwrite fills every physical sector of the drive with a pattern. If you enable this option, you can specify a sanitize method. If you select the restricted option, the drive is unavailable until after the sanitize operation completes successfully. If you select the unrestricted option, the drive is recoverable if the operation fails. See ["Sanitize erase methods"](#) for more information.
- Sanitize Block Erase – (SSD only) Sanitize block erase sets the blocks on the drive to a vendor-specific value, removing all user data. If you enable this option, you can specify a sanitize method. If you select the restricted option, the drive is unavailable until after the sanitize operation completes successfully. If you select the unrestricted option, the drive is recoverable if the operation fails. See ["Sanitize erase methods"](#) for more information.
- Three-pass erase – The three-pass erase method writes random data on the drive for the first and second passes and then zeroes for the third pass. This operation is supported on all drives and is the default erase option for all drives.
- Two-pass erase – The two-pass method writes random data on the drive for the first pass then writes zeroes for the second pass. This operation is supported on all drives.
- One-pass erase – The one-pass method writes zeroes on the drive. This operation is supported on all drives.





IMPORTANT:

- Sanitize erase operations cannot be stopped after starting, and the drive will continue to sanitize after a hot-plug or server reboot.
- During the sanitize erase operation, the drive is unusable until after the process is complete.
- One-, two-, and three-pass erase patterns can be stopped after the erase has begun, but data stored on the drive may not be recoverable even after stopping the erase. Hot-plugging the drive will cancel the erase process.
- A physical drive will remain offline after the erase process has completed if the one-, two-, or three-pass erase pattern is used. To bring the drive online and make it available for configuration, select the drive and click the **Enable Erased Drive** button.
- If the drive reboots during the one-, two-, or three-pass erase, the erase process may not persist following the reboot.

Sanitize erase methods

When you select a sanitize erase option, HPE SSA allows you to specify whether you want the drive to be unavailable until after successful completion or recoverable in case the operation fails. You can choose either of the following methods if your drive supports the method and if you have a license key for the method:

- **Restricted** – Using the restricted sanitize method means that until a drive successfully completes the sanitize operation, it will be unusable. If a restricted sanitize operation fails, you are only allowed to start another sanitize operation, or, if the drive is under warranty, you can return it to HP Enterprise. This method may take several hours.
- **Unrestricted** – Using the unrestricted sanitize method means that the drive will be recoverable in the case that the sanitize erase operation fails. User data may still be present on the drive. Not all drives support this sanitize method.

Consider the following about the sanitize erase methods listed above:

- When a sanitize erase operation is initiated, some drives may report an Estimated Maximum Erase Time. This is an estimate of the time the drive will take to complete the sanitize erase. Note that the drive may take longer than the estimated time.
- Some controllers may not support the sanitize erase options or may require a license key to enable the feature.

Erasing a drive

Procedure

1. Open HPE SSA. For more information, see "[Using the HPE SSA GUI](#)."
2. Open the **Configure** panel by choosing a device and clicking **Configure** in the **Actions** panel.
3. In the **Configure** panel, select **Unassigned Drives** under the Controller Devices heading.
4. From the list of unassigned drives, select the drive or drives that you want to erase, and click the **Erase Drive** button at the bottom of the screen.

A dialog displays the warning messages about the erase feature, the erase pattern options supported by the drive, and the drive that you have selected.

5. Review the messages carefully and select an erase pattern option for the drive. If you want to perform a sanitize erase, select a sanitize erase option and then specify whether you want the sanitize erase to be restricted or unrestricted (if both options are supported by your drive).



6. Click **OK** at the bottom of the screen to proceed or click **Cancel** to return to the previous screen.
7. If you click **OK**, warning messages may appear. Review the messages and click **Yes** to proceed or **No** to close the dialog.
8. If you click **Yes**, you receive a message stating that the erase process has started. Click **Finish** to close the dialog.
In the unassigned drive list, the Information Status Message icon appears on the drive being erased.
9. View the status and informational messages about the erase process by selecting the drive and clicking the **View Details** button. Click the **Refresh** button to update the erase progress.

Some drives may display an erase progress status of 10% or 20%, even up until completion of the erase process.

Managing FLS

Flexible Latency Scheduler provides the ability to control drive latency while still providing the benefit of hard drive optimization. It acts on a hard drive by inspecting a list of requests issued to a drive. FLS changes the controller logic when submitting requests to a rotating drive based on the longest outstanding command latency for a host request on that drive. FLS is a global option for the controller, to be applied to all drives in an effort to reduce the maximum observed latency from a host request.

The FLS attempts to put a cap on the high latency that can be experienced on some rotating disks (hard drives) under highly random workloads. The low setting puts a latency cap of 250 ms on any request, after which the controller schedules the request immediately.

The high setting uses a shorter latency cap of 50 ms, and the very high settings use caps of 30 ms or 10 ms. The target can be any valid individual controller target that supports the flexible latency scheduler feature.

High settings result in generally lower maximum latencies for hard drive logical volumes, but lower throughput for patterns that have improved latency.

FLS is currently available with all HPE Smart Array SR Gen10 controllers except HPE Smart Array S100i SR Gen10 SW RAID.

FLS is available in the following settings.

Setting	Description
Disable	Controller will allow drives to optimize for throughput, resulting in higher maximum latencies for some workloads.
High, Very High	Controller will attempt to compel disks to complete operations in a timely manner, but still allow for some drive optimization. Doing so will lower the effective throughput for some workloads.
Low, Medium	Controller will compel a disk to complete operations at an earlier time compared to the high setting, resulting in lower maximum latencies, but lower throughput.

! **IMPORTANT:** Setting a higher level of FLS can result in a loss of throughput for some request patterns.

Enabling FLS

Procedure

1. Open HPE SSA. For more information, see ["Using the HPE SSA GUI."](#)
2. Select the controller.
3. Click **Advanced Controller Settings**.



4. Scroll down to the option HDD Flexible Latency Optimization.

5. Select one of the following:

- Disabled
- Low
- Middle (100ms)
- High
- Very High (30ms)
- Very High (10ms)

6. Click **OK**.

7. A summary page appears. Click **Finish** to exit.

Diagnostic tasks

The HPE SSA Diagnostics feature replaced the Array Diagnostic Utility supported by SmartStart v8.20 and earlier.

HPE SSA generates the following reports and logs:

- Array diagnostic report

This report contains information about all devices, such as array controllers, storage enclosures, drive cages, as well as logical, physical, and tape drives. For supported solid state drives, this report also contains SmartSSD Wear Gauge information.

- SmartSSD Wear Gauge report

This report contains information about the current usage level and remaining expected lifetime of solid state drives attached to the system.

- Serial output logs

This log details the serial output for the selected controller.

For each controller, or for all of them, you can select the following tasks:

- View Diagnostic Report
- Save Diagnostic Report
- View SmartSSD Wear Gauge Report
- Save SmartSSD Wear Gauge Report

For the view tasks, HPE SSA generates and displays the report or log. For the save tasks, HPE SSA generates a report without the graphical display.

For either task, you can save the report. In online and offline environments, HPE SSA saves the diagnostic report to a compressed folder, which contains an XML report, a plain text report, and a viewer file so you can display and navigate the report through a web browser.

Each HPE SSA Diagnostics report contains a consolidated view of any error or warning conditions encountered. It also provides detailed information for every storage device, including the following:



- Device status
- Configuration flags
- Firmware version numbers
- Physical drive error logs

HPE SSA Diagnostics never collects information about the data content of logical drives. The diagnostic report does not collect or include the following:

- File system types, contents, or status
- Partition types, sizes, or layout
- Software RAID information
- Operating system device names or mount points

Performing a Diagnostics task using HPE SSA

Procedure

1. Open HPE SSA.
2. Open the **Diagnostics** panel by doing one of the following:
 - Select a device and click **Diagnose** in the quick navigation menu.
 - Select an available device from the **Home** screen, and then click **Diagnose** under the available options.
3. Select a report type.
For this example, use the Array Diagnostic Report selection.
4. Select **Array Diagnostic Report**.
The **Actions** panel for Array Diagnostic Report appears.
5. Click one of the task buttons:
 - If you click **View Diagnostic Report**, the report appears. When you are finished viewing the current report, click **Close** or **Save**.
 - If you click **Save Diagnostic Report**, wait for the report to generate, and then click **Close Report** or **Save Report**.

More information

[Reported information](#)

[Identifying and viewing diagnostic report files](#)

[Identifying and viewing SmartSSD Wear Gauge report files](#)

Maintenance tasks

If your controller is validated to be "ready" for online firmware activation, you can update the controller firmware without having to reboot the server.

Before validating that your controller is ready for online firmware activation, ensure that you have installed the compatible firmware and driver. After you install the firmware and driver the first time, you do not need to install them again.



For this feature to be enabled for your controller, your system must meet the following requirements:

- It must use the Linux operating system.
- The firmware you want to install must support online firmware activation.
- Compatible OS driver and utilities are installed on the server.
- Controller is not enabled for encryption.
- SmartCache is disabled.
- Physical drives are not in a failed state.
- No more than 8 drives are attached to the controller, and a drive expander is not used.
- Logical drives are valid.
- Drive firmware update is not in progress.
- Server critical hardware updates are not pending.
- Hardware settings changes in the update firmware that require a reinitialization of the controller are not pending.

Checking online firmware activation readiness

Procedure

1. Open HPE SSA.
2. Select your controller from the Configure menu.
3. In the Actions menu, select **Check Online Firmware Activation Readiness**.
4. Click **OK**.

The results screen appears to verify whether your controller configuration is ready for future firmware updates that are enabled for online activation.

If your controller is not enabled, a list of reasons is displayed. For more information about online firmware activation readiness, see *HPE SmartArray SR Gen10 User Guide* at <http://www.hpe.com/info/SmartArrayGen10-docs> or the release notes for the controller firmware that is enabled for online activation.



HPE SSA CLI

Using the HPE SSA CLI

Access HPE SSA with one of the many methods available:

- **Accessing HPE SSA in the offline environment**
- **Accessing HPE SSA in the online environment**

Some HPE SSA CLI features are only available in the offline environment, such as setting the boot controller and boot volume.

HPE SSA CLI has two operating modes:

- **Console mode**

You can adjust several configuration parameters on several devices without having to restart HPE SSA each time.

- **Command mode**

You can make an isolated change of just one configuration parameter on one device.

For VMware ESXi 5.x, HPE SSA CLI works similarly to the Command mode. However, slight differences exist, requiring you to run HPE SSA CLI through the VMware vSphere esxcli command-line utility. For more information about using HPE SSA CLI for ESXi 5.x, see the VMware Utilities User Guide.

For more information about using HPE SSA CLI for Linux, see the following README.txt files:

- HPE SSA: /opt/smartstorageadmin/ssaccli/README.TXT
- SSACLI (32bit): /opt/smartstorageadmin/ssaccli/ssaccli-1.XX-X.X.XXXX.txt
- SSACLI (64bit): /opt/smartstorageadmin/ssaccli/ssaccli-1.XX-XX.XXXX_XX.txt

! **IMPORTANT: HPE Special Reminder:** Before enabling encryption on the Smart Array controller module on this system, you must ensure that your intended use of the encryption complies with relevant local laws, regulations and policies, and approvals or licenses must be obtained if applicable.

For any compliance issues arising from your operation/usage of encryption within the Smart Array controller module which violates the above mentioned requirement, you shall bear all the liabilities wholly and solely. HPE will not be responsible for any related liabilities.

Opening the CLI in Console mode

The syntax of the command required to open the HPE SSA CLI in Console mode depends on the operating system that you are using.

- For Microsoft Windows, enter the following text: `C:\Program Files\Smart Storage Administrator\ssaccli\bin\ssaccli.exe`

Alternatively, click **Start**, and select **Programs>Windows System>Smart Storage Administrator CLI>Smart Storage Administrator CLI**.

- For Linux, enter the following text: `# ssaccli`



After you have entered Console mode in either operating system, the screen displays the following message and console prompt:

```
HP Smart Storage Administrator CLI 1.XX.X.X
Detecting Controllers...Done.
Type "help" for a list of supported commands.
Type "exit" to close the console.
=>
```

The remaining examples in the HPE SSA CLI section of this guide are described as if entered in Console mode.

Opening the CLI in Command mode

To use Command mode, identify the appropriate **HPE SSA CLI command**. Then, add the command to the end of the text line that is used for **opening the CLI in Console mode**.

The following examples use `help` as the command:

- Using Microsoft Windows: `C:\Program Files\Smart Storage Administrator\ssaccli\bin\ssaccli.exe help`
- Using Linux: `[root@localhost root]# ssaccli help`

The remaining examples in the HPE SSA CLI section of this guide are described as if entered in Console mode.

CLI syntax

Whether entered in Command mode or Console mode, a typical HPE SSA CLI command line consists of three parts: a target device, a command, and a parameter with values if necessary. Using angle brackets to denote a required variable and plain brackets to denote an optional variable, the structure of a typical HPE SSA CLI command line is as follows:

```
<target> <command> [parameter=value]
```

The <target> variable

This variable provides the path to the device that you want to configure. The device can be a controller, an array, a logical drive, or a physical drive. The syntax used is as follows:

```
controller all | slot=# | wwn=# | chassisname="AAA" | serialnumber=# |
chassisserialnumber=# | [array=all|<id>] [logicaldrive all|#] [physicaldrive
all|allunassigned|[:]:#:#,[:]:#:#...|[:]:#:#-[:]:#:#]
```

For example:

```
controller slot=3
controller wwn=500805F3000BAC11
controller slot=2 array A
controller chassisname="A" array B logicaldrive 2
controller chassisname="A" physicaldrive 1:0
controller all
controller slot=2 array all
controller slot=3 physicaldrive 1:2-1:5
```



The <command> variable

The <command> variable can be any of the following words or phrases, corresponding to typical configuration tasks:

add
create
delete
modify
remove
set target

Nonconfiguration commands are also available:

diag
help
rescan
shorthand
show
version

A command often requires a parameter, and some parameters require a value, but the specific parameters and values that are valid for a given command depend on the target to which you apply the command.

To determine the valid parameters and values for a particular combination of target and command variables, you can **query the device**. You can also inspect the example commands in other sections of this guide.

Querying a device

If you do not know what values a parameter can have, you can sometimes query the device to find out by entering a ? as the value of the parameter.

Example command:

```
=> ctrl ch="Lab4" ld 1 modify raid=0 ss=?
```

A typical screen response in this case could be:

Available options are:

8
16 (current value)
32
64
128
256 (default)

To determine which parameters can be queried, use the **help feature of the CLI**.

Hiding warning prompts

When you enter a command for an operation that can potentially destroy user data, the CLI displays a warning and prompts you for input (a y or an n) before continuing the operation. This situation is undesirable when running batch file scripts. To prevent warning prompts from being displayed, use the term `forced` as a parameter.

Example command:




```
ctrl ch="Lab4" ld 1 delete forced
```

Under VMware ESXi 5.x, the forced flag is required for any command that has the option for a forced flag. If the forced parameter is not used, the system generates an error message notifying the user that the command could not be completed without this flag.

Keyword abbreviations

Several commonly used keywords in the HPE SSA CLI have acceptable abbreviations, as shown in the following table. For a complete list of abbreviations, enter `help shorthand` in the HPE SSA CLI.

Keyword	Abbreviation in HPE SSA CLI	Keyword	Abbreviation in HPE SSA CLI
adapterid	ai	nobatterywritecache	nbwc
allunassigned	au	numberparitygroups	npg
arrayaccelerator	aa	parallelscsi	ps
cacheration	cr	physicaldrive	pd
chassisname ¹	ch ¹	postprompttimeout	ppto
chassisserialnumber	csn	queuedepth	qd
chassisslot	chs	raid1writebuffering	r1wb
configurationmode	cm	rebuildpriority	rp
connectionname	cn	serialnumber	sn
controller	ctrl	spareactivationmode	sam
ctrlpath	cp	ssdoverprovisioningoptimization	ssdopo
degradedperformancemode	dpo	ssdphysicaldrive	ssdpd
drivetype	dt	stripesize	ss
drivewritecache	dwc	surfaceanalyseventnotify	saen
elevatorsort	es	surfacescandelay	ssd
enclosure	enc	surfacescanmode	ssm
exitonerror	eo	surfacescanpriority	sp
expandpriority	ep	tapedrive	td
inconsistencyrepairpolicy	irp	waitforcacheroom	wfcr
licensekey	lk	—	—
logicaldrive	ld	—	—
mnpdelay	mpd	—	—

The show command

The `show` command enables you to obtain information about a device.

Syntax:

```
<target> show [detail][[status]
```



When you specify a target that consists of several devices, the information in the output is normally less comprehensive than when you specify only one device as the target. You can use the `[detail]` parameter in this case to retain all the information usually given for individual devices.

An extra parameter is available for controller targets: `config`. This parameter is used as follows:

```
<target controller> show config [detail]
```

If you use the `config` parameter, the output includes information about each device connected to the controller.

Example 1:

```
=> ctrl slot=9 show
```

A typical output would be:

```
Smart Array P421 in Slot 9
Bus Interface: PCI
Slot: 9
Serial Number: PBKTV0XTAZZ005
RAID 6 (ADG) Status: Enabled
Controller Status: OK
Hardware Revision: A
Firmware Version: 0.02-106
Rebuild Priority: Medium
Expand Priority: Medium
Surface Scan Delay: 3 secs
Surface Scan Mode: Idle
Queue Depth: Automatic
Monitor and Performance Delay: 60 min
Elevator Sort: Enabled
Degraded Performance Optimization: Disabled
Inconsistency Repair Policy: Disabled
Wait for Cache Room: Disabled
Surface Analysis Inconsistency Notification: Disabled
Post Prompt Timeout: 15 secs
Cache Board Present: True
Cache Status: OK
Accelerator Ratio: 10% Read / 90% Write
Drive Write Cache: Disabled
Total Cache Size: 1024 MB
Total Cache Memory Available: 816 MB
No-Battery Write Cache: Disabled
Cache Backup Power Source: Capacitors
Battery/Capacitor Count: 1
```



Battery/Capacitor Status: OK
SATA NCQ Supported: True
Spare Activation Mode: Activate on drive failure
Controller Temperature (C): 40
Cache Module Temperature (C): 0
Capacitor Temperature (C): 0
Boot Controller: True
Primary Boot Volume: logicaldrive 1
Secondary Boot Volume: logicaldrive 1

Example 2:

=> ctrl all show

Because this target consists of several devices, the output will be brief. A typical output would be:

MSA1000 at dog (sn: P56350D9IP903J, csn: (9J3CJN71XDCH, wwn: 500805F3000BAC11)
Smart Array 5312 in Slot 3 (sn: P4AB5X9BFMLNTJ)
Smart Array 532 in Slot 2 (sn: P44940LDAORS4F)

Example 3:

=> ctrl ch="lab4" show config

The output in this case will have detailed information because the target consists of only one device. A typical output would be:

MSA1000 at dog (sn: P56350D9IP903J, csn: (9J3CJN71XDCH, wwn: 500805F3000BAC11)
array A (SAS, Unused Space: 20091 MB)
logicaldrive 1 (219 MB, RAID 6(ADG), OK)
physicaldrive 1:1:3 (port 1:box 1:bay 3, SAS, 4.3 GB, OK)
physicaldrive 1:1:4 (port 1:box 1:bay 4, SAS, 9.1 GB, OK)
physicaldrive 1:1:5 (port 1:box 1:bay 5, SAS, 9.1 GB, OK)
physicaldrive 1:1:6 (port 1:box 1:bay 6, SAS, 9.1 GB, OK)
physicaldrive 1:1:7 (port 1:box 1:bay 7, SAS, 9.1 GB, OK)
physicaldrive 1:1:9 (port 1:box 1:bay 9, SAS, ??? GB, failed, spare)
unassigned
drive 1:1:1 (port 1:box 1:bay 1, SAS, 36 GB, OK)
physicaldrive 1:1:2 (port 1:box 1:bay 2, SAS, 36 GB, OK)
physicaldrive 1:1:8 (port 1:box 1:bay 8, SAS, 9.1 GB, OK)
physicaldrive 1:1:10 (port 1:box 1:bay 10, SAS, 9.1 GB, OK)
physical
physicaldrive 1:1:11 (port 1:box 1:bay 11, SAS, 9.1 GB, OK)



Displaying the current versions of the application layers

The `version` command displays the current versions of the application and other software layers.

Syntax:

```
version
```

Example command:

```
version
```

A typical output would be:

```
HP SSA CLI Version: 1.XX.X.X
```

```
SoulAPI Version: X.X.XX.X
```

The help command

To get help with the CLI, enter `help` at the CLI prompt, and then enter one or more help items, as follows:

```
=> help <item1> [item2] [item3]
```

A help item can be any of the following:

- A **CLI command**
- An HPE SSA CLI **keyword or keyword abbreviation**
- A CLI parameter
- A term commonly used in HPE SSA, such as migrate, extend, or cache
- The word `shorthand` (gives a list of abbreviations for keywords in the CLI)

The help feature of the HPE SSA CLI behaves like a browser search engine in that each item that you add to the help input string reduces the amount of help output text. For example, `help controller` produces extensive information, while `help controller remove` restricts the help output to information about how the `remove` command applies to controllers.

Typical procedures

The following sections describe some common HPE SSA CLI procedures.

Setting the boot controller

This procedure applies to booting in Legacy BIOS Boot Mode.

The boot controller is the first controller that the system examines (after power-up) to find a bootable logical drive or volume. The controller boots from the first logical drive by default. Hewlett Packard Enterprise recommends leaving the default setting in place. This setting prevents an error when booting after replacing the boot volume or controller. To assign a designated boot volume, see "**Setting the boot volume.**"

Syntax:

```
<target> modify [bootcontroller=enable|disable]
```

where `<target>` is a controller with a volume that can be set as the boot volume.

Example command:

```
controller slot=1 modify bootcontroller=enable
```



Setting the boot volume



CAUTION: Hewlett Packard Enterprise recommends setting the boot volume only in certain scenarios to prevent a system boot error.

This procedure is available in the online environment and only applies to booting in Legacy BIOS Boot Mode.

The boot volume is the volume that contains the OS and its support files. To work properly, the boot volume must be accessible by the designated boot controller. See "[Setting the boot controller](#)."

Scenario: When the bootable volume is assigned as Logical Drive 2 and a nonbootable logical drive is assigned as Logical Drive 1, a boot volume can be set without resulting in system error.

Syntax:

```
<target> modify [bootvolume=primary|secondary|none]
```

where <target> is a controller and a logical drive.

Example commands:

```
controller slot=1 ld 1 modify bootvolume=primary
controller slot=1 ld 2 modify bootvolume=secondary
controller slot=1 ld 1 modify bootvolume=none
```

Setting the target

If you must perform several operations on a given target device, you can simplify the required commands by setting the device as the default <target> for the CLI operations.

