



Hewlett Packard
Enterprise

HPE Smart Storage Administrator

User Guide

Abstract

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

Part Number: 742648-007a
March 2016
Edition: 7

© Copyright 2013, 2016 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	5
About HPE SSA	5
Benefits of using HPE SSA	5
Support for configuration tasks	6
Array configuration	8
Array configuration guidelines	8
Minimum requirements	8
Native support for 64-bit and 32-bit operating systems	9
Operations	10
Accessing HPE SSA in the offline environment	10
Launching HPE SSA with Intelligent Provisioning (Gen8 or later)	10
Launching HPE SSA during POST (Gen8 or later)	10
Launching HPE SSA from an ISO image (all generations)	10
Accessing HPE SSA in the online environment	14
Launching HPE SSA on a local server	14
Launching HPE SSA on a local server to configure a remote server	15
Launching HPE SSA on a remote server to configure a local server	16
Using the HPE SSA GUI	16
Icon and key stroke legend	17
Navigating the GUI	18
Configuration tasks	23
Diagnostic tasks	58
Using the HPE SSA CLI	60
Opening the CLI in Console mode	60
Opening the CLI in Command mode	60
CLI syntax	61
Typical procedures	65
Encryption commands	84
Using HPE SSA scripting	89
Capturing a configuration	90
Using an Input script	90
Creating an HPE SSA script file	90
Script file options	95
XML support	111
HPE SSA scripting warning messages	115
HPE SSA scripting error messages	115
Troubleshooting	121
HPE Smart Storage Administrator Diagnostics Utility CLI	121
About the utility	121
Reported information	121
Installing the utility	123
Launching the utility in CLI mode	123
Diagnostic report procedures	123
SmartSSD Wear Gauge report procedures	124
512e Physical drive support	124
Drive arrays and fault-tolerance methods	125
Drive arrays	125
Effects of a hard drive failure on logical drives	127
Fault-tolerance methods	128
Diagnosing array problems	136
Diagnostic tools	136
Troubleshooting resources	136

Optional components	138
HPE Secure Encryption.....	138
HPE Smart Array Advanced Pack	139
About SAAP	139
Acronyms and abbreviations.....	140
Documentation feedback	142
Index	143

Overview

About HPE SSA

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 (on page 6). 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 (on page 121).

Starting with HPE SSA and ProLiant Gen8 servers and server blades, HPE SSA is accessible both offline and online:

- Accessing HPE SSA in the offline environment (on page 10)
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 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 and boot volume.
- 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.

Benefits of using HPE SSA

HPE Smart Storage Administrator is an advanced utility that enables you to perform many complex configuration tasks. Previously, other HPE configuration utilities, including Array Configuration Utility and Option ROM Configuration for Arrays, were recommended for storage configuration.

Although ACU is still supported, HPE SSA replaces it starting with ProLiant Gen8 servers.

There are many GUI differences between ACU and HPE SSA. Some of the main differences include the following:

- The tabs in ACU are integrated into the new HPE SSA **Configure** menu.
- The 2-panel format in ACU is replaced with a 3-panel format in HPE SSA.
- HPE SSA includes a quick navigation menu, to access configuration or diagnostics options for a specific controller.
- With a Smart Array Gen8 or later controller, most options do not require an SAAP license.
- HPE SSA does not include a Wizards feature.

HPE SSA includes a stand-alone CLI diagnostics utility ("HPE Smart Storage Administrator Diagnostics Utility CLI" on page 121).

HPE SSA provides full-range support for standard and advanced configuration tasks ("Support for configuration tasks" on page 6). Some of these advanced tasks are not available in all HPE SSA interface formats (GUI, CLI, and Scripting).

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

- GUI, CLI, and Scripting interfaces are available.

- English, French, German, Italian, Japanese, Simplified Chinese, and Spanish languages are supported.
- Applications can be executed using the following tools:
 - Any bootable media, such as a software CD
 - Any supported Host operating system with a web browser (Microsoft Windows 2003 is not supported. For a list of supported operating systems, see the Hewlett Packard Enterprise website (<http://www.hpe.com/info/ossupport>).)
 - Intelligent Provisioning, which is embedded in ProLiant Gen8 and later servers
- 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

The following table uses these symbols:

+ —The controller supports this task through HPE SSA.

— —The controller supports this task through HPE SSA only with a registered SAAP license key.

—The controller supports this task through HPE SSA only with a registered SmartCache license key.

To identify controller-specific feature support and SAAP requirements, see the controller user guide or the Hewlett Packard Enterprise website (<http://www.hpe.com/info/enterprise/docs>).

For more information, see "About SAAP."