After you have set the target, any command that you enter in the CLI without a specified <target> is automatically applied to the set target. If you must also perform operations on other devices, you can still do so at any time by specifying the <target> for each of those operations as usual. You can also change the set target or clear it completely. The set target is automatically cleared when you close the CLI.



IMPORTANT: You cannot use the `set target` command in batch file scripts.

Syntax:

```
set target <target>
```

where <target> is a controller, array, or logical drive.

Example commands:

```
=> set target ctrl slot=3
=> clear target
```

Typical scenario

First, set a target as follows:

```
=> set target ctrl ch="Lab 4"
=> show target
controller chassisname="Lab 4"
```

As an example of how the set target command works, check the status of array A on this controller:

```
=> array A show
MSA1000 at Lab 4
array A
```



Interface Type: SAS
Unused Space: 7949 MB
Status: OK

The controller does not need to be specified because it is currently the set target.

Now clear the target, reset it, and enter a few commands for the new set target:

```
=> clear target
=> set target ctrl slot=3
=> array A add drives=1:7,1:8,1:9
=> array B add spares=1:10,1:11
=> ctrl slot=4 ld 3 modify ss=64
=> modify rp=high
```

This sequence includes a command for a different target (the controller in slot 4) as a demonstration. The next command in the sequence (the one for modifying the rebuild priority) applies to the controller in slot 3, not the one in slot 4. This is because the command does not specify a `<target>` for the rebuild priority, so the default set target is used instead.

Identifying devices

You can enter a command that causes the LEDs on target devices to flash, enabling you to identify the devices. After an hour, the LEDs stop flashing. You can also stop LED flashing by entering the `off` command.

Syntax:

```
<target> modify led=on|off
```

Example commands:

```
=> ctrl ch="Lab 4" modify led=on
=> ctrl ch="Lab 4" array A modify led=off
```

Deleting target devices

Syntax:

```
<target> delete [forced]
```

where `<target>` can be a controller, array, or logical drive. Except in the case of controllers, you can delete several devices simultaneously if they are of similar type by using the `all` keyword.

Because deleting a target device can result in data loss, the screen displays a warning prompt unless you include the `forced` parameter.

Example commands:

```
=> ctrl ch="Lab 4" delete forced
=> ctrl slot=3 ld all delete
```

Generating a diagnostic report

Hewlett Packard Enterprise provides several ways for administrators to generate diagnostic information, including all three HPE SSA formats and the standalone **HPE Smart Storage Administrator Diagnostics Utility CLI**.

For VMware ESX 5.x, diagnostic reports are generated differently and require you to run another utility, SSADUESXI, through the VMware vSphere CLI. For more information, see the HPE VMware Utilities User Guide.

In HPE SSA CLI, the `diag` command outputs diagnostic information about a specified controller or all controllers on the system. The `ssdrpt` option generates the SmartSSD Wear Gauge report.



Syntax

```
<target> diag <file=filename> [ris=on|off] [ssdrpt=on|off] [xml=on|off]  
[zip=on|off]
```

where the following is designated:

- <target> is a controller or all controllers.
- <file=filename> designates the target file in which the diagnostic information is saved.
- [ris=on|off] determines whether RIS information is or is not included. The `off` value is provided for backward compatibility and is ignored.
- [ssdrpt=on|off] specifies to generate or not generate the Smart SSD Wear Gauge report. The default is `off`.
- [xml=on|off] outputs diagnostic information in formatted XML. The `off` value is provided for backward compatibility and is ignored.
- [zip=on|off] compresses the output to a zipped file. Default behavior is uncompressed. The `off` value is provided for backward compatibility and is ignored.

Example commands

```
ctrl all diag file=c:\allcontrollers.zip  
ctrl slot=4 diag file=c:\ctrl_slot4.zip  
ctrl ch="mybox" diag file=mybox.zip ssdrpt=on
```

Erasing a physical drive

Syntax:

To begin the drive erase process:

```
<target> modify [erasepattern= zero|random_zero|random_random_zero|crypto|  
block|overwrite] [unrestricted=on|off]
```

Where

- the target can be any valid physical drive
- block is for SSDs only
- overwrite is for HDDs only
- crypto, if implemented, is for both SSDs and HDDs

To stop the drive erase process:

```
<target> modify [stoperase]
```

Where

- the target can be any valid physical drive
- stoperase applies to zero, random_zero, and random_random_zero erase patterns only.
- crypto, block, or overwrite drive erase processes cannot be stopped.

Example commands:

```
=> ctrl slot=3 physicaldrive 1e:1:1 modify enableerasedrive
```



```
=> ctrl slot=3 physicaldrive 1e:1:1 modify erase erasepattern=zero
=> ctrl slot=3 pd 1e:1:1 modify stoperase
=> ctrl slot=3 physicaldrive 1e:1:1 modify erase erasepattern=crypto
unrestricted=on
```

Setting sanitize lock

You can enter a command that sets the sanitize lock setting for any valid controller that supports this feature. This feature applies to SATA drives only.

Syntax:

```
<target> sanitizelock=freeze|anti-freeze|none
```

Example commands:

```
controller slot=5 modify sanitizelock=anti-freeze
```

NOTE: If you are enabling the freeze or anti-freeze setting for the first time or modifying the settings, you must reboot the server, and you might need to power cycle or hot-plug the drives.

Rescanning the system

A rescan detects devices that have been added to the system since the previous rescan or since the HPE SSA CLI was started, whichever is more recent.

Syntax:

Use the word `rescan` directly at the HPE SSA CLI prompt, without any target device or parameters.

Example command:

```
=> rescan
```

Entering or deleting a license key

Some advanced configuration tasks (available only on certain controller models) can be performed only after software is installed on the controller and a license key is registered to activate the software.

Syntax:

```
<target> add [lk=xxxxx-xxxxx-xxxxx-xxxxx-xxxxx]
```

where the target is any valid controller. The hyphens are optional.

To delete a license key, use a standard `delete` command but use the license key (not the controller) as the target:

```
<target> delete
```

Example commands:

```
=> ctrl slot=5 lk=12345-65432-78787-43434-24680 delete
=> ctrl slot=4 add lk=9876543210222224444466666
```

Optimizing controller performance for video

On some controller models, you can optimize the controller performance for video.

Syntax:

```
<target> modify dpo=enable
<target> modify elevatorsort=disable
<target> modify irp=enable
```




```
<target> modify queuedepth=automatic
```

```
<target> modify mnpd=60
```

where the target is any valid controller.

The `queuedepth` parameter can also have a numerical value from 2 to 32, and the `mnpd` parameter can have any value from 0 (disabled) to 60.

To disable the video performance optimization feature, reverse the disable and enable values noted above, set the queue depth to automatic, and set `mnpd` to 0, as in the second example command.

Example commands:

```
=> ctrl slot=5 modify dpo=enable
=> ctrl slot=5 modify elevatorsort=disable
=> ctrl slot=5 modify irp=enable
=> ctrl slot=5 modify queuedepth=16
=> ctrl slot=5 modify mnpd=25
=> ctrl slot=3 modify dpo=disable
=> ctrl slot=3 modify elevatorsort=enable
=> ctrl slot=3 modify irp=disable
=> ctrl slot=3 modify queuedepth=automatic
=> ctrl slot=3 modify mnpd=0
```

Creating a logical drive

Syntax:

```
<target> create type=ld [parameter=value]
```

`<target>` is usually a controller, but it can be an array if you are creating an additional logical drive on an existing array.

If you want to create a logical drive on a group of physical drives that are not yet assigned to an array, you do not have to build the array first. In the CLI, unlike in the GUI, the array is created automatically at the same time as the logical drive.

The standard parameters used when creating a logical drive are described in the following table. If you do not specify a particular parameter, the CLI uses the appropriate default value.

Parameter	Acceptable values	Comments
drives	[#:]#:#, [#:]#:#, ... [#:]#:#-[#:]#:# all allunassigned	The default setting is <code>all</code> .
raid	0 1 1adm 1+0 1+0adm 5 50 6 60 ?	The default setting is the highest RAID level that the controller model and number of drives can support, except 50 or 60 (which must be specified deliberately).
numberparitygroups	2 #	The default value is 2.
ss	8 16 32 64 128 256 512 1024 default ?	Units are KB. The default setting depends on the RAID level.

Table Continued



Parameter	Acceptable values	Comments
size	# min max maxmbr ?	This parameter determines the preferred size of the logical drive. Units are MB. ¹ The default setting is max.
sectors	32 63 default ?	The default setting depends on the operating system.
aa	enable disable ?	The default setting is enable.
drivetype	sas satalogical sata saslogical parallelscsi ss_sas ss_sata ?	—

If you specify an array as the target, you can omit the `drives` parameter, because the drives are already implicitly defined by the array ID. This feature is useful if you are entering the command directly into the CLI console because you do not need to remember which drives belong to the array. When you write a batch file, however, it is often easier to specify every drive in the array than to parse out the array ID.

When you use the `drives` parameter you can list each drive individually, you can specify a drive range, or you can specify both a range and some individual drives. A drive range can span ports, boxes, and bays. If you specify individual drives, they do not have to form a continuous sequence. If you specify a range, the CLI automatically excludes from the target any drive in the range that is unavailable (for example, a drive is excluded if it already belongs to an array, is a spare, has too small a capacity, or has failed).

If you want to specify an existing array by its drives rather than by array ID, then all of the drives that you specify must belong to the same array, and none of the drives in the array can be omitted.

Example commands:

```
ctrl slot=5 create type=ld drives=1:0,1:1,1:2,1:3 raid=adg
ctrl slot=5 create type=ld drives=1:1-1:4 raid=adg
ctrl slot=5 create type=ld drives=1:7,1:10-2:5,2:8-2:12 raid=adg
ctrl slot=5 array A create type=ld size=330 raid=adg
```

The following pair of commands demonstrates how the `drives` parameter can be used in a batch file to create two logical drives on the same array, one of 330 MB and the other of 450 MB:

```
ctrl slot=2 create type=ld drives=1:1-1:6 size=330 raid=adg
ctrl slot=2 create type=ld drives=1:1-1:6 size=450 raid=5
```

Sample scenario

Consider a situation in which you want to create two arrays. One of these arrays needs two logical drives, while the other needs only one.

First, determine what physical drives are available and what their properties are:

```
=> ctrl ch="Lab 4" pd all show
```

For this sample scenario, the screen response is:

```
MSA1000 at Lab 4
unassigned
physicaldrive 1:1:12 (port 1:box 1:bay12, SAS, 36.4 GB, OK)
physicaldrive 1:1:13 (port 1:box 1:bay13, SAS, 9.1 GB, OK)
physicaldrive 1:1:14 (port 1:box 1:bay14, SAS, 9.1 GB, OK)
```



Knowing this information, you can now create the first array with one logical drive:

```
=> ctrl ch="Lab 4" create type=ld drives=1:12
```

Now, verify that the array has been created:

```
=> ctrl ch="Lab 4" pd all show
```

In this case, the screen response is:

```
MSA1000 at Lab 4
array A
physicaldrive 1:1:12 (port 1:box 1:bay12, Parallel SCSI, 36.4 GB, OK)
unassigned
physicaldrive 1:1:13 (port 1:box 1:bay13, Parallel SCSI, 9.1 GB, OK)
physicaldrive 1:1:14 (port 1:box 1:bay14, Parallel SCSI, 9.1 GB, OK)
```

Create the second array on the two remaining physical drives. Before creating this array, determine what RAID options are available for these drives:

```
=> ctrl ch="Lab 4" create type=ld drives=1:13,1:14 size=300 raid=?
```

The response in this case is:

Available options are:

```
0
1 (default value)
```

Now, create the new array:

```
=> ctrl ch="Lab 4" create type=ld drives=1:13,1:14 size=300 raid=1
```

It is not strictly necessary to specify the RAID level in this example because it is the highest possible level for this scenario and will be used by default. However, it is included in the command as an example.

Now, verify that the array has been formed:

```
=> ctrl ch="Lab 4" pd all show
```

The screen response is:

```
MSA1000 at Lab 4
array A
physicaldrive 1:1:12 (port 1:box 1:bay12, SAS, 36.4 GB, OK)
array B
physicaldrive 1:1:13 (port 1:box 1:bay13, SAS, 9.1 GB, OK)
physicaldrive 1:1:14 (port 1:box 1:bay14, SAS, 9.1 GB, OK)
```

To create a second logical drive on array B, you can specify the array (method A) or you can specify every physical drive in the array (method B).

```
=> ctrl ch="Lab 4" array B create type=ld size=900 (method A)
```

```
=> ctrl ch="Lab 4" create type=ld drives=1:13,1:14 size=900 (method B)
```

Finally, verify that the logical drives have all been created correctly:

```
=> ctrl ch="Lab 4" ld all show
```

```
MSA1000 at Lab 4
array A
```



```
logicaldrive 1 (33.9 GB, RAID 0, OK)
array B
logicaldrive 2 (298 MB, RAID 1, OK)
logicaldrive 3 (896 MB, RAID 1, OK)
```

Moving a logical drive

HPE SSA enables you to move a single logical drive from one array to another array.

When moving the logical drive, you can choose one of the following destinations:

- An existing array
- A new array you create

To move a logical drive to an existing array, use the following information.

Syntax:

```
<target> modify [newarray=]
```

Command examples:

```
ctrl slot=1 ld 3 modify newarray=C
ctrl slot=1 ld 3 modify newarray=? // shows the available arrays
```

To move a logical drive and create a new array, use the following information.

Syntax:

```
<target> modify [drives=[#: ]#:#,[#: ]#:#,[#: ]#:#-[#: ]#:#,...]
```

Command examples:

```
ctrl slot=1 ld 3 modify drives=2e:1:1-2e:1:4
ctrl slot=1 ld 3 modify drives=? // shows the available drives
```

You can use the `drivetype` parameter to specify the drive interface type. When selecting all physical drives, and multiple drive types exist, you must specify the drive interface type. Mixed drives are not allowed on the same array or logical drives. If all drives on a controller are the same type, this parameter is not needed.

For `drivetype`, use one of the following valid options:

```
[drivetype=sas | sata | saslogical | parallelscsi | ss_sas |
ss_sata | ?]
```

Command examples:

```
controller slot=5 ld 1 modify drives=? drivetype=sas
controller slot=5 array A modify drives=? drivetype=ss_sas
```

Viewing enclosure information

The `enclosure` keyword is a valid `target` for the `show` command, and it outputs the storage enclosure information.

Syntax:

```
enclosure [ all | port:box | serialnumber=xxx ] show [ detail | status ]
```

where the `target` can be any valid storage controller.

Example commands:

```
controller slot=5 enclosure all show
```



```
controller slot=5 enclosure 4E:1 show detail
controller slot=5 enclosure serialnumber=UAB123456 show status
```

Viewing SSD physical drives

The `ssdphysicaldrive` keyword is a valid target for the `show` command, and it outputs the solid state physical drives on the specified controller.

Syntax:

```
<target> ssdpd all show [detail]
```

where the target is any valid controller.

Example commands:

```
controller slot=5 ssdpd all show
controller slot=5 ssdpd all show detail
```

Viewing SSD information

The `ssdinfo` keyword provides a summary for solid state drives for the specified controller.

Syntax:

```
<target> show ssdinfo [detail | summary]
```

Command examples:

```
controller slot=1 show ssdinfo
controller all show ssdinfo
controller all show ssdinfo detail
controller all show ssdinfo summary
```

Verifying online firmware activation readiness for current configuration

The `ofaready` keyword indicates whether the current controller configuration allows online firmware activation.

Syntax:

```
<target> show ofaready
```

For more information about online firmware activation readiness, see *HPE SmartArray SR Gen10 User Guide* at <http://www.hpe.com/info/SmartArrayGen10-docs> or the release notes for the controller firmware that is enabled for online activation.

SmartCache in SSACLI

SmartCache allows a caching logical drive to be created on a solid state physical drive, and to be associated with an existing data logical drive to provide a performance increase.

Enabling SmartCache:

To create a new SmartCache array for existing data logical drive 1 using the solid state drive 1e:1:10:

```
=> ctrl slot=1 create type=ldcache drives=1e:1:10 dataId=1
```

The command above creates a new SmartCache array, for example, "array B." Only one SmartCache array may exist on a controller. This example creates an additional caching logical drive using the existing SmartCache array for existing data logical drive 2:

```
=> ctrl slot=1 array B create type=ldcache dataId=2
```

Disabling SmartCache:



To disable SmartCache, you must delete the caching logical drive created when enabling SmartCache. To determine which logical drive is the caching logical drive, use the command `ld all show`.

```
=> ctrl slot=0 ld all show
```

The results indicate which logical drive is the SmartCache.

```
=> ld all show
```

```
Smart Array P421 in Slot 2
```

```
array A
```

```
logicaldrive 1 (68.3 GB, RAID 1, OK, cached)
```

```
SmartCache array B
```

```
logicaldrive 2 (16.0 GB, RAID 0, OK, SmartCache for Logical Drive 1)
```

Once the SmartCache is identified, you can delete the logical drive, disabling SmartCache.

```
=> ctrl slot=0 ld 2 delete
```

HPE SSD Smart Path

Hewlett Packard Enterprise SSD Smart Path enables I/O requests to bypass Smart Array firmware and access SSDs directly. This process accelerates reads for all RAID levels and writes for RAID 0.

The controller device driver analyzes each I/O request for the following Smart Path eligibility requirements:

- Read requests for all supported RAID levels
- Write requests for all RAID 0 volumes
- Size and location parameters, resulting in an operation that can be satisfied by a single physical disk I/O

Target:

The target can be any valid controller.

Example commands:

```
controller slot=0 create type=ld size=5000 raid=5 drives=2I:1:2-2I:1:7  
ssddmartpath=enable
```

```
controller slot=0 array a modify ssddmartpath=disable
```

Rapid Parity Initialization methods

RAID levels that use parity (RAID 5, RAID 6 (ADG), RAID 50, and RAID 60) require that the parity blocks be initialized to valid values. Valid parity data is required to enable enhanced data protection through background surface scan analysis and higher performance write operations. Two initialization methods are available:

- **Default** – Initializes parity blocks in the background while the logical drive is available for access by the operating system. A lower RAID level results in faster parity initialization.
- **Rapid** – Overwrites both the data and parity blocks in the foreground. The logical drive remains invisible and unavailable to the operating system until the parity initialization process completes. All parity groups are initialized in parallel, but initialization is faster for single parity groups (RAID 5 and RAID 6). RAID level does not affect system performance during rapid initialization.

Rapid Parity Initialization is available only for supported controllers and in arrays composed of supported physical drives.

Example of creating a logical drive with Rapid Parity Initialization:



```
=> ctrl slot=3 create type=ld drives=1e:1:1-1e:1:4  
parityinitializationmethod=rapid
```

SSD Over Provisioning Optimization

Access to solid state devices can be optimized by deallocating all used blocks before any data is written to the drive. The optimization process is performed when the first logical drive in an array is created, and when a physical drive is used to replace a failed drive. Not all controllers support this option.

This feature is disabled by default in the CLI.

When creating the first logical drive in an array, the `ssdoverprovisioningoptimization` parameter specifies to initialize solid state drives supporting the rapid parity initialization feature. The only valid option is "on".

Example:

```
ctrl slot=1 create type=ld drives=1I:1:1 ssdoverprovisioningoptimization=on  
ctrl slot=1 create type=ld drives=1I:1:2 ssdopo=on
```

Managing spare drives

Assigning one or more online spare drives to an array enables you to postpone replacement of faulty drives. However, it does not increase the fault-tolerance level of any logical drives in the array. For example, a logical drive in a RAID 5 configuration suffers irretrievable data loss if two of its physical drives are simultaneously in a failed state, regardless of the number of spare drives assigned to the array.

Any drive that you want to use as a spare must meet the following criteria:

- It must be an unassigned drive or a spare for another array.
- It must be of the same type as existing drives in the array (for example, SATA or SAS).
- It must have a capacity no less than that of the smallest drive in the array.

Syntax:

```
<target> add spares=[#:]#:#, [#:]#:#, [#:]#:#-[#:]#:#, ...|allunassigned  
[forced]
```

```
<target> remove spares=[#:]#:#, [#:]#:#, [#:]#:#-[#:]#:#, ...|all
```

where <target> is an array (or logical drive, if the array contains only one logical drive). The `forced` parameter represses any warning message prompts. If you specify a drive range, any drives in the range that do not meet the previous criteria are not used.

Example commands:

```
=> ctrl slot=3 array B add spares=1:6  
=> ctrl slot=4 array all add spares=1:5,1:7  
=> ctrl slot=5 array A add spares=1:1-1:5  
=> ctrl slot=5 array A remove spares=1:1-1:5
```

Setting the spare activation mode

The spare activation mode feature enables the controller firmware to activate a spare drive under the following conditions:

- When a data drive reports a predictive failure (SMART) status
- When a data drive fails; this mode is the default.



In normal operations, and for older controllers, the firmware starts rebuilding a spare drive only when a data drive fails. With the predictive failure activation mode, rebuilding can begin before the drive fails, reducing the likelihood of data loss that could occur if an additional drive fails.

Use the `spareactivationmode` keyword to toggle the spare activation mode for the controller between drive failure and drive predictive failure.

Syntax:

```
<target> modify spareactivationmode=[ failure | predictive ]
```

Example commands:

```
controller slot=1 modify spareactivationmode=predictive
controller slot=1 modify spareactivationmode=failure
```

Spare Management Mode in SSACLI

The `sparetype` keyword allows you to specify the spare type as "dedicated" (default) or "autoreplace." A dedicated spare temporarily takes over for a failed drive and can be shared between arrays. An autoreplace spare replaces a failed drive and cannot be shared between arrays.

Command example:

To add a spare with spare type of autoreplace:

```
=> ctrl slot=1 array A add spares=1e:1:5 sparetype=autoreplace
```

Expanding an array

You can increase the storage space on an array by adding physical drives. Any drive that you want to add must meet the following criteria:

- It must be an unassigned drive.
- It must be of the same type as existing drives in the array (for example, SATA or SAS).
- It must have a capacity no less than the capacity of the smallest drive in the array.



IMPORTANT: An array expansion, logical drive extension, or logical drive migration takes about 15 minutes per gigabyte. While this process is occurring, no other expansion, extension, or migration can occur simultaneously on the same controller.

Syntax:

```
<target> add drives=[#:]#:#, [#:]#:#, [#:]#:#-[#:]#:#, ...|allunassigned
[forced]
```

Where `<target>` is an array (or a logical drive, if the array contains only one logical drive). The `forced` parameter represses any warning message prompts. If you specify a drive range, any drives in the range that do not meet the previous criteria are not used.

If you add an odd number of drives to an array that contains a RAID 1+0 logical drive, you are prompted to convert the RAID 1+0 logical drive to RAID 5 or RAID 6 (ADG). Adding the `forced` parameter to the command prevents this prompt from appearing.

Example commands:

```
=> ctrl slot=3 array A add drives=1:0,1:1
=> ctrl slot=4 ld 1 add drives=allunassigned
=> ctrl slot=5 array A add drives=1:1-1:5
```



Shrinking an array

Some controllers may not support this option or may require a license key to enable this feature.

You can shrink the size of an array by removing a drive from an existing array. Observe the following criteria:

- After the shrink, the array must have sufficient capacity to contain all of the configured logical volumes.
- You may not remove drives from the array if the resulting number of drives does not support the fault tolerance (RAID level) of any existing logical drive. For example, if you have an array with four physical drives and a RAID 5 logical drive, you can remove, at most, one drive, because RAID 5 requires at least three physical drives.
- If the array contains a RAID 1+0 logical drive, you can remove only an even number of drives.
- If the array contains a compound RAID (RAID 50 or RAID 60) logical drive, you can remove drives only in multiples of the number of parity groups. For example, an array with 10 physical drives and a RAID 50 logical drive can be shrunk by removing two or four drives only.

Syntax:

```
<target> remove drives=[#:]#:#-[#:]#:#
```

where <target> is an array, and the specified physical drives are being removed to shrink the array.

For example, in an existing array (array a), six drives (1e:1:4-1e:1:9) are in use. With all criteria met, you can shrink the array to four drives by removing the last two drives with the command: <array a> remove drives=1e:1:8-1e:1:9

Example commands:

```
=> array a remove drives=1e:1:12-1e:1:14
```

```
=> array b remove drives=1c:1:6-1c:1:7
```

Moving an array

Some controllers may not support this option or may require a license key to enable this feature.

You can move an array by designating different physical drives for the array. To move the array, each of the physical drives where the array will reside must meet the following criteria:

- It must be an unassigned drive.
- It must be of the same type as the physical drives currently in the source array (for example, SATA or SAS).
- The destination drives must have sufficient capacity to hold all the logical drives present in the source array.

Like array creation and expansion, the useable space in all drives is reduced to the size of the smallest physical drive in the destination disk set.

Moving an array automatically removes any previously assigned spare drives. If spares are assigned to the existing array, they must be designated for the array when it is moved.

Syntax:

```
<target> modify drives=[#:]#:#-[#:]#:#
```

where <target> is an array, and the specified physical drives are the new destination for the array.

For example, in an existing array, three 72-GB SAS drives (1e:1:4-1e:1:6) are the source. Another drive of the same size is the spare (1e:1:9). With all criteria met, you can move the array to three different 72-GB SAS drives by specifying the new destination drives in the command (1e:1:12-1e:1:14). To maintain the same spare drive, be sure to designate the spare drive for the moved array.

Example commands:



```
=> array a modify drives=1e:1:12-1e:1:14
=> array b modify drives=1c:1:6-1c:1:7
```

Replacing an array

Some controllers may not support this option or may require a license key to enable this feature.

You can replace an array by designating different physical drives for the array. To replace the array, each of the physical drives where the array will reside must meet the following criteria:

- The destination drives must all be the same type (for example, SATA or SAS), but do not have to be same type as the drives in the source array.
- A destination drive does not have to be an unassigned drive, but if it is assigned, it must be a drive assigned to the array being replaced.
- The destination drives must have sufficient capacity to hold all the logical drives present in the source array.

Like array creation and expansion, the useable space in all drives is reduced to the size of the smallest physical drive in the destination disk set.

Replacing an array automatically removes any previously assigned spare drives. If spares are assigned to the existing array, they must be designated for the array when it is moved.

Syntax:

```
<target> modify drives=[#:]#:#-[#:]#:#
```

where <target> is an array, and the specified physical drives are the new destination for the array.

For example, in an existing array, three 72-GB SAS drives (1e:1:4-1e:1:6) are the source. Another drive of the same size is the spare (1e:1:9). With all criteria met, you can move the array to three different 72-GB SAS or SATA drives by specifying the new destination drives in the command (1e:1:12-1e:1:14). To maintain the same spare drive, be sure to designate the spare drive for the moved array.

Example commands:

```
=> array a modify drives=1e:1:12-1e:1:14
=> array b modify drives=1c:1:6-1c:1:7
```

Setting Auto RAID 0

Auto Array RAID 0 creates a single RAID 0 volume on each physical drive specified, enabling the user to select multiple drives and configure as RAID 0 simultaneously. Each array contains one physical drive and one RAID 0 logical drive. For more information about RAID 0, see "[RAID 0—No fault tolerance](#)."

Example command:

```
=> ctrl slot=1 create type=arrayr0 drives=allunassigned
```

Extending a logical drive

If the operating system supports logical drive extension, you can use any unassigned capacity on an array to enlarge one or more of the logical drives on the array.

ⓘ IMPORTANT: An array expansion, logical drive extension, or logical drive migration takes about 15 minutes per gigabyte. While this process is occurring, no other expansion, extension, or migration can occur simultaneously on the same controller. Controllers that support flash-backed write cache only do not support this process.

Syntax:

```
<target> modify size=#|max|? [forced]
```



where <target> is a logical drive.

If the operating system does not support logical drive extension, carrying out this command makes data on the logical drive unavailable. Therefore, the CLI displays a warning prompt as a safeguard in case you are using such an operating system. To prevent the prompt from appearing, use the `forced` parameter.

Example commands:

```
=> ctrl slot=3 ld 1 modify size=max
=> ctrl slot=4 ld 1 modify size=?
=> ctrl slot=3 ld 2 modify size=500 forced
```

Migrating a logical drive

This command enables you to adjust the stripe size (data block size) or RAID level of a selected logical drive. For more information, see "**Selecting a RAID method.**"

Consider the following factors before performing a migration:

- For some RAID-level migrations to be possible, you might need to expand the array by adding one or more drives.
- For migration to a larger stripe size to be possible, the array might need to contain unused drive space. This extra space is necessary because some of the larger data stripes in the migrated array are likely to be filled inefficiently.

! **IMPORTANT:** An array expansion, logical drive extension, or logical drive migration takes about 15 minutes per gigabyte. While this process is occurring, no other expansion, extension, or migration can occur simultaneously on the same controller. Controllers that support flash-backed write cache only do not support this process.

Syntax:

```
<target> modify [raid=0|1+0|1|5|6|50|60|adg|1adm|10adm|?] [ss=8|16|32|64|128|
256|512|1024|default|?]
```

where <target> is a logical drive.

The following limitations apply to this command:

- You cannot simultaneously query the RAID level and the stripe size of any given logical drive.
- If you do not specify a RAID level for a query or migration, the CLI uses the existing value by default.
- If you do not specify a stripe size, the CLI uses the default stripe size value for the RAID level that you specify.

Example commands:

```
=> ctrl slot=3 ld 1 modify raid=1
=> ctrl slot=4 ld 2 modify ss=16
=> ctrl slot=2 ld 3 modify raid=5 ss=16
```

Changing the Rebuild Priority setting

The Rebuild Priority setting determines the urgency with which the controller treats an internal command to rebuild a failed logical drive.

- At the low setting, normal system operations take priority over a rebuild.
- At the medium setting, rebuilding occurs for half of the time, and normal system operations occur for the rest of the time.



- At the medium high setting, rebuilding is given a higher priority over normal system operations.
- At the high setting, the rebuild takes precedence over all other system operations.

If the logical drive is part of an array that has an online spare, rebuilding begins automatically when drive failure occurs. If the array does not have an online spare, rebuilding begins when the failed physical drive is replaced.

Syntax

```
<target> modify rp=high|mediumhigh|medium|low|?
```

where <target> is a controller.

Example command

```
=> ctrl slot=3 modify rp=high
```

Changing the Expand Priority setting

The Expand Priority setting determines the urgency with which the controller treats an internal command to expand an array.

- At the low setting level, normal system operations take priority over an array expansion.
- At the medium setting, expansion occurs for half of the time, and normal system operations occur for the rest of the time.
- At the high setting, the expansion takes precedence over all other system operations.

Syntax:

```
<target> modify ep=high|medium|low|?
```

where <target> is a controller.

Example command:

```
=> ctrl slot=3 modify ep=high
```

Setting the surface scan mode

The `surfacescanmode` keyword sets the surface mode for the controller. The available modes are `disable`, `high`, or `idle`. If you specify the `idle` mode, you must also specify a surface scan delay value.

Syntax:

```
<target> modify [ surfacescanmode=disable | idle | high | ? ]
```

Example commands:

```
controller slot=1 modify surfacescanmode=high
```

```
controller slot=1 modify surfacescanmode=disable
```

```
controller slot=1 modify surfacescanmode=idle surfacescandelay=3
```

Changing the surface scan delay time

The setting for the surface scan delay determines the time interval for which a controller must be inactive before a surface scan analysis is started on the physical drives that are connected to it.

Surface scan analysis is an automatic background process that ensures that you can recover data if a drive failure occurs. The scanning process checks physical drives in fault-tolerant logical drives for bad sectors, and in RAID 5 or RAID 6 (ADG) configurations, it also verifies the consistency of parity data.

Syntax:



```
<target> modify ssm=idle ssd=#
```

where <target> is a controller and # is a number between 1 and 30. This number determines the delay time in seconds, but you do not need to include units with the command.

Example command:

```
=> ctrl sn=P56350D9IP903J modify ssm=idle ssd=3
```

Parallel Surface Scan

This commands allows the user to set the parallel surface scan count for the controller. A value of "1" will disable the feature, and the maximum count value is "16".

Target:

The target can be any valid controller.

Example commands:

```
controller slot=1 modify surfacescanmode=idle surfacescandelay=3
parallelsurfacescancount=16

controller slot=1 modify parallelsurfacescancount=16

ctrl slot=3 modify pssc=16
```

Re-enabling a failed logical drive

If a logical drive has failed and the data on it is invalid or non-recoverable, you can re-enable the logical drive so that it can be reused. This process preserves the structure of the logical drive and merely deletes data, whereas a `delete` command applied to a logical drive deletes the logical drive structure as well as the data.

Syntax:

```
<target> modify reenable [forced]
```

Example command:

```
=> ctrl slot=3 ld 1 modify reenable forced
```

Changing the controller cache ratio

The controller cache ratio setting determines the amount of memory allocated to read and write operations. Different types of applications have different optimum settings. You can not change the ratio if the controller has a flash-backed write cache and if there are logical drives configured on the controller.

Syntax:

```
<target> modify cr=#/#|?
```

where <target> is a controller, and #/# is the cache ratio in the format read percentage/write percentage.

Example command:

```
=> ctrl slot=3 modify cr=25/75
```

Enabling or disabling the drive cache

On controllers and drives that support physical drive write cache, you can use this command to enable or disable the write cache for all drives on the controller.



CAUTION: Use a backup power supply to minimize the possibility of data loss if a write process occurs during a power failure.



Syntax:

```
<target> modify drivewritecache=enable|disable [forced]
```

where <target> is a controller that supports drive write cache for configured or unconfigured drives.

Example command:

```
=> ctrl slot=5 modify dwc=enable
```

Enabling or disabling the array accelerator

If the controller has an array accelerator, you can disable or enable it for specified logical drives.

NOTE: Disabling the array accelerator for a logical drive reserves use of the accelerator cache for other logical drives on the array. This feature is useful if you want the other logical drives to have the maximum possible performance (for example, if the logical drives contain database information).

Syntax:

```
<target> modify aa=enable|disable|?
```

where <target> is a logical drive.

Example command:

```
=> ctrl slot=3 ld 1 modify aa=enable
```

Enabling a script to exit on error

When a script encounters an error during execution, the `exitonerror` keyword enables the script to decide whether to continue or to exit the application and return an exit code. When `exitonerror` is enabled, the last command line entry will cause the application to exit.

Syntax:

```
set [exitonerror=enable | disable]
```

Example commands:

```
set exitonerror=enable
```

```
set eoe=disable
```

```
show exitonerror
```

Using the Split Mirror command

The "splitmirror" command splits original RAID 1, RAID10 or ADM mirrored volumes into individual RAID0 volumes. The command will also split volumes into RAID1 or RAID10 volumes if ADM, and can be used to recombine the volumes. Optionally, the command will hide the new volume from the OS.

Syntax

```
<target> splitmirror action=[splitwithbackup|remirror|rollback|
activatebackup]
```

where <target> is the array.

Actions

The "splitwithbackup" action splits an array consisting of one or more RAID 1, RAID 1+0, RAID 1 (ADM), or RAID 10 (ADM) and creates two arrays: A primary array and a backup array. The new arrays have the following characteristics:



- If the original array contained RAID 1 or RAID 1+0 drives, the primary array will contain RAID 0 drives.
- If the original array contained RAID 1 (ADM) drives, the primary array will contain RAID 1 drives.
- If the original array contained RAID 10 (ADM) drives, the primary array will contain RAID 1+0 drives.

The backup array always contains RAID 0 logical drives. The primary array continues to be fully accessible to the OS while the backup array is hidden from the operating system.

The "remirror" action remirrors the array and preserves the existing data. Discard the backup array. This option recreates the original mirrored array with the contents of the primary array.

The "rollback" action remirrors the array and rollback to the contents of the backup array. Discard existing data. This option recreates the mirrored array but restores its contents to the point in time when the backup array was created. It is not recommended that this option be performed online, or while the logical drive to be rolled back is mounted or in use by the OS.

The "activatebackup" action activates the backup array. This option makes the backup array fully accessible to the operating system. The target can be any valid individual array target.

Example command

```
=> ctrl slot= 5 array A splitmirror action=split with backup
=> ctrl slot= 5 array A splitmirror action=remirror
=> ctrl slot= 5 array A splitmirror action=rollback
=> ctrl slot= 5 array A splitmirror action=activatebackup
```

Enabling SmartCache Write-back

⚠ CAUTION: Specifying the write-back cache write policy could result in data loss in the event of failure of the cache volume. When using a RAID 0 cache volume, a single SSD failure might result in data loss.

❗ IMPORTANT: If a demo license key expires, all SmartCache volumes configured with write-back cache write policy are converted to write-through. When this happens, the logical drive details show different values for Cache Write Policy and Cache Write Policy Requested. If the license is reinstalled, the SmartCache volumes are restored to their original write-back cache write policy.

Syntax:

```
<target> create [type=ldcache] [drives=##:##:##] [dataId]=#
[writethrough|writeback|?]
```

where <target> is any valid SSD drive and existing non-cached logical drive for the dataId.

Example command:

```
ctrl slot=1 create type=ldcache drives=1i:1:1,1i:1:2 dataId=1
ctrl slot=1 create type=ldcache drives=1i:1:1,1i:1:2 dataId=1
writepolicy=writethrough
```

Managing FLS

❗ IMPORTANT: Setting a higher level of FLS can result in a loss of throughput for some request patterns.

This setting controls the flexible latency scheduler.

Syntax:



```
<target> modify [latency=low|high|disable|?]
```

Where the <target> is any valid individual controller target that supports the flexible latency scheduler feature.

Example command:

```
<controller slot=5 modify latency=low>
```

Power mode

❗ **IMPORTANT:** A reboot may be required after switching power modes in order to optimize savings and performance.

❗ **IMPORTANT:** When the power mode is set to Balanced, future controller configuration changes may require a reboot for optimal performance.

Syntax:

```
<target> modify [powermode=minpower|balanced|maxperformance|?]
```

Where <target> is any valid controller supporting this feature.

This command sets the power for the controller.

- A value of `minpower` sets static settings to the lowest possible value and dynamically reduces power based on the workload.
- A value of `balanced` sets statistic settings based on the configuration and dynamically reduces power based on the workload (recommended).
- A value of `maxperformance` sets static settings to the highest possible value and does not dynamically reduce power (default).

The target can be any valid individual controller target that supports the power mode feature.

Example command:

```
controller slot=5 modify powermode=balanced
```

Encryption commands

❗ **IMPORTANT: HPE Special Reminder:** Before enabling encryption on the Smart Array controller module on this system, you must ensure that your intended use of the encryption complies with relevant local laws, regulations and policies, and approvals or licenses must be obtained if applicable.

For any compliance issues arising from your operation/usage of encryption within the Smart Array controller module which violates the above mentioned requirement, you shall bear all the liabilities wholly and solely. HPE will not be responsible for any related liabilities.

The following section describes common CLI commands for Secure Encryption, used in conjunction with HPE SSA. A separate license is required to use Secure Encryption. For more information about Secure Encryption, see the *HPE Smart Array SR Secure Encryption Installation and User Guide*.

For a complete list of supported commands, enter `help`.

Logging in and out

These commands allow the user to log in or out of Secure Encryption. Log in to the encrypting controller as either the "crypto" user or as an additional authorized user. The user will be prompted for the password if it is not provided on the command line. Password example text is indicated with "XXXXXX."



Target:

The target can be any valid controller supporting encryption.

Example commands:

Setting	Command
To log in	<pre>=> ctrl slot=7 login user=crypto => ctrl slot=7 login user=crypto password=XXXXXX => ctrl slot=7 login user=user</pre>
To log out	<pre>=> ctrl slot=7 logout</pre>

Toggling encryption on a controller

These commands allow the user to enable or disable encryption on a controller. Use the `enableencryption` command to initialize the controller by setting a password and a local master encryption key. To enable or disable encryption, change the encryption mode (local or remote), and enable or disable creation of future plaintext volumes, the master key must be specified.

The `login` command is used for authenticating either the `crypto` user or another user to the system in order to perform encryption functions.

Password and master key example text is indicated with "XXXXXX."

Target:

The target can be any valid controller supporting encryption.

Example commands:

The following command indicates all command options for enabling or disabling encryption. The table below provides specific examples.

```
<target> modify encryption=[on|off] [eula=yes] [localkeymanagermode=on|off]
[mixedvolumes=on|off] [masterkey=MASTERKEYSTRING] [user=crypto]
[password=PASSWORD STRING]
```

Setting	Command
Enable encryption (first time)	<pre>=> ctrl slot=7 enableencryption encryption=on eula=yes masterkey=XXXXXX localkeymanagermode=on mixedvolumes=on password=XXXXXX</pre>
Enable encryption (after initial setup)	<pre>=> ctrl slot=7 modify encryption=on eula=yes localkeymanagermode=on or =>ctrl slot=7 modify encryption=on eula=yes localkeymanagermode=off masterkey=XXXXXX</pre>
Disable encryption	<pre>=> ctrl slot=7 modify encryption=off</pre>

Express Local Encryption

This command enables initial encryption settings with defaults set. The `eula=yes` must be specified.

Target: The target can be any valid controller supporting encryption.



Example command:

```
=> ctrl slot=1 expresslocalencryption eula=yes
```

Setting the password and recovery options

These commands allow the user to set passwords and password recovery measures.

Setting a new password

Use the `suser` parameter for specifying the user account the password is being set for.

Use the `spassword` parameter to specify the new password.

The `user` and `password` parameters can be specified for the current user/password to allow this operation to be performed in a single command without having to log in first.

Password recovery settings

Set the password recovery question and answer with the `setrecoveryparams` command.

If the crypto password is lost, the `recoverpassword` command can be used to reset it.

The `user` and `password` parameters can be specified for the current user/password to allow this operation to be performed in a single command without having to log in first.

Target: The target can be any valid controller supporting encryption.

Example commands

Setting	Command
Setting the password	<pre>=> ctrl slot=7 setpasswd suser=crypto spassword=XXXXXX</pre>
Setting the password and bypassing login	<pre>=> ctrl slot=7 setpasswd suser=crypto spassword=XXXXXX user=crypto password=XXXXXX</pre>
Setting the password recovery question	<pre>=> ctrl slot=7 setrecoveryparams question=[Question string] answer=[Answer string]</pre>
Setting the password recovery question and bypassing login	<pre>=> ctrl slot=7 setrecoveryparams question=[Question string] answer=[Answer string] user=crypto password=XXXXXX</pre>
Recovering a lost password	<pre>=> ctrl slot=7 recoverpassword=> ctrl slot=7 recoverpassword answer=[Answer string] password=[New password]</pre>

-configfile command

This command allows specifying a file to read that contains lines for `key=value` pairs. Currently, only the `password=` key is valid.

Target:



The target can be any valid controller supporting encryption.

Example command:

```
ssaccli ctrl slot=3 login user=crypto -configfile [mypasswordfile].txt
```

Setting the controller password

The following commands modify controller password settings. If not currently logged in, a user name and password can be provided on the command line using the `user=` and `password=` parameters.

Target:

The target can be any valid controller supporting encryption.

Setting	Command
Setting the controller password	<pre>=> controller slot=7 setctrlpasswd spassword=XXXXXXX => controller slot=7 setctrlpasswd spassword=XXXXXXX user=crypto password=XXXXXXX</pre>
Removing the controller password	<pre>=> controller slot=7 removectrlpasswd => controller slot=7 removectrlpasswd user=crypto password=XXXXXXX</pre>
Suspending the controller password	<pre>=> controller slot=7 modify ctrlpasswd=suspend</pre>
Resuming the controller password	<pre>=> controller slot=7 modify ctrlpasswdstate=resume => controller slot=7 modify ctrlpasswdstate=resume user=crypto password=XXXXXXX</pre>

Encoding a plaintext logical drive

This command allows the user to encode an existing plaintext logical drive. If not currently logged in, a user name and password can be provided on the command line using the `user=` and `password=` parameters.

Target:

The target can be any valid controller supporting encryption.

Example command:

```
=> controller slot=7 array A encode  
=> controller slot=7 ld 1 encode preservedata=yes  
=> controller slot=7 ld 1 encode preservedata=no  
=> controller slot=7 ld 1 encode preservedata=yes user=crypto password=XXXXXX
```



Encryption key commands

The `login` command is used for authenticating either the `crypto` user or another user to the system in order to perform encryption functions.

Password and Master Key example text is indicated with "XXXXXX."

Target:

The target can be any valid controller supporting encryption.

Example commands:

Setting	Command
Setting the Master KeyCache	<code>controller slot=7 setmasterkeycache masterkeycache=[on off] [retrycnt=##] [retryinterval=##] [user=crypto] [password=XXXXXX]</code>
Rescanning keys	<code>controller slot =7 rescankeys</code>
Rekeying a controller or logical drive (multiple scenarios)	<code>controller slot=7 rekey</code> <code>controller slot=7 array A rekey</code> <code>controller slot=7 ld # rekey</code> <code>controller slot=7 ld # rekey user=crypto password=XXXXXX</code> <code>controller slot=7 ld all rekey</code>
Importing keys	<code>controller slot=7 import masterkey=[XXXXXX] [user=crypto] [password=XXXXXX]</code>
Setting or changing the Master Key	<code>controller slot=7 setmasterkey masterkey=XXXXXX</code>

Managing volatile keys

Volatile encryption keys

This command allows the user to turn on or off volatile encryption keys for a specified logical drive. Control of volatile keys is only available when the `crypto` user is logged in.

Target:

The target can be any valid logical drive that supports encryption.

Example commands:

```
controller slot=7 create type=ld drives=1e:1:1 volatileencryptionkeys=on
ctrl slot=7 ld 1 modify volatileencryptionkeys=on
ctrl slot=7 ld 1 modify volatileencryptionkeys=off
```

Volatile encryption key actions

This command sets the volatile key action. Options for this command are backup, restore, or remove.

Target:

The target can be any valid logical drive that supports encryption.

Example commands:

```
controller slot=7 ld 1 modify volatileencryptionkeyaction=backup
```



```
ctrl slot=7 ld 1 modify volatileencryptionkeyaction=restore
ctrl slot=7 ld 1 modify volatileencryptionkeyaction=remove
```

Encryption Key Manager Authentication

This command allows the user to bypass providing the controller password at boot.

The `user` and `password` parameters can be specified for the current user/password to allow this operation to be performed in a single command without having to log in first.

This command requires the following settings to be in place before continuing:

- HPE Secure Encryption must be configured to run in remote key management mode.
- The boot time controller password must be set.
- Local Key Cache must be enabled, with the number of access attempts count set at a value greater than 0.
- The Crypto Officer or an additional user must be logged in.

Target:

The target can be any valid controller supporting encryption.

Example command:

```
=> controller slot=7 modify keymanagerauthentication=on
=> controller slot=7 modify keymanagerauthentication=on user=crypto
password=XXXXXX
```

Instant secure erase

This command allows the user to instantly and securely erase an encrypted logical drive.

The `user` and `password` parameters can be specified for the current user/password to allow this operation to be performed in a single command without having to log in first.

Target:

The target can be any valid controller supporting encryption.

Example command:

```
=> controller slot=7 ld 1 modify instantsecureerase
=> controller slot=7 ld 1 modify instantsecureerase user=crypto
password=XXXXXX
```

Toggling the firmware lock

This command allows the user to lock or unlock the firmware.

If the firmware is locked, it cannot be flashed.

Target:

The target can be any valid controller supporting encryption.

Example command:

```
=> controller slot=7 modify fwlock=[on|off]
```



Unlocking volumes

This command allows the users to unlock encrypted volumes that are locked due to the controller password. Failing to supply the correct controller password during boot will cause encrypted volumes to be inaccessible until they are unlocked.

Target:

The target can be any valid controller supporting encryption.

Example command:

```
=> controller slot=7 modify unlockvolumes ctrlpassword=XXXXXX
```

Clearing an encryption configuration

This command clears the encryption configuration.

Target:

The target can be any valid controller supporting encryption.

Example command:

```
=> controller slot=7 clearencryptionconfig
```

Viewing the EULA agreement

This command allows the user to view the End User Licensing Agreement.

Target:

The target can be any valid controller supporting encryption.

Example command:

```
controller slot=7 show eula
```



HPE SSA Scripting

Using HPE SSA scripting

The HPE SSA Scripting application has two scripting modes:

- Capture mode for **capturing a configuration**

HPE SSA Scripting inspects the configuration of all internal and external array controllers connected to the server and then writes a script file describing this configuration.

- Input mode for **configuration**

HPE SSA Scripting reads the array configuration described in a specified script file. See "**Creating an HPE SSA script file**." HPE SSA Scripting then applies this configuration to a target system.

Capturing a configuration

To capture the configuration of a system, enter the following command at the system command line prompt:

```
ssascripting -c FILENAME [-internal | -external] -e FILENAME
```

The `-c` switch information, if provided, specifies the capture file name. The capture filename is optional. If no capture file is specified, HPE SSA Scripting names the default capture file `ssaoutput.ini` and places it in the HPE SSA Scripting working directory.

The `-internal` and `-external` switches limit capture to internal or external controllers.

The `-e` switch information allows the user to rename an error file. By default, HPE SSA Scripting names the error file `error.ini` and places it in the HPE SSA Scripting working directory.

Using an Input script

To use an Input script to configure or reconfigure a system, first locate a suitable HPE SSA Scripting script or see "**Creating an HPE SSA Scripting script file**."

Then, enter the following command at the system command line prompt:

```
ssascripting -i FILENAME [-internal | -external] [-reset] -e FILENAME
```

The `-i` switch information specifies the input file to use.

The `-internal` and `-external` switches limit configuration operations to internal or external controllers.

The `-reset` flag destroys any existing data and overwrites the current configuration with the configuration specified in the script.

The `-e` switch information allows the user to rename an error file. By default, HPE SSA Scripting names the error file `error.ini` and places it in the HPE SSA Scripting working directory.

Creating an HPE SSA Scripting script file

To create a valid HPE SSA Scripting script file, use one of the following methods:

- Modify the **sample custom input script**.
- Create a Capture file for **capturing a configuration**.



You can create a capture file from any server that has HPE SSA Scripting loaded, and then modify the values of the options in the file as necessary for the target system. This method is useful for applying a standard configuration to several servers that have similar storage resources.

- Write an original script.

Each line of text in an HPE SSA Scripting script file is in the format option=value and can be written in uppercase or lowercase letters. For information about possible option values and the minimum configuration information that a valid script must have, see the **sample custom input script**.

You can add blank lines and comments to any script to make it easier to read and understand. To create a comment, enter a semicolon, and then enter the comment text. HPE SSA Scripting ignores all text on the same line after a semicolon.

Script file requirements

The script file used by HPE SSA Scripting is a text file that contains options and parameters required to configure array controllers. In order to be successfully implemented, scripts must follow the guidelines below:

- Ordering of the options and parameters is important. Options need to be read in the sequence Control>Controller>Array>Logical Drive.
- Lines of the script file are written in the form option=value or left blank.
- Semicolons are used for comments within script files. If the first character in the line is a semicolon, HPE SSA Scripting ignores everything up to the next line.
- Some script options require the installation of a valid license key.

Sample custom input script

The sample script in this section gives all possible values for each option.

- If an **option** is shown in bold type, you must enter a value for that option when writing your own script.
- Default options are not universal across controllers and firmware.

You can use this script as a template for your own script.

Action = Configure|Reconfigure

Method = Custom

Controller = All|Slot [N][:N] | WWN [N] | First | SerialNumber [N] | IOCabinet [N],IOBay [N],IOChassis [N],Slot [N],Cabinet [N],Cell [N]

ClearConfigurationWithDataLoss = Yes|**No**|Forced

LicenseKey = XXXXX-XXXXX-XXXXX-XXXXX-XXXXX

DeleteLicenseKey = XXXXX-XXXXX-XXXXX-XXXXX-XXXXX

ReadCache = 0|10|20|25|30|40|50|60|70|75|80|90|100

WriteCache = 0|10|20|25|30|40|50|60|70|75|80|90|100

RapidParityInitialization = Enable|Disable

RebuildPriority = **Low**|Medium|Mediumhigh|High

ExpandPriority = **Low**|Medium|High

SurfaceScanDelay = N

SurfaceScanDelayExtended = N




```

SurfaceScanMode = Idle|High|Disabled
MNPDelay = 0..60
IRPEnable = Enable|Disable
DPOEnable = Enable|Disable
ElevatorSortEnable = Enable|Disable
QueueDepth = 2|4|8|16|32|Automatic
DriveWriteCache = Enable|Disable
NoBatteryWriteCache = Enable|Disable
PreferredPathMode = Auto|Manual
BootVolumePrimary = Logical Drive Number|None
BootVolumeSecondary = Logical Drive Number|None
HBAMode = Enable|Disable
PowerMode = MinPower|Balanced|MaxPerformance
Latency = Disable|Low|High

; Array Options
; There can be multiple array specifications in the file
Array= A|B|C|D|E|F|G|...Z|AA|AB|AC... | Next
Drive = Port:ID... | Box:Bay... | Port:Box:Bay,... | N | *
OnlineSpareMode = Dedicated | AutoReplace
OnlineSpare = Port:ID,... | Box:Bay,... | Port:Box:Bay | None | N
SplitMirror = SplitWithBackup|Rollback|Remirror|ActivateBackup

; Caching Array Options
; There can be only one Caching Array specification in the file
CachingArray = A|B|C|D|E|F|...Z|AA|AB|AC...
Drive= Port:ID,... | Box:Bay,... | Port:Box:Bay,...

; Logical Drive Options
; There can be multiple logical drive specifications in the file
; The maximum strip size depends on the number of drives in an array and the
size of the controller cache
LogicalDrive= 1|2|3... max Volumes | Next
Repeat = 0... max Volumes
RAID = 0|1|10|5|6|ADG|50|60
Size = N|Max|MAXMBR
SizeBocks = N
NumberOfParityGroups = N
Sectors = 32|63
StripSize = 8|16|32|64|128|256|512|1024

```



```

ArrayAccelerator = Enable|Disable
SSDOverProvisioningOptimization = Enable|Disable
OPTIONAL: Renumber = N
OPTIONAL: SetBootVolumePrimary = Enable
OPTIONAL: SetBootVolumePrimary = Enable
OPTIONAL: SetBootVolumeSecondary = Enable

; Caching Logical Drive Options
CachingLogicalDrive= 1|2|3... max Volumes/2
RAID = 0|1 ;FW and Controller dependent
Size = N
CachedLogicalDrive = Logical Drive Number

```

Sample custom capture script

Running HPE SSA Scripting in capture mode captures the configuration of all array controllers on the server in a capture file. The capture file can be used in input mode with minor edits to replicate the configuration to other array controllers.

Use of the "-internal" or "-external" command line switches effectively filters the controllers captured. For example, if the "-external" switch is used, then any controllers internal to the host computer are not captured in the output file.

Example:

```

; Date captured: Wed Jan 15 15:33:14 2014
; Version: 1:60:0:5

Action= Configure
Method= Custom

; Controller Specifications
; Controller HP Smart Array P800, Firmware Version 5.20
Controller= Slot 9
ReadCache= 25
WriteCache= 75
RebuildPriority= Medium
ExpandPriority= Medium
SurfaceScanDelay= 3
DriveWriteCache= Disabled
MNPDelay= 60
IRPEnable= Disabled
DPOEnable= Disabled
ElevatorSortEnable= Enabled
QueueDepth= Automatic
; Unassigned Drives

```



```

; 2E:2:1 (300 GB), 2E:2:8 (146 GB), 2E:2:9 (146 GB), 2E:2:10 (146 GB), 2E:
2:11 (300 GB), 2E:2:12 (300 GB), 2E:2:13 (146 GB), 2E:2:14 (146 GB), 2E:2:15
(146 GB), 2E:2:16 (146 GB), 2E:2:17 (146 GB), 2E:2:18 (146 GB)

; Array Specifications
Array= A
; Array Drive Type is SAS
; 2E:1:1 (146 GB)
Drive= 2E:1:1
OnlineSpare= No

; Logical Drive Specifications
LogicalDrive= 1
RAID= 0
Size= 139979
Sectors= 32
StripeSize= 128
ArrayAccelerator= Enabled

```

Sample encryption script

Initial encryption configuration

When the controller is new and unconfigured, or after the user has executed the "EncryptionClearConfig" command to delete the encryption configuration, the user must initialize the encryption parameters before first use of the controller.

```

Action= Configure
Method= Custom
Controller= SLOT 1

AcceptEULA=yes
EncryptionCryptoPasswordSet=Password.12345
EncryptionMasterKey= MyKeyName
EncryptionKeyManager= Local
AllowPlainText= True
Encryption= Enable

EncryptionUserPasswordSet= UPassword.9995 ; optional
ControllerPassword= MyBootTimePasswd ; optional
FirmwareLock= On ; optional
EncryptionRecoveryQuestion=This is my question ; optional
EncryptionRecoveryAnswer=This is my answer ; optional

```

Encryption configuration

When the controller encryption settings are configured, HPE SSA Scripting can be used to change some encryption parameters. The actual parameters to be changed vary based on the controller capabilities and the version of controller firmware.



```
Action= Reconfigure
```

```
Method= Custom
```

```
Controller= SLOT 1
```

```
; first, log in
```

```
EncryptionUser= crypto
```

```
EncryptionPassword= Password.12345
```

```
; now change the parameters
```

```
FirmwareLock = Off
```

```
AllowPlainText= false
```

When creating volumes, you can set an additional volume option to determine whether the volume will be encrypted or plaintext. You must first log in as the Crypto Officer user to use this feature. By default, volumes are encrypted.

```
Action= Configure
```

```
Method= Custom
```

```
Controller= SLOT 1
```

```
EncryptionUser= crypto
```

```
EncryptionPassword= Password.12345
```

```
Array= A
```

```
Drives= 1
```

```
LogicalDrive= 1
```

```
RAID= 0
```

```
PlainText= Enable
```

This example converts an existing plaintext volume to cyphertext, without preserving data.

```
Action= Reconfigure
```

```
Method= Custom
```

```
Controller= SLOT 1
```

```
EncryptionUser= crypto
```

```
EncryptionPassword= Password.12345
```

```
Array= A
```

```
LogicalDrive= 1
```

```
Encode= DestroyData
```

Script file options

Options in HPE SSA Scripting script files are divided into the following categories:

- **Control category**
- **Controller category**



- **Array category**
- **Logical Drive category**

Each category has several scripting options, but you do not always need to assign values to every option. HPE SSA Scripting can use default values in some instances, while in other instances, a listed option might not be relevant for a particular configuration or scripting mode.

The options for each category are listed in the following table and described in more detail in the remainder of this section.



Category	Options	Description
Control	Action	These options define the overall behavior of HPE SSA Scripting when it processes scripts and creates configurations. Control options can occur only once in a script file and must be listed first.
	Method	
Controller	BootVolumePrimary	Options in this category specify the controller that is to be configured (or the controller that had its configuration captured). Although the Controller option must begin this section of the script, you can script other options in this category in any order.
	BootVolumeSecondary	
	Controller	
	CacheState	You can use one script file to configure all controllers in a system, and you can configure the controllers identically or individually. If you define each controller configuration individually, enter the option values for one controller and its arrays and logical drives before specifying the option values for another controller.
	ClearConfigurationWithDataLoss	
	DeleteLicenseKey	
	DPOEnable	
	DriveWriteCache	
	ElevatorSortEnable	
	ExpandPriority	
	HBAMode	
	IRPEnable	
	Latency	
	LicenseKey	
	MNPDelay	
	NoBatteryWriteCache	
	ParallelSurfaceScanCount	
	PowerMode	
	QueueDepth	
	RaidArrayId	
	ReadCache	
	RebuildPriority	
	RapidParityInitialization	
	SurfaceScanDelay	
	SurfaceScanDelayExtended	
	SurfaceScanMode	
	WriteCache	

Table Continued



Category	Options	Description
Array	Array	These options describe an array that is to be configured on the controller that was previously specified in the script. (If no controller was previously specified, HPE SSA Scripting stops processing the script and creates an error file.) Although the Array option must begin this section of the script, you can script the other options in this category in any order.
	ArrayR0	
	CachingArray	
	Drive	
	DriveType	
	Join	
	OnlineSpare	
	OnlineSpareMode	
	SmartPath	
	SplitMirror	
Logical Drive	ArrayAccelerator	These options describe a logical drive that is to be configured on an array that was previously specified in the script. (If no array was previously specified, HPE SSA Scripting stops processing the script and creates an error file.) Although the LogicalDrive option must begin this section of the script, you can script the other options in this category in any order.
	LogicalDrive	
	NumberOfParityGroups	
	RAID	
	Renumber	
	Repeat	
	ResourceVolumeOwner	
	Sectors	
	ShrinkSize	
	Size	
	SizeBlocks	
	StripSize	
	StripeSize	

Control category

The Control category has the following options:

- **Action mode**
- **Method mode**

Action mode

You must specify an Action mode:



- In Configure mode, you can create new arrays, but you cannot modify existing arrays. The controller must be connected to unassigned physical drives for this mode to be available.
- In Reconfigure mode, you can modify existing arrays. For example, you can set up an array expansion, a logical drive extension, or a migration. These procedures do not destroy data, unless you specifically want the data to be deleted. In this mode, HPE SSA does not change an existing option setting unless you specifically script a different value for that option.

If you use the `-reset` command line switch, the existing controller configuration is cleared with data loss as the first step in the configuration process. This command line switch is not compatible with Reconfigure mode.

Method mode

The default value for this option is Custom.

In Auto mode, HPE SSA Scripting can perform an expansion, extension, or migration without user intervention if the values that you set for other options imply that such an operation is necessary.

Controller category

The Controller category has the following options:

- **Controller**
- **CacheState**
- **ClearConfigurationWithDataLoss**
- **DeleteLicenseKey**
- **DPOEnable**
- **DriveWriteCache**
- **ElevatorSortEnable**
- **FLS**
- **IRPEnable**
- **LicenseKey**
- **MNPDelay**
- **NoBatteryWriteCache**
- **Power modes**
- **QueueDepth**
- **RapidParityInitialization**
- **ReadCache**
- **RebuildPriority**
- **SurfaceScanDelay**
- **SurfaceScanDelayExtended**
- **Video performance options**



Controller

You must enter a value for this option because it identifies the controller that you want to configure:

- `All`—Configure all detected internal and external controllers in the system.
- `Slot N[:N]`—Configure the internal controller in slot number N. External controllers can be identified by appending the Port Number.
- `WWN N`—Configure the external controller that has the World Wide Name N.
- `SerialNumber N`—Configure the shared storage controller that has serial number N.
- `IOCabinet[N], IOBay[N], IOChassis[N], Slot[N], Cabinet[N], Cell[N]`—Configure the controller in the Integrity server that has the slot path information defined by this sequence of identifiers.

CacheState

This option enables you to flush the cache or disable flushing of the cache. Values are `FlushEnable` and `FlushDisable`.

You can use this option to prevent stale cache issues.

ClearConfigurationWithDataLoss

Clearing the configuration results in data loss because it deletes all logical volumes and arrays on the controller. If you clear a configuration, you can write commands later in the script file to create a new configuration from the liberated drive capacity.

The values for this option are as follows:

- `Yes` The configuration will be cleared. All arrays and logical drives on the controller will be deleted. This option will fail if there are mounted logical volumes in the OS.
- `Forced` The configuration will be cleared. All arrays and logical drives on the controller will be deleted.
- `No` The configuration will not be cleared. This is the default option.

DPOEnable

The `DPOEnable` option controls the behavior of the controller Degraded Mode Performance Optimization. This option is used to tune controller performance for video applications and requires the installation of a valid license key. Command options are `Enable` or `Disable`.

DriveWriteCache

This option controls the settings of the write cache for all connected physical disks. For this option, the setting is `Enable` or `Disable`. Not all physical disks or controllers support this option.

DriveWriteCache[Unconfigured|Configured]

This option controls the settings of the write cache for all configured or unconfigured logical drives. For this option, the setting is `Enable` or `Disable`. Not all physical disks or controllers support this option.

To modify the parameters in the script, you must uncomment the lines.

ElevatorSortEnable

The `ElevatorSortEnable` option controls the behavior of the controller cache write Elevator sort algorithm. This option is used to tune controller performance for video applications and requires the installation of a valid license key. Command options are `Enable` or `Disable`.



FLS

❗ **IMPORTANT:** Setting a higher level of FLS can result in a loss of throughput for some request patterns.

This command modifies the Flexible Latency Schedule feature. Command options are `Low`, `High`, or `Disable`.

Example: `Latency=[Low | High | Disable]`

IRPEnable

The `IRPEnable` option controls the behavior of the controller Inconsistency Repair Policy. This option is used to tune the controller performance for Video applications and requires the installation of a valid license key. Command options are `Enable` or `Disable`.

LicenseKey, DeleteLicenseKey

These options enable you to enter a 25-character license key to activate or uninstall some controller features. Hyphens can be entered, but are not required.

MNPDelay

The `MNPDelay` option controls the behavior of the controller Monitor and Performance Analysis Delay and is expressed in values ranging from 0 to 60. This option is primarily used to tune controller performance for video applications and requires the installation of a valid license key.

NoBatteryWriteCache

With this option, the controller can enable write cache when no energy pack is present or when the energy pack fails. Values are `Enable` or `Disable`. The default is `Disable`.

Some controllers do not support this option.

ParallelSurfaceScanCount

The `ParallelSurfaceScanCount` command specifies the number of parallel surface scans which can be active at any time. The minimum and maximum values are controller and firmware release dependent. This feature is not supported by all controllers.

Example: `ParallelSurfaceScanCount = N ;number of scans`

Power modes

❗ **IMPORTANT:** A reboot may be required after switching power modes in order to optimize savings and performance.

❗ **IMPORTANT:** When the power mode is set to `Balanced`, future controller configuration changes may require a reboot for optimal performance.

This option modifies the controller power mode.

- A value of `minpower` sets static settings to the lowest possible value and dynamically reduces power based on the workload.
- A value of `balanced` sets statistic settings based on the configuration and dynamically reduces power based on the workload (recommended).
- A value of `maxperformance` sets static settings to the highest possible value and does not dynamically reduce power (default).

Example: `PowerMode= [MinPower | Balanced | MaxPerformance]`



QueueDepth

QueueDepth controls the behavior of the cache write queue. This option is used to tune controller performance for video applications. Command values are 2, 4, 8, 16, 32, or Automatic.

RapidParityInitialization (RPI)

This option enables Rapid Parity Initialization. The RPI command is not supported on all controllers and is dependent on firmware and hard disk support. Command options are `Enable` or `Disable`.

ReadCache, WriteCache

Enter a number between 0 and 100 to specify the percentage of cache that is to be allocated to drive reads or writes. The default values for both options are controller firmware-dependent.

The allowable cache ratios depend on the controller model and whether it has battery-backed write cache.

RebuildPriority, ExpandPriority

These options establish the priority for rebuild and expand functions. Rebuild has four possible values: Low, Medium, MediumHigh, and High. MediumHigh is only available when Rapid Parity Initialization is enabled.

Expand has only three possible values: Low, Medium, and High.

These options are not required.

SurfaceScanDelay

Enter a number between 0 and 30 to specify the duration of the surface scan delay in seconds. This option is not required. If you do not specify a value for this option, the delay remains unchanged. A value of 0 disables the scan.

SurfaceScanDelayExtended

Enter a number between 0 and 300 to specify the duration of the surface scan delay in tenths of seconds. This option is not required. If you do not specify a value for this option, the delay remains unchanged. A value of 0 disables the scan.

If this parameter and SurfaceScanDelay are both in the input file, this parameter takes precedence.

SurfaceScanMode

This parameter specifies the Surface Scan Mode with the following values:

- `Idle`—The delay interval is set with the normal values from `SurfaceScanDelay` or `SurfaceScanDelayExtended`.
- `High`—The surface scan enters a mode guaranteed to make progress despite the level of controller I/O.
- `Disabled`—No mode is selected.

Video performance options

To optimize the controller performance for video data storage, set values for the following options as indicated:

`DPOEnable = NoElevatorSortEnable = YesIRPEnable = No`

In addition, do the following:

- Set the `MNPDelay` to any integer value from 1 to 60 (units are minutes). If you want to disable this option, set a value of zero instead.
- Set the `QueueDepth` to one of the following values: `2 | 4 | 8 | 16 | 32 | Automatic`

Array category

The Array category has the following options:



- **Array**
- **ArrayR0**
- **CachingArray**
- **Drive**
- **DriveType**
- **OnlineSpareMode**
- **OnlineSpare**
- **SplitMirror**

Array

Enter a letter or pair of letters to identify the array that is to be created or reconfigured, and observe these additional limitations:

- In Configure mode, HPE SSA Scripting creates a new array. The value that you specify for the array option must be the next available letter or pair of letters in the sequence, according to the number of existing arrays on the controller. AA follows Z, and BA follows AZ.
- In Reconfigure mode, HPE SSA Scripting can either create a new array or reconfigure an existing array. In this case, the value that you specify can identify an existing array, or it can correspond to the next available array letter or pair of letters in the existing configuration.

ArrayR0

The option ArrayR0 creates a single drive array with a RAID0 volume for each argument. When ArrayR0 is specified, DriveType is the only other valid command in the array section.

Command options:

- X:Y, . . . : Port:Id for controllers that use Port/Id drive numbering schemes. For Box/Bay numbering schemes, use Box:Bay.
- X:Y:Z, . . . : Port:Box:Bay for SAS controllers
- N: N number of drives will be used
- *: Use all unconfigured drives

CachingArray

This option specifies the Caching Array ID.

- In Configure mode, HPE SSA Scripting creates a new caching array. The value that you specify for the array option must be the next available array letter in the existing configuration. AA follows Z, and BA follows AZ.
- In Reconfigure mode, the array letter identifies an existing array, or it can identify the next available array letter in the existing configuration to create a new array.

Drive

You can use this option in the input file to specify new physical drives for the array. Use this option to build a new array or to expand, shrink, or move an existing array.

Observe the following guidelines:



- If you are expanding an array, each drive that you add must have a capacity no less than that of the smallest drive already in the array. The added drives and the existing drives in the array must all be the same type (for example, SAS or SATA).
- If you are moving or shrinking an array, HPE SSA Scripting compares the current drives with the requested drives, and then determines whether you are moving or shrinking. Shrinking or moving an array is supported only in Custom method mode.
- If the value of the **ClearConfigurationWithDataLoss** option is Yes, you can use the `Drive` option to remove drives from an array.

Custom method mode—To specify the drives to be used in the array (different arrays on the same controller can use different methods), choose one of the following methods:

- To specify individual drives, use the applicable convention (port:ID, box:bay, or port:box:bay).
- To specify only the number of drives to use (not which specific drive IDs to use), enter that number as the value for this option. For example, if you enter `drive=3`, HPE SSA Scripting uses the first three available drives to build or expand the array that you define in the remainder of the script. HPE SSA Scripting automatically determines which drives are suitable to use.
- To use all available drives, enter an asterisk as the value for this option. An array that is configured using this method cannot have a spare.

DriveType

This option specifies the interface type for drives that HPE SSA Scripting must use to build the array.

Choose one of the valid drive types:

[SCSI | SAS | SATA | SATASSD | SASSSD]

The value is often used with the wildcard (*) or number argument for `Drive`.

OnlineSpareMode

The value for this option specifies the behavior of the spare drives when they are activated.

For controllers that support Auto Replace Drive, setting this value to **AutoReplace** enables a rebuilt spare drive to become a data drive in the array. When the failed data drive is replaced, it assumes the role of the former spare, eliminating the need for a second array rebuild.

Behavior	Description
Dedicated	Default value for spares
AutoReplace	The spare drive becomes a data drive when the rebuild is complete. AutoReplace spares cannot be shared across arrays.

OnlineSpareMode is only supported on select Smart Array Px2x controllers and later.

OnlineSpare

The value for this option determines whether the array specified previously in the script will be configured with spare drives.

The interface type for drives and spare drives must match, for example, all SAS or all SATA.



- `X:Y, . . .`: Port:Id for controllers that use Port/Id drive numbering schemes. For controllers using Box/Bay numbering schemes, use Box:Bay.
- `X:Y:Z, . . .`: Port:Box:Bay for SAS controllers
- `None`: No spares will be added to the array and any existing spares on the array will be removed.
 - `Configure mode`: if the `OnlineSpare` option is not specified, the default value is `None`.
 - `Reconfigure mode`: If the `OnlineSpare` option is not specified, the spare state of the array will remain unchanged.
- `N`: N number of spares will be added to the array. HPE SSA Scripting will choose the best available, unused physical drives for the spares.

SmartPath

Hewlett Packard Enterprise SSD Smart Path enables I/O requests to bypass Smart Array firmware and access SSDs directly. This process accelerates reads for all RAID levels and writes for RAID 0.

The controller device driver analyzes each I/O request for the following Smart Path eligibility requirements:

- Read requests for all supported RAID levels
- Write requests for all RAID 0 volumes
- Size and location parameters, resulting in an operation that can be satisfied by a single physical disk I/O

The command values are `Enable` or `Disable`.

Example:

```
SSDSmartPath= Enable | Disable
```

SplitMirror

This option splits the original mirrored volumes into individual RAID0 volumes, optionally hiding the new volume from the OS. `SplitMirror` also recombines volumes. The command will only execute if scripting is started with the `-offline` command line option and the join target volumes are not in use by the OS.

`SplitMirror` operates on the following volumes:

- RAID 1
- RAID 10
- RAID 1 ADM
- RAID 10 ADM

`SplitMirror` uses the following commands:

- `SplitWithBackup`: The mirrored array is split into two new arrays and the newly created array is hidden from the OS.
- `Rollback`: The two arrays are recombined using the backup array as source data. Any changes made to the online array are lost.
- `Remirror`: The two arrays are recombined using the online array as source data.
- `ActivateBackup`: The backup volume is made visible to the OS.



Logical Drive category

The Logical Drive category has the following options:

- **ArrayAccelerator**
- **LogicalDrive**
- **CachingLogicalDrive**
- **CachedLogicalDrive**
- **NumberOfParityGroups**
- **RAID**
- **Renumber**
- **Repeat**
- **SetBootVolumePrimary**
- **SetBootVolumeSecondary**
- **Sectors**
- **Size**
- **SizeBlocks**
- **SmartCache Write-back**
- **SSDOPO**
- **StripSize**

ArrayAccelerator

This option specifies whether the array accelerator is enabled or disabled for the specified logical drive. The default value is Enabled.

CachingLogicalDrive

The value that you enter for this option specifies the ID number of the caching logical drive that is to be created or modified. The first logical drive on an array must have an ID of 1 (not 0), and logical drive numbering must be contiguous.

- In Configure mode, HPE SSA Scripting accepts only the ID number of the next possible logical drive.
- In Reconfigure mode, HPE SSA Scripting also accepts the ID number of any existing logical drive.

Caching logical drives and cached data drives must be on the same Array Controller.

CachedLogicalDrive

The value that you enter for this option specifies the data logical drive ID number to associate with the Caching logical drive.

In Configure mode, HPE SSA Scripting accepts the ID number of any existing logical drive.

Caching logical drives and cached data drives must be on the same Array Controller.

LogicalDrive

The value that you enter for this option specifies the ID number of the logical drive that is to be created or modified. The first logical drive on an array must have an ID of 1 (not 0), and logical drive numbering must be contiguous.



- In Configure action mode, HPE SSA Scripting accepts only the ID number of the next possible logical drive.
- In Reconfigure action mode, HPE SSA Scripting also accepts the ID number of any existing logical drive.

NumberOfParityGroups

When you create a RAID 50 or RAID 60 configuration, you must also set the number of parity groups.

You can use any integer value greater than 1 for this setting, with the restriction that the total number of physical drives in the array must be exactly divisible by the number of parity groups.

The maximum number of parity groups possible for a particular number of physical drives is the total number of drives divided by the minimum number of drives necessary for that RAID level (three for RAID 50, four for RAID 60).

RAID

The value that you enter for this option specifies the RAID level of the logical drive.

- When the Action mode is Configure, HPE SSA Scripting automatically selects the highest RAID level that the controller and drive configuration can support except RAID 50 or RAID 60. To specify RAID 50 or 60 for a controller that supports either of these RAID levels, use the Custom setting. In this case, you must also specify the **number of parity groups**.
- When the Action mode is Reconfigure, the default value is the existing RAID level for that logical drive. If you specify a different RAID setting, then HPE SSA Scripting either ignores the new setting, or attempts to migrate the logical drive to the specified RAID level (when Method mode is Custom).

HPE SSA Scripting supports the following values for RAID levels:

- 60—RAID 60
- 50—RAID 50
- ADG—RAID ADG is the equivalent of RAID 6
- 6—RAID 6
- 5—RAID 5
- 10ADM—RAID 10 with 3-way mirroring
- 10—RAID 10 (mirroring with 2+ disks)
- 1ADM—RAID 1 with 3-way mirroring
- 1—RAID 1 (mirroring with 2 disks)
- 0—RAID 0

Renumber

This option rennumbers the logical drive to N.

Typically used after a `Join` command, this option ensures the proper logical volume numbering is set, for example, the boot volume is ID 1.

Do not mix `Renumber` and `Join` operations in the same script, because HPE SSA Scripting does not update its internal state after a `renumber`. If other commands are mixed with `Renumber`, commands targeting specific volumes might target the incorrect volume.



Repeat

The value you enter for this option specifies the number of times that HPE SSA Scripting is to repeat this logical drive configuration.

Use one of the following values:

- **N**—In Configure mode, HPE SSA Scripting creates N new logical drives.
- **MAX**—HPE SSA Scripting creates the maximum number of logical drives possible. The number of drives created depends on the number of existing drives and the maximum number of logical drives supported by the controller.

You must specify the logical drive ID as **Next**. This **Size** option controls the size of each logical drive, or if the size is set to **MAX**, the size of the volumes is set to consume all available space on the array.

SetBootVolumePrimary

This option sets the current logical volume to be the primary boot volume on the current controller through the command **Enable**. This option is not supported on all array controllers.

SetBootVolumeSecondary

This option sets the current logical volume to be the secondary boot volume on the current controller through the command **Enable**. This option is not supported on all array controllers.

Sectors

This option specifies the number of sectors that are to comprise each track. Enter 32 to disable MaxBoot or 63 to enable it.

- For new logical drives, the default setting is 32.
- For an existing logical drive, the default setting is the existing setting.

Logical drive performance is likely to decrease with MaxBoot enabled.

Size

This parameter specifies the size of the logical drive.

Use the following values to set the size:

- **N**—The size in MB
- **MAX**—Uses all available unused space in the array for the logical drive. This value is the default.
- **MAXMBR**—Creates the largest size of volume supportable by a 32-bit MBR (2TiB)

In Reconfigure mode, the default setting is the existing size of the logical drive. If you enter a larger value, HPE SSA Scripting extends the logical drive to the new size if there is unused drive capacity on the same array, as long as the operating system supports logical drive extension. You cannot reduce the size of the logical drive.



CAUTION: Back up all data before extending a logical drive.

SizeBlocks

This option specifies the size of the logical drive in 512- or 4096-byte blocks. Use this option to set an exact size without rounding.

When HPE SSA Scripting captures a configuration, it reports volume sizes rounded down to the nearest MB. When the capture is played back, the size is reduced to fit within the already rounded MB size. This option prevents the loss of volume size over successive generations of captures.



However, if the size does not fit the geometry of the array, caused by different disks or RAID levels, the size will be rounded down to fit the geometry.

SmartCache Write-back

⚠ CAUTION: Specifying the write-back cache write policy could result in data loss in the event of failure of the cache volume. When using a RAID 0 cache volume, a single SSD failure might result in data loss.

❗ IMPORTANT: If a demo license key expires, all SmartCache volumes configured with write-back cache write policy are converted to write-through. When this happens, the logical drive details show different values for Cache Write Policy and Cache Write Policy Requested. If the license is reinstalled, the SmartCache volumes are restored to their original write-back cache write policy.

This option enables write-back cache support. Command options are `Writethrough` or `Writeback`.

Example: `CacheWritePolicy=[Writethrough | Writeback]`

SSDOPO (SSDOverProvisioningOptimization)

This option sets the over provisioning optimization parameter at volume creation time through the command `Enable`. This option is not supported on all array controllers, and only supported on some SSDs. This option is not valid in reconfigure mode.

StripSize

Beginning with HPE SSA, HPE SSA CLI, and HPE SSA Scripting version XXXX, the `StripSize` option replaces the `StripeSize` option. When distributing data across multiple physical drives (striping), the **strip size** is the amount of data that is written to each physical drive. The **full stripe size** refers to the combined size of all the strips across all physical drives, excluding parity-only drives.

You can enter a numerical value for this option to specify the size of the data strips (in kilobytes), or you can leave this option blank and allow HPE SSA Scripting to use a default value.

Strip sizes available for a RAID level vary based on the controller and the controller firmware level. The maximum strip size changes dynamically and is reduced for arrays with a large number of data drives or with smaller controller cache sizes. The controller must be able to read an entire strip of data at a time into cache memory during transformation. Available memory is the limiting factor.

The following strip sizes are available:

- 8—8 KB
- 16—16 KB
- 32—32 KB
- 64—64 KB
- 128—128 KB
- 256—256 KB
- 512—512 KB
- 1024—1024 KB

The default strip size value depends on the action mode:



- In Configure action mode, the default value is determined by the RAID level that you specified earlier in the script.
- In Reconfigure action mode, the default value for this option is the strip size that is already configured for the logical drive. If you enter a value that is different from the existing strip size, HPE SSA Scripting attempts to migrate the logical drive to the strip size that you specify. (If you intend to migrate the logical drive, back up all data before starting the migration procedure.)

Encryption support commands

-
- ❗ **IMPORTANT: HPE Special Reminder:** Before enabling encryption on the Smart Array controller module on this system, you must ensure that your intended use of the encryption complies with relevant local laws, regulations and policies, and approvals or licenses must be obtained if applicable.

For any compliance issues arising from your operation/usage of encryption within the Smart Array controller module which violates the above mentioned requirement, you shall bear all the liabilities wholly and solely. HPE will not be responsible for any related liabilities.

- ❗ **IMPORTANT:** The following commands require the purchase and configuration of Secure Encryption before executing.
-

The following commands have been added to HPE SSA Scripting versions 1.60.0.0 and later to support data encryption features on compatible controllers.

Encryption controller category

AllowPlainText

This command allows future plaintext volumes to be created when encryption is enabled. Existing volumes are not modified. To toggle this command, use `Yes` or `No` values.

AcceptEULA

The command `Yes` accepts the Encryption End User License Agreement. The EULA must be accepted before Encryption can be enabled.

ControllerPassword

This command sets the boot time controller password. Enter the password as the value.

ControllerPasswordMode

This command sets the boot password mode, allowing it to be suspended or deleted.

This command has the following values:

- `Remove`
- `Suspend`
- `Resume`

Encryption

This command enables or disables controller encryption through the values `Enable` or `Disable`. Before encryption can be enabled, the following values must be set:

- **Crypto user password**
- **Key management mode**



- **Initial Master Encryption Key**

- **EULA** (accepted)

EncryptionClearConfig

This command deletes all secrets and Critical Security Parameters from the controller. This command does not delete or modify any existing logical volumes. However, if the volumes are encrypted, they will be placed into an offline state by the controller.

If any encrypted logical volumes exist on the controller, this command must be run offline.

To run this command, enter the value `KeyName`. The `KeyName` is required if encrypted logical volumes exist on the controller. Otherwise, this argument can be omitted.

EncryptionCryptoPasswordSet

This command sets the encryption password for the Crypto user role. The command value is the password.

EncryptionKeyManager

This command sets the Key Manager mode to either Local or Remote. If set to Remote Key Manager mode, ensure the external configuration is set up properly on the server's iLO configuration screens, and the system configuration is added to the remote key manager.

Command values:

- Local key manager mode: `Local`
- Remote key manager mode: `Remote`

EncryptionKeyManagerAuthentication

When enabled, `EncryptionKeyManagerAuthentication` allows the firmware to bypass prompting the user for the boot time controller password.

This command requires the following settings to be in place before continuing:

- HPE Secure Encryption must be configured to run in remote key management mode.
- The boot time controller password must be set.
- Local Key Cache must be enabled, with the number of access attempts count set at a value greater than 0.
- The Crypto Officer or an additional user must be logged in.

Use the following commands to toggle `EncryptionKeyManagerAuthentication`:

- `Enable`
- `Disable`

EncryptionRekeyVolume

This command rekeys the logical volume when the value `Yes` is entered. This results in a volume transformation.

EncryptionMasterKey

This command sets the initial Master Encryption Key name. The command value is the key name.



EncryptionPassword

This command provides the encryption password for the role selected under **EncryptionUser**. The command value is the user password.

EncryptionRecoveryQuestion

This command sets the controller recovery question. The controller recovery question is used when the user has forgotten the Crypto user password and wishes to change it. Enter the question as the value.

EncryptionRecoveryAnswer

This command sets the controller password recovery answer. This is the answer to the question set through the **EncryptionRecoveryQuestion command**. Enter the answer as the command value.

EncryptionRekeyController

This command rekeys the controller when the value `Yes` is entered.

EncryptionUser

The commands `User` or `Crypto` set the encryption user role. The selected role determines which encryption commands are available. Currently the only two roles supported are `User` and `Crypto`.

EncryptionUserPasswordSet

This command sets the encryption password for the `User` role. The command value is the password.

ExpressLocalEncryption

`ExpressLocalEncryption` performs a setup of the controller encryption configuration.

`Enable`

The following items are enabled:

- Controller encryption
- Local key manager mode
- Random crypto password, not recoverable
- Random master key name, not recoverable
- Future plaintext volumes not allowed

To disable express encryption, the entire encryption configuration must be cleared, which will require deleting any configured encrypted volumes.

FirmwareLock

This command toggles the array controller firmware through the values `Lock` and `Unlock`. Firmware must be unlocked prior being updated.

Encryption logical volume category

InstantSecureErase

This command performs an instant cryptographic erase on the current logical volume. All user data is lost upon entering the command value `Yes`.

Encode

This command converts a plaintext logical volume to an encrypted volume, optionally converting the existing data or destroying it. The command values are `PreserveData` or `DestroyData`.



If `PreserveData` is entered, a volume transformation occurs.

EncryptionVolatileKeys

`EncryptionVolatileKeys` configures the volatile key status of the selected volume. Keys are not automatically backed up when enabling this feature. If the script is enabled and the system experiences a power outage, the encryption keys will be lost and there will be no access to the volume.

- `Enable`
- `Disable`

EncryptionVolatileKeyAction

`EncryptionVolatileKeyAction` modifies the status of the volatile keys for the volume. Volatile keys must be enabled and the controller configured to run in remote key management mode. These operations may take some time, depending on the workload and responsiveness of the remote key manager.

- `Backup`—Backs up the current encryption key to the remote key manager
- `Restore`—Restores the current encryption key from the remote key manager. The key must be backed up before it can be restored.
- `Remove`—Removes the current encryption key from the remote key manager. The key must be backed up before it can be removed.

Plaintext

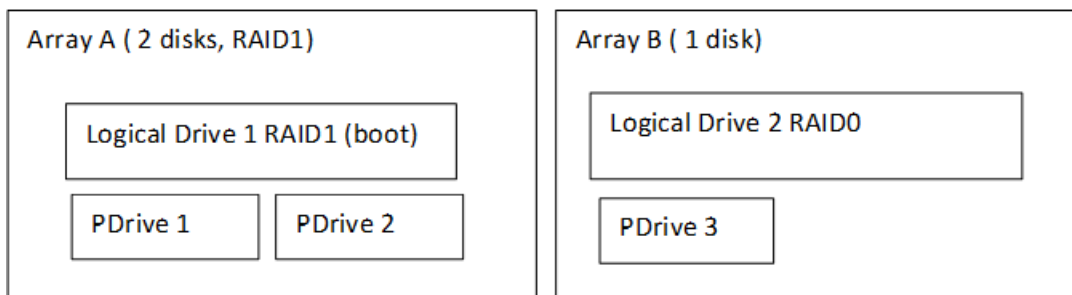
This command specifies whether to create a new logical volume in either plaintext or encrypted form. To create a plaintext volume when encryption is enabled, the `AllowPlainText` option must have been set during initial configuration of the array controller. The command values are `Enable` or `Disable`, in reference to whether encrypted forms are enabled. The default form is encrypted, or enabled.

Split/Join scripting

The `SPLITMIRROR` command enables you to split or join arrays and their associated drives. It is possible to execute the command only after starting the script with the `-offline` command in the CLI. The target volumes for the join operation cannot be in use by the operating system.

Initial configuration example

In this example, two arrays are each set up with a logical drive, and three physical drives (PDrives) between them.



Performing the split operation

Use the following syntax for the Split operation:

`SplitMirror = SplitWithBackup`

When the command executes, the mirrored array is split into two new arrays with one becoming a backup array.



In the example, Backup Array C will be hidden from the operating system. To split the RAID1 logical drive in the example configuration, the command would be used as follows, with the controller in Slot 6:

Action= Reconfigure


Method= Custom

Controller= SLOT 6

Array= A

SplitMirror= SplitWithBackup

The `SPLITMIRROR` command creates a new logical drive (3) and a new array (C) from the original Array. The configuration is shown in the following diagram:



Executing the join command

Scripting must be launched with the `-offline` command line option to enable this command. This command has three options:

- SplitMirror = Remirror
- SplitMirror = Rollback
- SplitMirror = ActivateBackup

Example 1: SplitMirror = Remirror

This form of the command does the following:

- Recreates the original array
- Retains data from Logical Drive 1
- Removes Logical Drive 3, along with all data

To keep the modifications made to Logical Drive 1 after the `SPLITMIRROR`:

Action= Reconfigure

Method= Custom

Controller= SLOT 6

Array= A

SplitMirror= Remirror

After the command executes, Array C is removed and physical drive 2 (PDrive2) is now a mirror of PDrive1, with the original contents of PDrive2 lost.

Example 2: SplitMirror = Rollback

This form of the command does the following:



- Recreates the original array
- Retains data from Logical Drive 3 (that is, the original data at the time of the split)
- Removes Logical Drive 1, along with all data

To return to the original configuration at the time of the SPLITMIRROR:

Action= Reconfigure

Method= Custom

Controller= SLOT 6

Array= A

SplitMirror= Rollback

After the command executes, physical drive 1 (PDrive1) becomes a mirror of PDrive2 (the original data at the time of the split), with the modifications to PDrive1 lost.

If this example is attempted while the system is online (that is, booted from Logical Drive 1), a blue-screen or kernel panic will likely result. Perform this command offline and ensure the volumes involved in the join are not mounted or otherwise used by the Operating System.

Example 3: SplitMirror = ActivateBackup

- Exposes the hidden logical drive to the operating system
- Retains the data from Logical Drive 1
- Retains the data from Logical Drive 3

To keep the modifications made to Logical Drive 1 and also access the original logical volume data at the time of the SPLIT, enter the following:

Action= Reconfigure

Method= Custom


Controller= SLOT 6

Array= A

SplitMirror= ActivateBackup

After the command executes, Logical Volume 3 is exposed to the Operating System.

There might be implications with the operating system with respect to the volume name, unique IDs, and so forth, which may require manual intervention by the System Administrator.



Basic scripting scenarios

Configure a two-disk array

Task



Create a simple two-disk array with a 100 GB RAID 1 logical volume.

Hardware configuration

- Storage system
 - An unconfigured HPE Smart Array controller in slot 1
- If the controller contains an existing configuration, the script fails unless you add the -reset parameter to the command line to clear the existing configuration.
- Two available physical drives at the following location: 11:4:8 and 11:4:7

Script

```
Action=Configure
Method=Custom
Controller=Slot1
Array=A
Drive=11:4:8, 11:4:7
LogicalDrive=1
RAID=0
Size=100000
```

Summary

The script combines the two physical drives into a single 100 GB logical drive.

Reconfigure to modify logical drive size

Task

Reconfigure an existing 100 GB logical volume to modify its size to 200 GB.

Hardware configuration

- Storage system
- An HPE Smart Array in slot 1
- An existing logical drive with a capacity equal to or larger than 200 GB

Script

```
Action=Reconfigure
Method=Custom
Controller=Slot1
Array=A
LogicalDrive=1
RAID=0
Size=200000
```

Summary



The script modifies the size of the logical drive only, because multiple transformations of a logical drive cannot be performed at the same time. For example, HPE SSA Scripting must complete the size expansion before beginning another transformation, such as a RAID level change.

To modify existing arrays and volumes, the HPE SSA GUI may be a more intuitive tool for some users. The GUI provides graphical feedback that guides the user and displays only available options.

Reconfigure to modify RAID level

Task

Reconfigure an existing RAID 0 logical volume to be a RAID 1 logical volume.

The selected RAID level is supported only if the minimum number of physical drives are available.

Hardware configuration

- Storage system
- An HPE Smart Array in slot 1
- An existing logical drive

Script

```
Action=Reconfigure
```

```
Method=Custom
```

```
Controller=Slot1
```

```
Array=A
```

```
LogicalDrive=1
```

```
RAID=1
```

Summary

The script modifies the RAID level of the logical drive only, because multiple transformations of a logical drive cannot be performed at the same time. For example, HPE SSA Scripting must complete the RAID level change before beginning another transformation, such as a size expansion.

To modify existing arrays and volumes, the HPE SSA GUI may be a more intuitive tool for some users. The GUI provides graphical feedback that guides the user and displays only available options. For example, the HPE SSA GUI only displays optional RAID levels based on the minimum number of required physical drives.

Configure by disk type and count

Task

Create an array by specifying the type and count of disks for HPE SSA Scripting to use. This example use two SAS drives. This method does not require you to specify physical disks by location.

Hardware configuration

- Storage system
- An HPE Smart Array in slot 1
- Two available SAS physical disks

This script fails if two unused SAS drives are not available on the Smart Array.

Script

```
Action=Configure
```



```
Method=Custom
Controller=Slot1
Array=A
Drive=2
DriveType=SAS
LogicalDrive=1
RAID=0
```

Summary

This script does not specify a size, but uses the `MAX` default value to consume the fully available capacity of the physical drives.

HPE SSA Scripting does not support arrays with mixed drive types. An array must comprise drives of the same type, for example all SAS or all SATA drives. Different arrays can have different drive types.

Configure two arrays, one with a spare

Task

Create two arrays with one online spare on a Smart Array. One array uses two physical drives in a RAID 1 logical volume. The second array uses the remaining physical disks to create 20 identically sized RAID 5 logical volumes.

Hardware configuration

- Storage system
- An unconfigured HPE Smart Array in slot 1

If the controller contains an existing configuration, the script fails unless you add the `-reset` parameter to the command line to clear the existing configuration.

- Numerous available physical drives at unspecified locations

Script

```
Action=Configure
Method=Custom
Controller=Slot1
Array=A
Drive=2
OnlineSpare=1
LogicalDrive=1
RAID=1
Size=MAX
Array=B
Drive=*
LogicalDrive=NEXT
Repeat=20
RAID=5
Size=MAX
```



Summary

For both arrays, this script uses the MAX value to consume the fully available capacity of the physical drives.

For Array B, the Repeat option specifies the number of times to repeat the configuration. Instead of using the number 2, the script uses the NEXT value to specify the ID of the second logical drive. When using the Repeat option, you must also use the NEXT value for the LogicalDrive option.

This script will fail if the controller does not have the required number of disks connected.

Advanced commands

Controller=ALL

You can use Controller=All instead of the Slot=N command.

Advantages

- You do not need to determine where controllers are installed.
- You can target all controllers at once.

Examples

Change read cache to 100% on all controllers (only configured controllers, those with volumes, accept this command):

```
Controller=ALL
```

```
ReadCache=100
```

Clear the configuration on all controllers:

```
Controller=ALL
```

```
ClearConfigurationWithDataLoss=Yes
```

Controller=First

You can use Controller=First instead of the Slot=N command.

Advantages

- You do not need to determine where controllers are installed.
- You can target the first controller found, a common practice with embedded controllers, because they are always slot 0. (Controllers are sorted by slot number, lowest to highest, and internal controllers before external ones.)

Example

Configure a volume on the embedded controller (or the controller in the lowest slot number, if an embedded controller does not exist):

```
Controller=First
```

```
Array=A
```

```
Drive=1E:1:1, 1E:1:2
```

```
LogicalDrive=1
```

```
RAID=1
```



-internal or -external

You can use the command line option `-internal` or `-external` instead of the default option, which is both internal and external controllers.

Advantage

This option filters the discovered controllers, restricting the command to only internal or only external controllers.

Example

When using the `Controller=First` command, if the system had no internal controllers (due to hardware failure or discovery failure) but had an external controller for shared storage, the script would attempt to create a volume on the external controller. To prevent that action, the `-internal` command line argument restricts discovery to only internal controllers:

```
# ssascripting -i input.txt -internal
Controller=First
Array=A
Drive=1E:1:1, 1E:1:2
LogicalDrive=1
RAID=1
```

Array=Next

You can use `Array=Next` instead of the `Array=A` command.

Advantage

You do not need to know how many arrays exist. The scripting software selects the next valid array identifier.

Example

Add one array to the configuration, regardless of how many arrays already exist:

```
Controller=Slot1
Array=Next
Drive=1E:1:1, 1E:1:2
LogicalDrive=2
RAID=1
```

Hewlett Packard Enterprise recommends pairing the `Array=Next` command with **`LogicalDrive=Next`**, because you would not create an array without also creating a logical drive, and this additional command does not require you to know how many logical drives already exist.

LogicalDrive=Next

You can use `LogicalDrive=Next` instead of the `LogicalDrive=N` command.

Advantage

You do not need to know how many logical drives exist. The scripting software selects the next valid identifier.

Example

Add one logical drive to Array A, regardless of how many logical drives already exist:

```
Controller=Slot1
Array=A
Drive=1E:1:1, 1E:1:2
```



```
LogicalDrive=Next
```

```
RAID=1
```

Repeat=N | MAX

You can use this command as part of the logical drive specification.

Advantage

You can quickly divide the available space on an array into multiple volumes without specifying each volume individually.

Example

Create the maximum number of logical volumes the controller supports (each volume is the same size, and the volumes, in total, use all available free space):

```
Controller=Slot1
```

```
Array=A
```

```
Drive=*
```

```
LogicalDrive=Next
```

```
Repeat=Max
```

```
RAID=1
```

Drive=N

You can use `Drive=N` instead of specifying the actual paths (`Drive=1E:1:1, 1E:1:2`).

Advantage

You can specify the number of drives needed for the configuration, and the scripting software selects the first available drives based on increasing path. This command is most useful when all drives in the configuration are the same size.

Example

Add a logical drive to the configuration using four physical drives:

```
Controller=Slot1
```

```
Array=A
```

```
Drive=4
```

```
LogicalDrive=1
```

```
RAID=5
```

Drive=*

You can use `Drive=*` instead of specifying the actual paths (`Drive=1E:1:1, 1E:1:2`).

Advantage

You can use all available drives for the configuration without knowing their number or actual path.

For this command to execute, all configuration rules must be met. For example, if only two physical drives exist, the scripting software uses both drives, but does not create volumes that require more than two drives, such as a RAID 5 volume.

Example

Add a logical drive to the configuration using all unassigned drives:

```
Controller=Slot1
```

```
Array=A
```



```
Drive=*
LogicalDrive=1
RAID=5
```

In this example, three or more physical drives must be available to support the RAID level.

OnlineSpare=N

You can use OnlineSpare=N instead of specifying the actual paths (OnlineSpare=1E:1:1, 1E:1:2).

Advantage

You can specify the number of spares needed for the configuration, and the scripting software selects the first available drives based on increasing path. This command is most useful when all drives in the configuration are the same size.

Example

Add a logical drive to the configuration using two spare drives:

```
Controller=Slot1
Array=A
Drive=1E:1:1, 1E:1:2
OnlineSpare=2
LogicalDrive=1
RAID=1
```

OnlineSpare=*

You can use OnlineSpare=* instead of specifying the actual paths (OnlineSpare=1E:1:1, 1E:1:2).

Advantage

You can use all available, unassigned drives as spares for the configuration without knowing the number of drives available or their actual paths.

Observe the following guidelines:

- You cannot exceed the controller settings for maximum number of spares allowed in an array, typically eight spares.
- You cannot combine this command with **Drive=***, because the script processes the Drive specification first and uses all available drives for that action, leaving no unassigned drives for the OnlineSpare request.

Example

Add a logical drive to the configuration using all unassigned drives as spares:

```
Controller=Slot1
Array=A
Drive=1E:1:1, 1E:1:2
OnlineSpare=*
LogicalDrive=1
RAID=1
```

DriveType=SCSI | SAS | SATA | SATASSD | SASSSD

You can specify what type of drive to use in a configuration.



Advantage

You can ensure that only a certain type of drive is used in a configuration where multiple drive types exist.

Example

Create one volume using all existing SAS solid state drives, and create one volume using all existing SATA drives:

```
Controller=Slot1
Array=A
Drive=*
DriveType=SASSSD
LogicalDrive=1
RAID=0
Array=B
Drive=*
DriveType=SATA
LogicalDrive=2
RAID=0
```

Size=MAXMBR

You can use this command to limit the size of the volume to 2^{32} blocks (2 TiB), which is the largest bootable volume supported on MBR-based systems.

Advantage

You configure the largest possible, bootable volume size regardless of the free space available, provided that the free space is at least 2 TiB.

Example

Create an array with a max sized bootable volume from a 3 TB disk at 1l:1:1:

```
Controller= Slot 1
Array= Next
Drive= 1l:1:1
LogicalDrive=Next
RAID= 0
Size= MAXMBR
```

SizeBlocks=N

You can use this option to capture the exact number of 512-byte blocks in a volume, and then to recreate the exact size when playing back a capture.

Advantage

The `Size` value is rounded to the nearest Mebibyte. When you capture the same system repeatedly, this rounding can result in the eventual loss of a small amount of volume size. This option eliminates the rounding.

Example

Create an array with volume using exactly 20,971,520 512-byte blocks (exactly 10 GiB) using the disk at 1l:1:1:

```
Controller= Slot 1
```




```
Array= Next
Drive= 1I:1:1
LogicalDrive=Next
RAID= 0
SizeBlocks= 20971520
```

Advanced scripting scenarios

Clear the configuration on all internal controllers

Task

Clear the current configuration on all internal controllers in the system.

Hardware configuration

Unspecified

Script

```
# ssascripting -i input.txt -internal
Controller=ALL
ClearConfigurationWithDataLoss=Yes
```

Summary

By combining the advanced commands to specify `-internal` controllers first, and then `ALL` controllers, the user can clear configurations without affecting external controllers attached externally and without having to identify the internal controllers by slot or clear them one at a time.

Add a simple volume on systems with mixed controller and physical disk configurations

Task

Add a simple, one-disk RAID 0 volume on each system, using one script. However, do not add a new volume to systems booting from an external controller.

Hardware configuration

The administrator has 500 systems, some with embedded controllers, some with controllers in expansion slots, and some booting from an external controller, such as an HPE StorageWorks MSA1500cs Modular Smart Array.

Each system has a different physical disk configuration, some disks in an internal drive cage and some in an external storage box. Each system can have a different number of arrays and volumes.

Script

```
# ssascripting -i input.txt -internal
Controller=First
Array=Next
Drive=1
LogicalDrive=Next
RAID=0
```

Summary



By combining the advanced commands to specify `-internal` controllers only, and then `First` controllers, the user does not add volumes to external controllers or to additional internal controllers on the same system. The `Next` command eliminates the need to identify specific arrays and volumes.

Create multiple arrays on unconfigured systems with mixed physical disk configurations

Task

Add two arrays, each with one volume, to each system. Configure the first array as an OS volume with RAID 1 and a single spare. Configure the second array to use all remaining disks in a RAID 5 volume with no spares.

Hardware configuration

The administrator has 500 unconfigured systems with embedded controllers. Each system has a different physical disk configuration, some disks in an internal drive cage and some in an external storage box.

Script

```
Controller=Slot0
Array=A
Drive=2
OnlineSpare=1
LogicalDrive=1
RAID=1
Array=B
Drive=*
LogicalDrive=2
RAID=5
```

Summary

The `Slot0` specification targets embedded controllers only. Each array is configured using available drives. Array A uses two available drives, and the script eliminates the need to specify the drives by address. Array B uses all remaining available drives, again without specific drive addresses, but also without the administrator having to specify the number of drives, which can vary for each system.

Create the maximum size bootable volume

Task

When more than 2TiB of disk space is available, use all possible physical disks to create an array, create the maximum sized boot volume, and then consume the remaining space in another volume.

Hardware configuration

The administrator has an unconfigured system with a controller.

Script

```
Controller= Slot 1
Array= A
Drive= *
LogicalDrive=1 ; consumes the first 2TiB
RAID= 0
Size= MAXMBR
```



```
LogicalDrive=2 ; consumes remainder of free space
```

```
RAID=0
```

Summary

Without specifying drives by name, the administrator instructs the controller to use available disk space, up to 2TiB, to create the first volume in an array. Any remaining disk space is consumed in a second volume automatically.

Create a configuration with a specific block size

Task

Identify the size of a logical volume in blocks, and then create a configuration for another logical volume that matches exactly, using block size instead of capacity.

Currently, all Smart Array volumes consist of 512-byte blocks. When you create a custom configuration that uses the `Size` parameter, the script specifies the input size in MiB, which rounds the requested value into some number of blocks equal to or slightly greater than the requested size; the exact value depends on the array configuration.

However, if you use the `SizeBlocks` parameter to create the logical volume, then the script specifies the input size in the exact number of blocks, and no rounding occurs. In some cases, the RAID configuration might adjust the number of blocks in the new logical drive.

Hardware configuration

The administrator has a system with a controller and a configured logical volume.

Script output

To identify the size of an existing logical volume, in blocks, observe the output when you capture the existing configuration. Scripting outputs the current size of volumes in block units in a comment.

```
; Logical Drive Specifications
```

```
LogicalDrive= 2
```

```
RAID= 5
```

```
Size= 279964
```

```
; SizeBlocks= 573367088
```

```
Sectors= 32
```

```
StripSize= 512
```

```
Caching= Enabled
```

You can then use the size, specified in blocks, to configure another logical volume that is an exact match in size. When using an input script, you can specify either `Size` or `SizeBlocks`, but not both.

Summary

To create a logical drive that uses the exact same configuration of another logical drive, use the `SizeBlocks` parameter to specify the size, thus avoiding the rounding that can occur with the `Size` parameter.

Configuring HPE SmartCache

Task

This script illustrates how to configure HPE SmartCache.

Hardware configuration

The administrator has a system with a controller, physical drives, and logical drives configured.

Script output

```
Action= Configure
```



```
Method= Custom
Controller= Slot 1
Array= A
Drive= 2E:1:1, 2E:1:2, 2E:1:3
LogicalDrive= 1
RAID= 5
Size= 139947
CachingArray= B
Drive= 6I:0:5
CachingLogicalDrive= 2
CachedLogicalDrive= 1
RAID= 0
Size= 13993
```

Summary

Caching Logical Drives and Cached data drives must be on the same Controller.

Array A is created with 3 SAS physical drives, containing a single RAID 5 data logical drive (Logical Drive 1). Caching Array B is created using a single SSD drive, containing a single RAID 0 caching logical drive (Logical Drive 2) and is associated with Logical Drive 1.

Configuring HPE Secure Encryption

Task

This script illustrates how to configure HPE Secure Encryption.

Hardware configuration

The administrator has a system with a new and unconfigured controller, or the user has executed "EncryptionClearConfig" command to delete the encryption configuration.

Script output

```
Action= Configure
Method= Custom
Controller= SLOT 1
AcceptEULA=yes
EncryptionCryptoPasswordSet=Password.12345
EncryptionMasterKey= MyKeyName
EncryptionKeyManager= Local
AllowPlainText= True
Encryption= Enable
EncryptionUserPasswordSet= UPassword.9995 ; optional
ControllerPassword= MyBootTimePasswd ; optional
Firmwarelock= On ; optional
EncryptionRecoveryQuestion=This is my question ; optional
```



```
EncryptionRecoveryAnswer=This is my answer ; optional
```

Summary

Encryption is now configured with a password, master key name, recovery question and answer, and settings for either local or remote key mode.

Modifying encryption configurations

Task

This script illustrates how to modify encryption settings in HPE Secure Encryption.

Configuration

The administrator has a system with a configured controller.

Script output 1

```
Action= Reconfigure
Method= Custom
Controller= SLOT 1 ;first, log in
EncryptionUser= crypto
EncryptionPassword= Password.12345 ; now change the parameters
FirmwareLock = Off
AllowPlainText= false
```

Summary: The actual parameters to be changed will vary based on the controller capabilities and the version of controller firmware.

Script output 2

```
Action= Configure
Method= Custom
Controller= SLOT 1
EncryptionUser= crypto
EncryptionPassword= Password.12345
Array= A
Drives= 1
LogicalDrive= 1
RAID= 0
PlainText= Enable
```

Summary: When creating volumes, there is an additional volume option to set to determine whether the volume will be encrypted or plaintext. Prior login as the crypto user is required to use this feature. By default, volumes are encrypted.

Script output 3

```
Action= Reconfigure
Method= Custom
Controller= SLOT 1
EncryptionUser= crypto
EncryptionPassword= Password.12345
```



```
Array= A
LogicalDrive= 1
Encode= DestroyData
```

Summary: This example encodes (converts from plaintext to cyphertext) an existing plaintext volume, without preserving data.

XML support

HPE SSA scripting supports an XML file format for input and output.

XML output

To create an XML output document, use an XML suffix with the output file name:

```
C:\ssascripting -c out.xml
```

The following text is an example XML output file.

```
<?xml version="1.0"?> <Config.document> <!-- Date captured: Tue Jun 09
10:03:08 2009 --> <!-- Version: 8:30:4 --> <Action>Configure</Action>
<Method>Custom</Method> <Controller ID="Slot 1">

<!-- Controller HP Smart Array P410, Firmware Version 1.99 -->

<ReadCache>25</ReadCache>
<WriteCache>75</WriteCache>
<RebuildPriority>Medium</RebuildPriority>
<ExpandPriority>Medium</ExpandPriority>
<SurfaceScanDelay>3</SurfaceScanDelay>
<DriveWriteCache>Disabled</DriveWriteCache>
<LicenseKey>35DRP-7RH6S-R89GR-4MX6N-8K48X</LicenseKey>
<!-- LicenseKeyType "Flex License" -->
<MNPDelay>60</MNPDelay>
<IRPEnable>Disabled</IRPEnable>
<DPOEnable>Disabled</DPOEnable>
<ElevatorSortEnable>Enabled</ElevatorSortEnable>
<QueueDepth>Automatic</QueueDepth>
<!-- Unassigned Drives 1I:4:5 (60.0 GB), 2I:2:1 (72 GB) -->
<Array ID="A">
<!-- Array Drive Type is Solid State SATA -->
<!-- Free space 0 GBytes -->
<!-- 1I:4:8 (120.0 GB), 1I:4:7 (120.0 GB) -->
<Drive>1I:4:8, 1I:4:7</Drive>
<OnlineSpare>No</OnlineSpare>
<LogicalDrive ID="1">
<Raid>1</Raid>
<Size>114439</Size>
```



```

<Sectors>32</Sectors>
<StripeSize>128</StripeSize>
<ArrayAccelerator>Enabled</ArrayAccelerator>
</LogicalDrive>
</Array>
<Array ID="B">
<!-- Array Drive Type is SAS -->
<!-- Free space 0 GBytes -->
<!-- 1I:4:6 (72 GB),2I:2:4 (72 GB),2I:2:3 (72 GB),2I:2:2 (72 GB)-->
<Drive>1I:4:6, 2I:2:4, 2I:2:3, 2I:2:2</Drive>
<OnlineSpare>No</OnlineSpare>
<LogicalDrive ID="2">
<Raid>5</Raid>
<Size>52478</Size>
<Sectors>32</Sectors>
<StripeSize>64</StripeSize>
<ArrayAccelerator>Enabled</ArrayAccelerator>
</LogicalDrive>
<LogicalDrive ID="3">
<Raid>5</Raid>
<Size>52478</Size>
<Sectors>32</Sectors>
<StripeSize>64</StripeSize>
<ArrayAccelerator>Enabled</ArrayAccelerator>
</LogicalDrive>
<LogicalDrive ID="4">
<Raid>5</Raid>
<Size>52478</Size>
<Sectors>32</Sectors>
<StripeSize>64</StripeSize>
<ArrayAccelerator>Enabled</ArrayAccelerator>
</LogicalDrive>
<LogicalDrive ID="5">
<Raid>5</Raid>
<Size>52478</Size>
<Sectors>32</Sectors>
<StripeSize>64</StripeSize>
<ArrayAccelerator>Enabled</ArrayAccelerator>

```



```

</LogicalDrive>
</Array>
</Controller> </Config.document>

```

XML input

XML input follows the same format as the document for **XML output**. Use the same parameter order as the standard input format.

The following example shows a simple input script in both standard and XML formats.

Standard format	XML format
Action= Configure	<?xml version="1.0"?>
Method= Custom	<Config.document>
Controller= Slot 1	<Action>Configure</Action>
Array=A	<Method>Custom</Method>
Drive= 1I:4:8, 1I:4:7	<Controller ID="Slot 1">
LogicalDrive= 1	<Array ID="A">
RAID= 0	<Drive>1I:4:8, 1I:4:7</Drive>
Size= 100000	<LogicalDrive ID="1">
	<Raid>0</Raid>
	<Size>100000</Size>
	</LogicalDrive>
	</Array>
	</Controller>
	</Config.document>

XML input file DTD

The following DTD outlines the parameters for an HPE SSA scripting XML input file.

```

<!DOCTYPE Config.document [
<!ELEMENT Config.document ( Action, Method, Controller+ ) >
<!ELEMENT Action ( Configure | Reconfigure ) >
<!ELEMENT Method ( Auto | Custom ) >

<!ELEMENT Controller, Array ) ChassisName? | ClearConfigurationWithDataLoss |
DPOEnable? | DriveWriteCache? | ElevatorSortEnable? | ExpandPriority? |
IRPEnable? | Initiator? | LicenseKey? | MNPDelay? | PreferredPathMode? |
QueueDepth? | ReadCache? | RebuildPriority? | SurfaceScanDelay? |
WriteCache? ) >
<!ATTLIST Controller ID PCDATA #REQUIRED >

```




```

<!ELEMENT ClearConfigurationWithDataLoss ( YES | NO ) NO >
<!ELEMENT DPOEnable ( YES | NO ) >
<!ELEMENT DriveWriteCache ( ENABLE | DISABLE ) >
<!ELEMENT ElevatorSortEnable ( YES | NO ) >
<!ELEMENT ExpandPriority ( HIGH | MEDIUM | LOW ) >
<!ELEMENT IRPEnable ( YES | NO ) >
<!ELEMENT LicenseKey ( #PCDATA ) >
<!ELEMENT MNPDelay ( #PCDATA ) >
<!ELEMENT PreferredPathMode ( AUTO | MANUAL ) >
<!ELEMENT ReadCache ( 0 | 10 | 20 | 30 | 40 | 50 | 60 | 70 | 80 | 90 | 100 )
>
<!ELEMENT RebuildPriority ( HIGH | MEDIUM | LOW ) >
<!ELEMENT SurfaceScanDelay ( #PCDATA ) >
<!ELEMENT QueueDepth ( #PCDATA ) >
<!ELEMENT WriteCache ( 0 | 10 | 20 | 30 | 40 | 50 | 60 | 70 | 80 | 90 | 100 )
>

<!ELEMENT Array, LogicalDrive+ ( Drive, OnlineSpare ) >
<!ATTLIST Array ID PCDATA #REQUIRED >
<!ELEMENT Drive ( ALL | #PCDATA ) >
<!ELEMENT OnlineSpare ( YES | NO | #PCDATA ) >

<!ELEMENT LogicalDrive ( Raid, Repeat | Size, Sectors, StripeSize,
ArrayAccelerator? ) >
<!ATTLIST LogicalDrive ID PCDATA #REQUIRED >
<!ELEMENT ArrayAccelerator ( ENABLE | DISABLE ) ENABLE >
<!ELEMENT Raid ( 0 | 1 | 1ADM | 10 | 10ADM | 4 | 5 | 6 | ADG | 50 | 60 ) >
<!ELEMENT Repeat ( MAX | #PCDATA ) >
<!ELEMENT Sectors ( 32 | 63 ) >
<!ELEMENT Size ( MAX | #PCDATA ) MAX >
<!ELEMENT StripeSize ( 8 | 16 | 32 | 64 | 128 | 256 | 512 ) >
]>

```

HPE SSA scripting warning messages

Warning code	Warning message	Comment or clarification
4000	License key is already installed	The input file specified a license key. However, a license key is already installed on the controller.
4001	Clear configuration command failed - configuration is already cleared	The <code>-reset</code> option was used on a controller that does not have a configuration.

Table Continued



Warning code	Warning message	Comment or clarification
4002	Chassis name is already set to this value	The chassis name is already set to the supplied value. The command is ignored.
4003	One or more controller commands were skipped because the controller is not configured	Certain commands require that the controller be configured before they can be sent to the controller.
4004	Using Repeat function	Some commands have been ignored because the Repeat parameter has been specified in the input file.
4005	The system must be rebooted for the firmware flash to complete	A controller has been flashed with new firmware. The new firmware requires a reboot to take effect.
4006	Unable to set the array accelerator for this volume. The cache board may be missing or have a bad status, or the controller may not support a cache	The controller may not support a cache (for example, some Software RAID controllers) or the cache may be missing or defective.
4007	The controller does not have an energy pack	The operation requested requires an energy pack.
4008	The controller energy pack status is Failed	The operation requires an energy pack and the existing energy pack has failed.
4009	This parameter cannot be modified	The input script attempted to modify a parameter which is not modifiable.
4010	Rapid Parity Initialization is not available	Rapid Parity Initialization is not available on the selected volume; there may already be an RPI in progress, or the firmware does not support RPI.
4011	The command is not available when MODE=CONFIG	—
4012	The command is not available when MODE=RECONFIG	—

HPE SSA scripting error messages

Error code	Error message	Comment or clarification
0	—	No error exists. The program completed successfully.
2053	Too many coinciding expansion, migration, or extension operations	The new configuration requires more transformations than are possible at one time. For example, you cannot expand a logical volume and transform its RAID level at the same time.
2056	Controller does not support license keys	The controller does not support license key entry or deletion.
2059	Invalid license key	The entered license key is not a valid license key.
2064	Controller does not support SSP	The controller does not support SSP functions.
2817	Invalid Action	The requested actions are invalid, for example, combining <code>-reset</code> with capture mode.
2818	Invalid Method	The method must be either Custom or Auto.

Table Continued



Error code	Error message	Comment or clarification
2819	Invalid Controller	An invalid controller value was specified.
2821	No controllers detected	No controllers were detected. This error applies to Input mode only.
2823	Invalid Rebuild Priority	The specified rebuild priority is not supported.
2824	Invalid Expand Priority	The specified expand priority is not supported. This error also occurs if the controller does not allow expansion, and therefore does not support expand priority.
2826	Array not specified	The Array command is missing from the script file. Some commands were found that require an Array to be specified.
2827	New Array ID does not match the next available Array ID	The array ID in the script file does not match the array ID of the next available array. For example, if the configuration has an Array A and the input file specifies Array C (without Array B), then the script generates this error.
2828	New Array ID already exists	The array ID specified in the script file (in Configure mode) already exists in the configuration. Configure mode can only create new arrays.
2829	Cannot create Array	The controller will not allow a new array to be created, either because the controller has no unassigned physical drives attached or because the maximum number of arrays or logical drives has been reached already.
2830	Cannot expand Array	The Array could not be expanded, either because the controller does not support expansion, or the current configuration does not allow expansion to occur on the array.
2831	Cannot change Array Spare	The spare state of the array could not be changed. This error can occur when you attempt to add or drop a spare and the current configuration does not allow a spare state change for the array.
2832	Invalid physical drive. Possible reasons for this include: a typing error, the drive already exists in an array, the drive type does not match, or it is a failed drive.	A specified physical drive is not a valid physical drive, or it cannot be placed in the array.
2833	Invalid Spare	A specified spare is not a valid spare drive, or it cannot be placed in the array as a spare.
2834	Invalid logical drive	The logical drive ID is not valid.
2836	New Logical Drive ID does not match the next available logical drive ID. If you are attempting to replicate a configuration that has non-consecutive logical drive numbers, then you must change the script file so that the logical drive numbers are consecutive.	The script file specifies a logical drive ID that is not the first unused ID in the sequence. For example, this message appears if the controller has only Logical Drive 1 and the script file specifies creation of Logical Drive 3 (omitting Logical Drive 2). A common cause of this error is that the input file specifies nonsequential logical drive numbers. In this case, change the logical drive numbers in the input file so that they are sequential.

Table Continued



Error code	Error message	Comment or clarification
2837	New Logical Drive ID already exists	This error occurs in Configure mode when the logical drive ID specified in the script file already exists in the configuration. In Configure mode, you can create new logical drives only.
2838	Cannot create Logical Drive	The array has no free space, or the maximum number of logical drives has been reached already.
2839	Cannot migrate Logical Drive RAID	The controller does not support RAID migration, or migration is not possible with the current controller configuration.
2841	Cannot extend Logical Drive	The controller does not support extension, or the current controller configuration cannot be extended. For example, extension is not possible if the array has no free space.
2842	Invalid RAID	The specified RAID level is invalid or is not possible with the current physical disk and array configuration.
2843	Invalid Size	The specified size is invalid or is not possible with the current configuration.
2844	Invalid Stripe Size	The specified stripe size is invalid, not supported by the current RAID level, or not possible with the current configuration.
2849	Invalid ClearConfigurationWithDataLoss parameter	The valid parameters are Yes and No (default).
2850	Controller does not support Chassis Name	The controller does not support setting a chassis name.
2851	Invalid Chassis Name	The entered chassis name is invalid. Use characters from the set a–z, A–Z, 0–9, !, @, #, *, (,), ,, -, ., +, :, ;, /, and [space]. The name cannot end with a space character or exceed the maximum number of characters allowed by the controller.
2852	Invalid SSP State	The requested SSP state is not a valid SSP state.
2853	Cannot change SSP settings	The SSP settings cannot be changed for the controller or logical drive.
2854	Invalid SSP Adapter ID	The adapter ID is not a valid adapter ID that was detected by the controller.
2857	Invalid Surface Scan Delay	—
2861	Controller does not support redundancy settings	The controller is not redundant or does not support redundancy settings.
2864	Invalid Preferred Path Mode	The specified value for the preferred path mode is not valid, or the controller is not available.
2865	Invalid Preferred Path	The specified preferred path is not a valid chassis slot for an available, active controller, or the controller is not available.
2866	Failure opening capture file	—
2867	Failure opening input file	—

Table Continued



Error code	Error message	Comment or clarification
2868	Failure opening error file	—
2877	There are no suitable spares available	HPE SSA Scripting found no drives that could be used as spares for the specified array.
2880	Invalid Physical Disk Type Specified	—
2882	Invalid MNP delay	The specified value for MNP delay is invalid.
2883	Invalid Cache Logical Drive	The specified caching logical drive ID is invalid.
2884	Invalid Caching Array	The specified caching array ID is invalid.
2885	Error creating Caching Logical Drive	—
2886	Error creating Caching Array	—
2887	Error modifying Caching Volume	—
2888	Unable to modify Caching Arrays	Caching arrays cannot be modified once created.
2889	The requested Array modification is invalid	Error specifying the list of drives when reconfiguring an array
3000	Invalid Option	The value of the option supplied to this parameter is invalid.
3001	Invalid Command	The command is invalid.
3002	Command Failed	The controller returned an error for a command.
3003	License Key Delete Failed	HPE SSA Scripting is unable to delete the license key.
3004	Invalid Sector Size	—
3005	Cannot delete Array	—
3006	Invalid Number of Parity Groups	—
3007	Chassis name is too long	—
3008	Chassis name is already in use	Another controller is already using the entered chassis name.
3009	Auto Configure failed	The Auto Configure mode was unable to complete auto configuration.
3010	Cannot extend logical drive, not enough free space for the requested size	—
3011	Cannot extend logical drive, requested size is too small	—
3012	Cannot specify both SIZE and SHRINKSIZE	The input file cannot specify both a SIZE and SHRINKSIZE parameter.
3013	Cannot shrink Array	The array shrink operation was not successful.
3014	Cannot move Array	The array move operation was not successful.
3015	Invalid operation - Advanced Pack support required	The requested operation requires a valid license key to be entered.

Table Continued



Error code	Error message	Comment or clarification
3016	Spare drives cannot be specified by a count in Reconfigure mode	When in Reconfigure mode, the requested spare drives must be specified by their addresses. A simple count cannot be used.
3017	Disk drives cannot be specified by a count in Reconfigure mode	When in Reconfigure mode, the requested data drives must be specified by their addresses. A simple count cannot be used.
3018	Invalid number of physical disks	—
3019	Cannot create Array - no physical disks specified	HPE SSA Scripting cannot create an array unless physical disks are specified in the input file DRIVE parameter.
3020	SSP must be enabled in order to perform this operation	For the specified operation, HPE SSA Scripting requires SSP to be supported and enabled.
3021	Invalid connection name	—
3022	The <code>connectionname</code> cannot be removed when the <code>hostmode</code> has a non-default value.	—
3023	Invalid Host Mode	—
3024	Invalid Adapter ID	—
3025	This controller does not have host mode modification capability	—
3026	You need to have administrator rights to continue	—
3027	Another instance of HPE SSA is already running (possibly a service). Please terminate the HPE SSA application before running HPE SSA scripting.	—
3028	Invalid Drive Cache setting. Valid options are ENABLE and DISABLE.	—
3029	Invalid or out of order Command	Verify the ordering of the commands in the input configuration file.
3030	Invalid or missing Array for Reconfigure	Expecting a valid array in Reconfigure mode
3031	Invalid or missing Filename for Firmware Update	The firmware filename supplied is not a valid filename.
3032	Firmware Update Failed	The firmware flash operation on the controller failed.
3033	This controller has been configured with a more recent version of software. To prevent data loss, configuration changes to this controller are not allowed. Please upgrade to the latest version to be able to continue to configure this controller.	Some changes to SOULAPI are not backward-compatible. This check prevents the use of an older version of software on a controller configured with a newer version, possibly overriding changes and destroying data.

Table Continued



Error code	Error message	Comment or clarification
3034	Operations on this Array are temporarily unavailable while the Array is transforming.	The user requested too many simultaneous changes. For example, the user added new disks to an array (expand array) and changed the size or RAID level of logical volumes on the array. The solution is for the user to wait until the array transformation is complete.
3035	Invalid value for NoBatteryWriteCache	A problem exists with the NOBATTERYWRITECACHE command. Check the supplied argument. Not all controllers support this operation.
3036	Cannot delete Logical Drive	A problem occurred during the attempt to delete the specified logical drive. The drive may be in use by the OS, it may not exist, or it may not be the last logical drive on the array.
3038	The SplitMirror command is not valid for the specified Array	The operation is not available. Older array controllers/firmware may not support this command.
3046	This command is only available when running offline	The operation requested cannot be performed online. Boot to offline HPE SSA Scripting and supply the -offline parameter.
3047	Caching Volume parameters are not valid for a data Volume	The parameters supplied are not applicable or valid.
3048	Invalid Spare Mode	—
3053	The Firmware Lock/Unlock failed	The controller firmware returned an error when attempting to lock or unlock the firmware.
3056	Unable to set the boot password	The controller firmware returned an error when attempting to set the boot password.
3057	Unable to set the boot password state	The controller firmware returned an error when attempting to set the boot password state.
3064	Instant Secure Erase failed	The controller firmware returned an error when attempting to perform an Instant Secure Erase operation on the selected logical volume.
3100	SSD Smart Path is not supported	—
3101	The SSD Smart Path command failed	—
3102	Unable to lock controller. Make sure other instances of SSA, SSACLI, or SSASCRIPING are not configuring this controller	—

Encryption error messages

Error code	Error message	Comment or clarification
3049	Encryption Login Failed	—
3050	Encryption Login command is not available	—
3051	The Encryption command is not available	—

Table Continued



Error code	Error message	Comment or clarification
3052	The Zeroize command failed	The controller firmware returned an error when attempting to set zeroize the encryption configuration.
3054	Unable to change the Key Manager Mode	The controller firmware returned an error when attempting to change the key manager mode.
3055	Unable to set the Key Encrypting Key	The controller firmware returned an error when attempting to set Key Encrypting Key name.
3058	The controller does not support encryption	An attempt was made to execute an encryption command on a controller which does not support encryption.
3059	Controller Encryption is disabled	An attempt was made to execute an encryption command when encryption was disabled.
3060	Login required for this Encryption command	A login is required before the encryption command can be attempted.
3061	Plaintext volumes are not allowed	The option to create future plaintext volumes was not selected when initially configuring the encryption configuration.
3062	Rekey of the encrypted volume failed	The controller firmware returned an error when attempting to rekey the selected logical volume.
3063	Plaintext to Ciphertext conversion of the volume failed	The controller firmware returned an error when attempting to convert the logical volume from plaintext to ciphertext.
3065	Error changing the Controller Encryption Mode	The controller firmware returned an error when attempting to change the encryption mode (on/off).
3066	Error - you must indicate acceptance of the End User License Agreement (EULA) using ACCEPTEULA=YES	—

Table Continued



Error code	Error message	Comment or clarification
3067	Unable to set the password	An attempt to set the encryption password failed.
3068	Unable to set the Recovery question/answer	An attempt to set the encryption recovery question/answer pair failed.



Troubleshooting

HPE Smart Storage Administrator Diagnostics Utility CLI

About the utility

The HPE Smart Storage Administrator Diagnostics Utility CLI collects all possible information about storage devices in the system, detects all problems, and provides a detailed configuration report in .zip format.

After downloading and installing the software, you can run the utility as a CLI in an online environment. The functionality in this utility is mirrored in the diagnostics features of the HPE Smart Storage Administrator, which can be run in an offline environment.

The utility generates two types of reports:

- Array diagnostic report

This report contains information about all devices, such as array controllers, storage enclosures, drive cages, as well as logical, physical, and tape drives. For supported solid state drives, this report also contains SmartSSD Wear Gauge information.

- SmartSSD Wear Gauge report

This report contains information about the current usage level and remaining expected lifetime of solid state drives attached to the system.

For more information, see "[**Reported information**](#)."

Reported information

The array diagnostic report provides detailed information about devices (array controllers, storage enclosures, drive cages, physical drives, logical drives, and tape drives).

For example, device information on a typical embedded controller might include:

- Software versions
- Errors
- Controller information:
 - Name
 - Attached devices
 - Description
 - PCI bus
 - PCI device
 - PCI function
- Drive information:
 - Interface
 - WWID
 - Drive model



- Serial number
- Firmware revision
- Total blocks

The SmartSSD Wear Gauge report contains information on the current usage level of and expected lifetime remaining for solid state drives attached to the system.

For discovered SSDs, the report summary page provides the following calculated totals:

- Total Solid State Drives with Wearout Status
- Total Smart Array Solid State Drives
- Total Non Smart Array Solid State Drives
- Total Solid State SAS Drives
- Total Solid State SATA Drives
- Total Solid State Drives

In addition to these totals, the summary page also displays the following tables:

- Solid State Drives with Wearout Status
- Solid State Drives with Less Than an Estimated 56 Days of Life Remaining
- Solid State Drives with Less Than 2% Usage Remaining
- Solid State Drives with Less Than 5% Usage Remaining
- Smart Array Controllers
- Non Smart Array Controllers

When the generated report is viewed in a browser, the report page displays the following fields in the SmartSSD Status table.



Field	Description
SSD Wear Status	<p>Indicates the SSD's wear status with one of the following messages:</p> <ul style="list-style-type: none"> • OK • Not Supported • The SmartSSD Wear Gauge log is full. Wear Gauge parameters are not available. • SSD has less than 5% usage remaining before wearout. • SSD has less than 2% usage remaining before wearout. • SSD has less than an estimated 56 days before it reaches the maximum usage limit for writes (wearout) and should be replaced as soon as possible. • SSD has less than 5% of usage remaining before wearout. It has less than an estimated 56 days before it reaches the maximum usage limit and should be replaced as soon as possible. • SSD has less than 2% of usage remaining before wearout. It has less than an estimated 56 days before it reaches the maximum usage limit and should be replaced as soon as possible. • SSD has reached the maximum rated usage limit for writes (wearout) and should be replaced immediately.
Power Cycles	Indicates the number of times the SSD has powered on from the powered off state
Power On Hours	Indicates the number of hours the SSD has been powered on
Estimated Life Remaining Based On Workload To Date	Indicates an estimate of the number of days the SSD has before SSD Utilization reaches 100%. This field is displayed only if the value shown for Usage Remaining is less than 100%.
Usage Remaining	Indicates the percentage of the SSD that has not worn out. Usage remaining is equal to the difference of 100 and the SSD Utilization percentage.
SSD Utilization	Indicates the percentage of the SSD that has worn out. This field is displayed only if the value shown for Usage Remaining is less than 100%.

Installing the utility

Procedure

1. Browse to the [HPE Smart Storage Administrator website](#).
2. Click **Download software**.
3. Select an OS.
4. Identify the preferred software and version, and then click **Download**.
5. Save, and then run, the executable file.

By default, the software installs at C:\Program Files\Smart Storage Administrator\.

Launching the utility in CLI mode



Procedure

1. Click **Start>All Programs>Windows System>Smart Storage Administrator Diagnostics Utility>Read Me**
2. Open a command prompt as administrator.
3. Change directory (`cd`) to the location where `ssaduccli.exe` is installed.
This directory is commonly `C:\Program Files\Smart Storage Administrator\ssaduccli\bin`.
4. Do one of the following:
 - Generate a diagnostic report with the following command:`ssaduccli -f adu-report.zip`
 - Generate a SmartSSD Wear Gauge report with the following command:`ssaduccli -ssd -f ssd-report.zip`

For more options, use the following command:

```
ssaduccli -help
```

Diagnostic report procedures

Viewing the diagnostic report

Procedure

1. **Launch the utility.**
2. Browse to the .zip file you created using the utility.
3. Open the HTML file to view the report.

Identifying and viewing diagnostic report files

The diagnostic report output archive contains the following files:

- `ADUReport.txt`—Diagnostic report in text format
- `ADUReport.xml`—Diagnostic report in XML format
- `ADUReportViewer.htm`—HTML viewer for XML diagnostic report
- `SlotX.txt` (`SlotX.old`)—Controller serial output log

The serial output log file(s) are only available if the Smart Array SAS/SATA Event Notification Service is installed and running.

Procedure

1. Extract `ADUReportViewer.htm` to a directory.
2. Open `ADUReportViewer.htm` in the browser.



SmartSSD Wear Gauge report procedures

Viewing the SmartSSD Wear Gauge report

Procedure

1. Launch the utility.
2. Browse to the .zip file you created using the utility.
3. Open the HTML file to view the report.

Identifying and viewing SmartSSD Wear Gauge report files

The SmartSSD Wear Gauge report output archive contains the following files:

- `SmartSSDWearGaugeReport.txt`—SmartSSD wear gauge report in text format
- `SmartSSDWearGaugeReport.json`—SmartSSD wear gauge report in JSON format
- `SmartSSDWearGaugeReport.htm`—HTML viewer for the JSON wear gauge report

Procedure

1. Extract the following files to a single directory:
 - `SmartSSDWearGaugeReport.json`
 - `SmartSSDWearGaugeReport.htm`
2. Open `SmartSSDWearGaugeReport.htm` in the browser.

512e Physical drive support

HPE Smart Storage Administrator is able to detect and correct performance issues caused by nonoptimal logical drive alignment for 512e physical drives.

The following scenarios indicate that drive support is needed:

- Multiple logical drives exist in a single array.
- An array consists of one or more 512e physical drives.
- At least one of the logical drives in the array is not aligned on a native block boundary. For current 512e drives, the native block boundary is 4K.

HPE SSA displays a warning indicating that the logical drive is not optimally aligned and that performance of the logical drive is not optimal. Additionally, if one or more of the following scenarios are met, the array presents a "Realign Logical Drive" button:

- There is enough free space in the array to move the logical drive to be aligned to the native 4K boundary.
- The controller can perform the transformation (requires a cache module with a fully charged energy pack).
- The controller does NOT have SmartCache enabled.




Drive arrays and fault-tolerance methods

Drive arrays


The capacity and performance of a single physical (hard) drive is adequate for home users. However, business users demand higher storage capacities, higher data transfer rates, and greater protection against data loss when drives fail.

Connecting extra physical drives (P_n in the figure) to a system increases the total storage capacity but has no effect on the efficiency of read/write (R/W) operations. Data can still be transferred to only one physical drive at a time.

With an array controller installed in the system, the capacity of several physical drives can be combined into one or more virtual units called **logical drives** (also called **logical volumes** and denoted by L_n in the figures in this section). Then, the read/write heads of all the constituent physical drives are active simultaneously, reducing the total time required for data transfer.



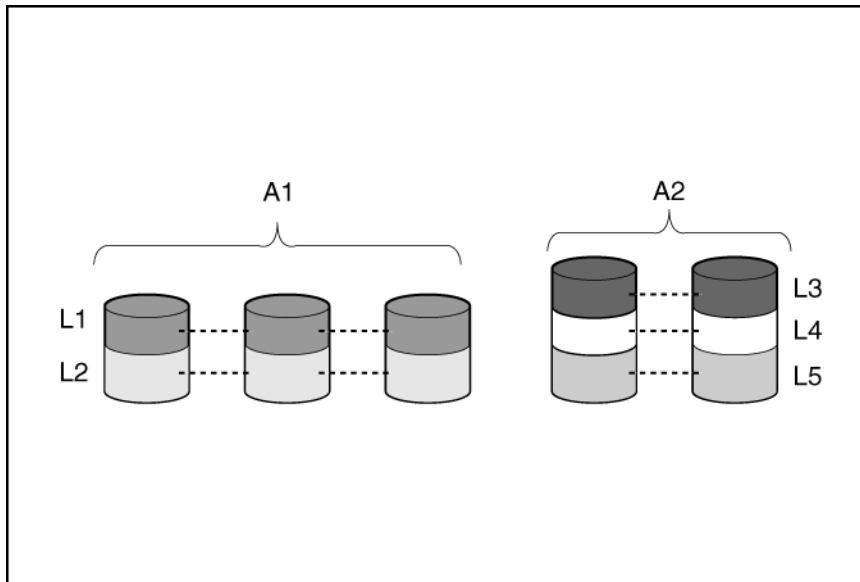
Because the read/write heads are active simultaneously, the same amount of data is written to each drive during any given time interval. Each unit of data is called a **block** (denoted by B_n in the figure), and adjacent blocks form a set of data **stripes** (S_n) across all the physical drives that comprise the logical drive.



For data in the logical drive to be readable, the data block sequence must be the same in every stripe. This sequencing process is performed by the array controller, which sends the data blocks to the drive write heads in the correct order.

A natural consequence of the striping process is that each physical drive in a given logical drive will contain the same amount of data. If one physical drive has a larger capacity than other physical drives in the same logical drive, the extra capacity is wasted because it cannot be used by the logical drive.

The group of physical drives containing the logical drive is called a **drive array**, or just **array** (denoted by A_n in the figure). Because all the physical drives in an array are commonly configured into just one logical drive, the term array is often used as a synonym for logical drive. However, an array can contain several logical drives, each of a different size.



Each logical drive in an array is distributed across all the physical drives within the array. A logical drive can also extend across more than one port on the same controller, but it cannot extend across more than one controller.

Drive failure, although rare, is potentially catastrophic. For arrays that are configured as shown in the previous figure, failure of any physical drive in the array causes every logical drive in the array to suffer irretrievable data loss. To protect against data loss due to physical drive failure, logical drives are configured with fault tolerance.

For any configuration except RAID 0, further protection against data loss can be achieved by assigning a drive as an **online spare** (or **hot spare**). This drive contains no data and is connected to the same controller as the array. When any other physical drive in the array fails, the controller automatically rebuilds information that was originally on the failed drive to the online spare. The system is thus restored to full RAID-level data protection, although it now no longer has an online spare. (However, in the unlikely event that another drive in the array fails while data is being rewritten to the spare, the logical drive will still fail.)

When you configure an online spare, it is automatically assigned to all logical drives in the same array. Additionally, you do not need to assign a separate online spare to each array. Instead, you can configure one hard drive to be the online spare for several arrays if the arrays are all on the same controller.

More information

[Fault-tolerance methods](#)

Effects of a hard drive failure on logical drives

When a drive fails, all logical drives that are in the same array are affected. Each logical drive in an array might be using a different fault-tolerance method, so each logical drive can be affected differently.

- RAID 0 configurations do not tolerate drive failure. If any physical drive in the array fails, all RAID 0 logical drives in the same array also fail.
- RAID 1 tolerates the failure of one drive.

- RAID 10 configurations tolerate multiple drive failures as long as no failed drives are mirrored to another failed drive.
- RAID 5 configurations tolerate one drive failure.
- RAID 50 configurations tolerate one failed drive in each parity group.
- RAID 6 configurations tolerate two failed drives at a given time.
- RAID 60 configurations tolerate two failed drives in each parity group.
- RAID 1 (ADM) and RAID 10 (ADM) configurations tolerate multiple drive failures if no more than two drives, mirrored to one another, fail.

Fault-tolerance methods

Several fault-tolerance methods exist. The methods most often used with Smart Array controllers are hardware-based RAID methods.


Alternative fault-tolerance methods are also available. However, hardware-based RAID methods provide a much more robust and controlled fault-tolerance environment, so these alternative methods are seldom used.

More information

[Alternative fault-tolerance methods](#)

RAID 0

A RAID 0 configuration provides data striping, but there is no protection against data loss when a drive fails. However, it is useful for rapid storage of large amounts of noncritical data (for printing or image editing, for example) or when cost is the most important consideration. The minimum number of drives required is one.



This method has the following benefits:


- Useful when performance and low cost are more important than data protection.
- Has the highest write performance of all RAID methods.
- Has the lowest cost per unit of stored data of all RAID methods.
- All drive capacity is used to store data (none allocated for fault tolerance).




RAID 1 and RAID 1+0 (RAID 10)

In RAID 1 and RAID 1+0 (RAID 10) configurations, data is duplicated to a second drive. The usable capacity is $C \times (n / 2)$ where C is the drive capacity with n drives in the array. A minimum of two drives is required.

When the array contains only two physical drives, the fault-tolerance method is known as RAID 1.



When the array has more than two physical drives, drives are mirrored in pairs, and the fault-tolerance method is known as RAID 1+0 or RAID 10. If a physical drive fails, the remaining drive in the mirrored pair can still provide all the necessary data. Several drives in the array can fail without incurring data loss, as long as no two failed drives belong to the same mirrored pair. The total drive count must increment by 2 drives. A minimum of four drives is required.



This method has the following benefits:

- It is useful when high performance and data protection are more important than usable capacity.
- This method has the highest write performance of any fault-tolerant configuration.




- No data is lost when a drive fails, as long as no failed drive is mirrored to another failed drive.
- Up to half of the physical drives in the array can fail.


RAID 1 (ADM) and RAID 10 (ADM)

In RAID 1 (Advanced Data Mirroring (ADM)) and RAID 10 (ADM) configurations, data is duplicated to two additional drives. The usable capacity is $C \times (n / 3)$ where C is the drive capacity with n drives in the array. A minimum of 3 drives is required.

When the array contains only three physical drives, the fault-tolerance method is known as RAID 1 (ADM).



When the array has more than six physical drives, drives are mirrored in trios, and the fault-tolerance method is known as RAID 10 (ADM). If a physical drive fails, the remaining two drives in the mirrored trio can still provide all the necessary data. Several drives in the array can fail without incurring data loss, as long as no three failed drives belong to the same mirrored trio. The total drive count must increment by 3 drives.




This method has the following benefits:

- It is useful when high performance and data protection are more important than usable capacity.
- This method has the highest read performance of any configuration due to load balancing.
- This method has the highest data protection of any configuration.
- No data is lost when two drives fail, as long as no two failed drives are mirrored to another failed drive.
- Up to two-thirds of the physical drives in the array can fail.

RAID 5

RAID 5 protects data using parity (denoted by $P_{x,y}$ in the figure). Parity data is calculated by summing (XOR) the data from each drive within the stripe. The strips of parity data are distributed evenly over every physical drive within the logical drive. When a physical drive fails, data that was on the failed drive can be recovered from the remaining parity data and user data on the other drives in the array. The usable capacity is $C \times (n - 1)$ where C is the drive capacity with n drives in the array. A minimum of three drives is required.




This method has the following benefits:

- It is useful when usable capacity, write performance, and data protection are equally important.
- It has the highest usable capacity of any fault-tolerant configuration.
- Data is not lost if one physical drive fails.

RAID 6

RAID 6 protects data using double parity. With RAID 6, two different sets of parity data are used (denoted by $P_{x,y}$ and $Q_{x,y}$ in the figure), allowing data to still be preserved if two drives fail. Each set of parity data uses a capacity equivalent to that of one of the constituent drives. The usable capacity is $C \times (n - 2)$ where C is the drive capacity with n drives in the array.

A minimum of 4 drives is required.




This method is most useful when data loss is unacceptable but cost is also an important factor. The probability that data loss will occur when an array is configured with RAID 6 (Advanced Data Guarding (ADG)) is less than it would be if it were configured with RAID 5.

This method has the following benefits:

- It is useful when data protection and usable capacity are more important than write performance.
- It allows any two drives to fail without loss of data.

RAID 50

RAID 50 is a nested RAID method in which the constituent drives are organized into several identical RAID 5 logical drive sets (parity groups). The smallest possible RAID 50 configuration has six drives organized into two parity groups of three drives each.



For any given number of drives, data loss is least likely to occur when the drives are arranged into the configuration that has the largest possible number of parity groups. For example, four parity groups of three drives are more secure than three parity groups of four drives. However, less data can be stored on the array with the larger number of parity groups.

All data is lost if a second drive fails in the same parity group before data from the first failed drive has finished rebuilding. A greater percentage of array capacity is used to store redundant or parity data than with non-nested RAID methods (RAID 5, for example). A minimum of six drives is required.

This method has the following benefits:

- Higher performance than for RAID 5, especially during writes.
- Better fault tolerance than either RAID 0 or RAID 5.
- Up to n physical drives can fail (where n is the number of parity groups) without loss of data, as long as the failed drives are in different parity groups.

RAID 60

RAID 60 is a nested RAID method in which the constituent drives are organized into several identical RAID 6 logical drive sets (parity groups). The smallest possible RAID 60 configuration has eight drives organized into two parity groups of four drives each.

For any given number of hard drives, data loss is least likely to occur when the drives are arranged into the configuration that has the largest possible number of parity groups. For example, five parity groups of four drives are more secure than four parity groups of five drives. However, less data can be stored on the array with the larger number of parity groups.

The number of physical drives must be exactly divisible by the number of parity groups. Therefore, the number of parity groups that you can specify is restricted by the number of physical drives. The maximum number of parity groups possible for a particular number of physical drives is the total number of drives divided by the minimum number of drives necessary for that RAID level (three for RAID 50, 4 for RAID 60).

A minimum of 8 drives is required.

All data is lost if a third drive in a parity group fails before one of the other failed drives in the parity group has finished rebuilding. A greater percentage of array capacity is used to store redundant or parity data than with non-nested RAID methods.

This method has the following benefits:

- Higher performance than for RAID 6, especially during writes.
- Better fault tolerance than RAID 0, 5, 50, or 6.
- Up to $2n$ physical drives can fail (where n is the number of parity groups) without loss of data, as long as no more than two failed drives are in the same parity group.

Comparing the hardware-based RAID methods

Not all controllers support all RAID levels.

Item	RAID 0	RAID 1+0	RAID 5	RAID 6 (ADG)	RAID 1(0) (ADM)
Alternative name	Striping (no fault tolerance)	Mirroring	Distributed Data Guarding	Advanced Data Guarding	Advanced Data Mirroring
Formula for number of drives usable for data (n = total number of drives in array)	n	$n/2$	$n-1$	$n-2$	$n/3$

Table Continued



Item	RAID 0	RAID 1+0	RAID 5	RAID 6 (ADG)	RAID 1(0) (ADM)
Percentage of drive space usable ¹	100%	50%	67% to 93%	50% to 96%	33%
Minimum number of physical drives	1	2	3	4	3
Tolerates failure of one physical drive	No	Yes	Yes	Yes	Yes
Tolerates simultaneous failure of more than one physical drive	No	Only if no two failed drives are in the same mirrored pair	No	Yes	Only if no three drives are in the same mirror group ²
Read performance	High	High	High	High	High
Write performance	High	Medium	Low	Low	Medium
Relative cost	Low	High	Medium	Medium	Very high

¹ Values for the percentage of drive space usable are calculated with these assumptions: (1) all physical drives in the array have the same capacity; (2) online spares are not used; (3) no more than 14 physical drives are used per array for RAID 5; and (4) no more than 56 drives are used with RAID 6 (ADG).

² Mirror groups include the physical drives in each mirror.

Selecting a RAID method

Not all controllers support all RAID levels. To determine the RAID capabilities of your controller, see the model-specific information for your controller on the [Hewlett Packard Enterprise website](#).

Most important criterion	Also important	Suggested RAID level
Fault tolerance	Cost effectiveness	RAID 6
	I/O performance	RAID 10 (ADM), RAID 1+0, RAID 50, RAID 60
Cost effectiveness	Fault tolerance	RAID 6
	I/O performance	RAID 5 (RAID 0 if fault tolerance is not required)
I/O performance	Cost effectiveness	RAID 5 (RAID 0 if fault tolerance is not required)
	Fault tolerance	RAID 10 (ADM), RAID 1+0, RAID 50, RAID 60

Alternative fault-tolerance methods

Your operating system may also support software-based RAID or controller duplexing.

- **Software-based RAID** resembles hardware-based RAID, except that the operating system works with logical drives as if they were physical drives. To protect against data loss caused by physical drive failure, each logical drive must be in a different array from the others.
- **Controller duplexing** uses two identical controllers with independent, identical sets of drives containing identical data. In the unlikely event of a controller failure, the remaining controller and drives will service all requests.



Neither of these alternative fault-tolerance methods supports online spares or automatic data recovery, nor do they support auto-reliability monitoring or interim data recovery.

If you decide to use one of these alternative methods, configure your arrays with RAID 0 for maximum storage capacity and refer to your operating system documentation for further implementation details.

Diagnosing array problems

Diagnostic tools

To troubleshoot array problems and generate feedback about arrays, use the following diagnostic tools:

- **Event Notification Service**

This utility reports array events to the Microsoft Windows system event log and IML. You can obtain the utility from the SmartStart CD or the **Hewlett Packard Enterprise website**. When prompted for product information, enter the server model name.

- **HPE Insight Diagnostics**

Insight Diagnostics is a tool that displays information about the system hardware configuration and performs tests on the system and its components, including drives if they are connected to Smart Array controllers. This utility is available on the **Hewlett Packard Enterprise website**.

- **POST messages**

Smart Array controllers produce diagnostic error messages (POST messages) at reboot. Many POST messages suggest corrective actions. For more information about POST messages, see the *ProLiant Servers Troubleshooting Guide*.

- **HPE Smart Storage Administrator**

For information about error messages, see the *ProLiant Servers Troubleshooting Guide*.

- **HPE Smart Storage Administrator Diagnostics Utility CLI**

This standalone diagnostic utility provides configuration and error information about array controllers, storage enclosures, drive cages, logical drives, physical drives, and tape drives. For any supported SSDs, the utility provides current usage level and remaining expected lifetime. For more information, see "**Smart Storage Administrator Diagnostics Utility CLI**."

Troubleshooting resources

Troubleshooting resources are available for HPE Gen10 server products in the following documents:

- *Troubleshooting Guide for HPE ProLiant Gen10 servers* provides procedures for resolving common problems and comprehensive courses of action for fault isolation and identification, issue resolution, and software maintenance.
- *Error Message Guide for HPE ProLiant Gen10 servers and HPE Synergy* provides a list of error messages and information to assist with interpreting and resolving error messages.
- *Integrated Management Log Messages and Troubleshooting Guide for HPE ProLiant Gen10 and HPE Synergy* provides IML messages and associated troubleshooting information to resolve critical and cautionary IML events.

To access the troubleshooting resources, see the Hewlett Packard Enterprise Information Library (**<http://www.hpe.com/info/gen10-troubleshooting>**).



Optional components

HPE Smart Array SR Secure Encryption

HPE Smart Array SR Secure Encryption is a controller-based, enterprise-class data encryption solution that protects data at rest on any drive attached to the HPE Smart Array controller. The solution is available for both local and remote deployments, and works in conjunction with HPE Smart Array SR SmartCache.

HPE Smart Array SR Secure Encryption functions to secure sensitive data, for example, all business data or data deemed mission critical and requiring extra levels of protection. Protection is accomplished through the use of data encryption keys and a simple deployment methodology. Many companies under government regulations require that sensitive data must be secure and uncompromised. HPE Secure Encryption provides a **FIPS 140-2** Level 1 validated enterprise class controller-based encryption solution for data-at-rest on all SAS/SATA drives. Not only is all data at rest protected, but any data that resides on the FBWC module, which is used as a write-caching mechanism, is also encrypted.

Using the Encryption Manager, you can access Secure Encryption features to perform tasks such as key management through the HPE Enterprise Secure Key Manager and managing passwords. For more information about optional encryption solutions for HPE SSA, see *HPE Smart Array SR Secure Encryption Installation and User Guide*.

To purchase a license key, see the Secure Encryption product page on the [**Hewlett Packard Enterprise website**](#).



HPE Smart Array Advanced Pack

More information

[Support for configuration tasks](#)

About SAAP

SAAP is a collection of additional and advanced controller features embedded in the firmware of select Smart Array controllers. For more information on the current features available through SAAP, see the [Hewlett Packard Enterprise website](#).

For more information about advanced features, see the SAAP QuickSpecs on the [Hewlett Packard Enterprise website](#).

To install the license key and activate SAAP, use one of the following methods:

- Installing a license key with HPE SSA CLI
- **[Installing a license key with HPE SSA Scripting](#)**

Some SAAP features are required for advanced configuration tasks in HPE SSA. For a list of these tasks, see "Support for configuration tasks."



Support and other resources

Accessing Hewlett Packard Enterprise Support

- For live assistance, go to the Contact Hewlett Packard Enterprise Worldwide website:
<https://www.hpe.com/info/assistance>
- To access documentation and support services, go to the Hewlett Packard Enterprise Support Center website:
<https://www.hpe.com/support/hpesc>

Information to collect

- Technical support registration number (if applicable)
- Product name, model or version, and serial number
- Operating system name and version
- Firmware version
- Error messages
- Product-specific reports and logs
- Add-on products or components
- Third-party products or components

Accessing updates

- Some software products provide a mechanism for accessing software updates through the product interface. Review your product documentation to identify the recommended software update method.
- To download product updates:

Hewlett Packard Enterprise Support Center

<https://www.hpe.com/support/hpesc>

Hewlett Packard Enterprise Support Center: Software downloads

<https://www.hpe.com/support/downloads>

My HPE Software Center

<https://www.hpe.com/software/hpesoftwarecenter>

- To subscribe to eNewsletters and alerts:
<https://www.hpe.com/support/e-updates>
- To view and update your entitlements, and to link your contracts and warranties with your profile, go to the Hewlett Packard Enterprise Support Center **More Information on Access to Support Materials** page:
<https://www.hpe.com/support/AccessToSupportMaterials>





IMPORTANT: Access to some updates might require product entitlement when accessed through the Hewlett Packard Enterprise Support Center. You must have an HPE Passport set up with relevant entitlements.

Customer self repair

Hewlett Packard Enterprise customer self repair (CSR) programs allow you to repair your product. If a CSR part needs to be replaced, it will be shipped directly to you so that you can install it at your convenience. Some parts do not qualify for CSR. Your Hewlett Packard Enterprise authorized service provider will determine whether a repair can be accomplished by CSR.

For more information about CSR, contact your local service provider.

Remote support

Remote support is available with supported devices as part of your warranty or contractual support agreement. It provides intelligent event diagnosis, and automatic, secure submission of hardware event notifications to Hewlett Packard Enterprise, which will initiate a fast and accurate resolution based on your product's service level. Hewlett Packard Enterprise strongly recommends that you register your device for remote support.

If your product includes additional remote support details, use search to locate that information.

Remote support and Proactive Care information

HPE Get Connected

<https://www.hpe.com/services/getconnected>

HPE Proactive Care services

<https://www.hpe.com/services/proactivecare>

HPE Datacenter Care services

<https://www.hpe.com/services/datacentercare>

HPE Proactive Care service: Supported products list

<https://www.hpe.com/services/proactivecaresupportedproducts>

HPE Proactive Care advanced service: Supported products list

<https://www.hpe.com/services/proactivecareadvancedsupportedproducts>

Proactive Care customer information

Proactive Care central

<https://www.hpe.com/services/proactivecarecentral>

Proactive Care service activation

<https://www.hpe.com/services/proactivecarecentralgetstarted>

Warranty information

To view the warranty information for your product, see the links provided below:

HPE ProLiant and IA-32 Servers and Options

<https://www.hpe.com/support/ProLiantServers-Warranties>

HPE Enterprise and Cloudline Servers

<https://www.hpe.com/support/EnterpriseServers-Warranties>

HPE Storage Products

<https://www.hpe.com/support/Storage-Warranties>



Regulatory information

To view the regulatory information for your product, view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products*, available at the Hewlett Packard Enterprise Support Center:

<https://www.hpe.com/support/Safety-Compliance-EnterpriseProducts>

Additional regulatory information

Hewlett Packard Enterprise is committed to providing our customers with information about the chemical substances in our products as needed to comply with legal requirements such as REACH (Regulation EC No 1907/2006 of the European Parliament and the Council). A chemical information report for this product can be found at:

<https://www.hpe.com/info/reach>

For Hewlett Packard Enterprise product environmental and safety information and compliance data, including RoHS and REACH, see:

<https://www.hpe.com/info/ecodata>

For Hewlett Packard Enterprise environmental information, including company programs, product recycling, and energy efficiency, see:

<https://www.hpe.com/info/environment>

Documentation feedback

Hewlett Packard Enterprise is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback (docsfeedback@hpe.com). When submitting your feedback, include the document title, part number, edition, and publication date located on the front cover of the document. For online help content, include the product name, product version, help edition, and publication date located on the legal notices page.

Websites

General websites

Hewlett Packard Enterprise Information Library

www.hpe.com/info/EIL

Subscription Service/Support Alerts

www.hpe.com/support/e-updates

Insight Remote Support

www.hpe.com/info/insightremotesupport/docs

For additional general support websites, see **[Support and other resources](#)**.

Product websites

HPE Smart Array

<http://www.hpe.com/info/SmartArrayGen10-docs>



HPE Integrated Lights-Out

www.hpe.com/info/ilo