Procedure	Smart Array G6 and G7 controllers	Smart Array Gen8 controllers	Smart Array Gen9 controllers (RAID mode)	Smart HBA Gen9 controllers ¹ (RAID mode)
Activate or delete license keys	+	+	+	+
Assign a RAID level to a logical drive	+	+	+	+
Assign spare drives to an array	+	+	+	+
Configure multiple systems identically	+ ²	+ ²	+ ²	+ ²
Configure a RAID 6 logical drive	—	+	+	
Configure a RAID 60 logical drive	—	+	+	
Configure RAID 1 (ADM) and RAID 10 (ADM)	—	+	+	
Copy the configuration of one system to multiple systems	— ²	+ ²	+	+
Create multiple logical drives per array	+	+	+	+
Create or delete arrays and logical drives	+	+	+	+
Enable or disable a physical drive write	+	+	+	+

Procedure	Smart Array G6 and G7 controllers	Smart Array Gen8 controllers	Smart Array Gen9 controllers (RAID mode)	Smart HBA Gen9 controllers ¹ (RAID mode)
cache				
Enable solid state drives to be used as caching devices, using SmartCache	n/a	#	#	
Enable optimized data path to solid state drives using SSA 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)	– ³	+	+	+
Move and delete individual LUNs	–	+	+	+
Optimize the controller performance for video	– ³	+	+	+
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	+	+	+	+

Procedure	Smart Array G6 and G7 controllers	Smart Array Gen8 controllers	Smart Array Gen9 controllers (RAID mode)	Smart HBA Gen9 controllers ¹ (RAID mode)
Removing a drive from an array (restripe the data on an array to occupy fewer physical drives, then remove the excess drives from the array)	+/- ³	+	+	+
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	-	+	+	+ no ADM support

¹Scripting is the most efficient method for this task.

²Having the latest firmware installed allows access all of the supported features.

³The task is supported only from the Configuration screen.

Array configuration

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 (on page 125)."

Minimum requirements

For minimum operating system requirements to run any HPE SSA format, see the Hewlett Packard Enterprise website (<http://www.hpe.com/info/ossupport>).

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 8.0 or later
- Google Chrome

For a list of supported controllers, see Smart Array RAID Controllers on the Hewlett Packard Enterprise website (<http://www.hpe.com/servers/smartarray>).

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.

Operations

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 (Gen8 or later) (on page 10)
- Launching HPE SSA during POST (Gen8 or later) (on page 10)
- Launching HPE SSA from an ISO image (all generations) (on page 10)

To access the HPE SSA CLI or HPE SSA Scripting in an offline environment ("[Accessing HPE SSA in the offline environment](#)" on page 10), 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 (Gen8 or later)

1. Boot the server.
2. Press **F10** to launch Intelligent Provisioning.
3. At the main screen, select **Perform Maintenance**.
4. At the **Maintenance** screen, select **HPE Smart Storage Administrator (HPE SSA)**.
The system launches the HPE SSA GUI.

Launching HPE SSA during POST (Gen8 or later)

For Gen8 servers:

1. Boot the server.
During POST, the system recognizes devices.
2. When the system recognizes a Smart Array controller, press **F5**.
The system launches the HPE SSA GUI, or if you are using Serial Console, the system launches the HPE SSA CLI.

For Gen9 and later servers:

1. Boot the server.
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 (all generations)

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 (on page 11)
- Burning the image to a CD or DVD (on page 11)
- Flashing the image to a USB memory key or SD card on a HPE UEFI bootable server ("[Flashing the image to a USB memory key or SD card on a UEFI bootable server](#)" on page 11)
- Installing the image on a PXE server (on page 12)

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 (<http://www.hpe.com/info/enterprise/docs>).
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.

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 (<http://www.hpe.com/info/enterprise/docs>).
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.

1. Download the ProLiant Offline HPE Smart Storage Administrator ISO image from the Hewlett Packard Enterprise website (<http://www.hpe.com/servers/ssa>).



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 (HPSSA)" 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 (on page 12).
2. Set up PXELinux.
3. Configure PXELinux (on page 12).
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 (<http://syslinux.zytor.com/wiki/index.php/PXELINUX>)

These instructions presume that you are using a Linux TFTP server and the TFTP package (<http://www.kernel.org/pub/software/network/tftp>). 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:

1. Download the ProLiant Offline HPE Smart Storage Administrator ISO image from the Hewlett Packard Enterprise website (<http://www.hpe.com/servers/ssa>).
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/hpssaoffline-2.40.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/hpssaoffline-2.40.iso
iso1mnt=/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/hpssaoffline-2.40.iso
iso1mnt=/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/hpssaoffline-2.40.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/hpssaoffline-1.60-16.0.iso
iso1mnt=/mnt/bootdevice

```

NFS volumes are mounted with the following options:

- -o ro
- nolock
- Windows operating systems:
`isol=smbfs://192.168.0.99/share/path/to/hpssacd/image/hpssaoffline-1.60-16.0.iso isolmnt=/mnt/bootdevice`
- Windows operating systems with logon credentials:
`isol=smbfs://user:password@192.168.0.99/share/path/to/hpssacd/image/hpssaoffline-1.60-16.0.iso isolmnt=/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.

Users familiar with ACU Scripting must now install the HPE SSA CLI application to obtain the scripting executable. The new HPE SSA scripting executable (hpssascripting) replaces the former executable (hpacuscripting) in all scripts.

For information about the minimum monitor settings and the version numbers of supported operating systems and browsers, see the README.txt file provided with the executable.

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:

1. Obtain the executable files from one of the following locations:
 - The Hewlett Packard Enterprise website (<http://www.hpe.com/support/hpesc>), under **Drivers & Downloads**.
When prompted for product information, enter the appropriate server or server blade model name.
 - The software CD that is provided with the controller
2. Follow the installation instructions provided with the executable.
3. After the executables are installed, launch each executable in the following manner:
 - GUI—Click **Start**, and then select **Programs>HPE System Tools>HPE Smart Storage Administrator>Setup HPE Smart Storage Administrator**.
Depending on your configuration scenario, choose one of the following options:
 - Launching HPE SSA on a local server (on page 14)
 - Launching HPE SSA on a local server to configure a remote server (on page 15)
 - Launching HPE SSA on a remote server to configure a local server (on page 16)
 - CLI—Click **Start**, and then select **Programs>HPE System Tools>HPE Smart Storage Administrator CLI**
 - Scripting—Run `hpssascripting.exe` (Windows) or `hpssascripting` (Linux).

Launching HPE SSA on a local server

Microsoft OS

1. Click **Start**, and then select **Programs>HPE System Tools>HPE Smart Storage Administrator>HPE 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 ("[Configuring a controller](#)" on page 24).
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:
 - o For local mode, enter: `hpssa- local`
 - o For remote mode, enter: `hpssa- start`
 HPE SSA launches in a browser window.
2. For a list of options, enter the following:


```
hpssa- h
```

Launching HPE SSA on a local server to configure a remote server

1. On the local server (host), click **Start**, and then select **Programs>HPE System Tools>HPE Smart Storage Administrator>HPE Smart Storage Administrator Preferences**.
The **HPE Smart Storage Administrator Preferences** screen appears.
 - o Optional: Under System Management Homepage Integration, click **Enable**, or leave **Disabled** (default).
 - o Optional: Under Preferred Language, select a language or leave as **System Default**.
 - o Click **Exit**.
2. On the remote server, open the browser.
3. Enter the following text into the address field of the remote browser (where *servername* is the name or IP address of the host):


```
http://servername:2301
```

 The login screen for the System Management Homepage opens.
4. Enter your login credentials:
 - o If you are using version 7.0.0 or later of the System Management Homepage, use your operating system user name and password.
 - o If you are using an earlier version of the System Management Homepage, use your WBEM user name and password.
 The System Management Homepage opens.
For more information about the System Management Homepage, see the following:
 - o The Hewlett Packard Enterprise System Management Homepage webpage (<http://www.hpe.com/info/smh>)
 - o The *HP System Management Homepage Installation Guide* on the Hewlett Packard Enterprise website (<http://www.hpe.com/info/enterprise/docs>)
5. Click **HPE Smart Storage Administrator** on the left side of the screen.
HPE SSA opens, scans the remote server, and detects controllers. When controller detection is complete, the controllers are available on the **Controller/Device** menu.

6. Configure a controller ("[Configuring a controller](#)" on page 24).
7. (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.

Launching HPE SSA on a remote server to configure a local server

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

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

- Optional: Under HP System Management Homepage Integration, click **Enable**, or leave **Disabled** (default).
 - Optional: Under Preferred Language, select a language or leave as **System Default**.
 - Click **Exit**.
2. On the server that you want to configure, connect to the Systems Insight Manager server (port: 280), and then log in.
 3. Select **Device Queries**.
 4. Under **Device by Type**, select **All Servers**.
 5. Connect to the server that is running HPE SSA.
 6. Under **Device Links**, select **System Management Homepage**.
- The login screen for the System Management Homepage opens.
7. Log in using your credentials:
 - If you are using version 7.0.0 or later of the System Management Homepage, use your operating system user name and password.
 - If you are using an earlier version of the System Management Homepage in Windows, use your Wbem user name and password.

The System Management Homepage opens.

For more information about the System Management Homepage, see the following:

- The Hewlett Packard Enterprise System Management Homepage (<http://www.hpe.com/info/enterprise/docs>)
 - The *HP System Management Homepage Installation Guide* on the Hewlett Packard Enterprise website (<http://www.hpe.com/info/enterprise/docs>)
8. Click **HPE Smart Storage Administrator** on the left side of the screen.
- HPE SSA opens, scans the remote server, and detects controllers. When controller detection is complete, the controllers are available on the **Controller/Device** menu.
9. Configure a controller ("[Configuring a controller](#)" on page 24).
- When configuration is complete, continue with the next step.
10. (Optional) To make newly created logical drives available for data storage, in a Windows OS, use the operating system disk management tools to create partitions and format the drives.

Using the HPE SSA GUI

Access HPE SSA with one of the many methods available:

- Accessing HPE SSA in the offline environment (on page 10)
- 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 (on page 18)."

Icon and key stroke legend

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 Drive
	Assigned Physical Drive
	Unassigned Physical Drive
	Unassigned Drives
	Spare Drive
	Transient Drive
	Storage Enclosure
	Tape Drive
	Locked
	License Manager/Encryption Manager
	Cache Manager

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

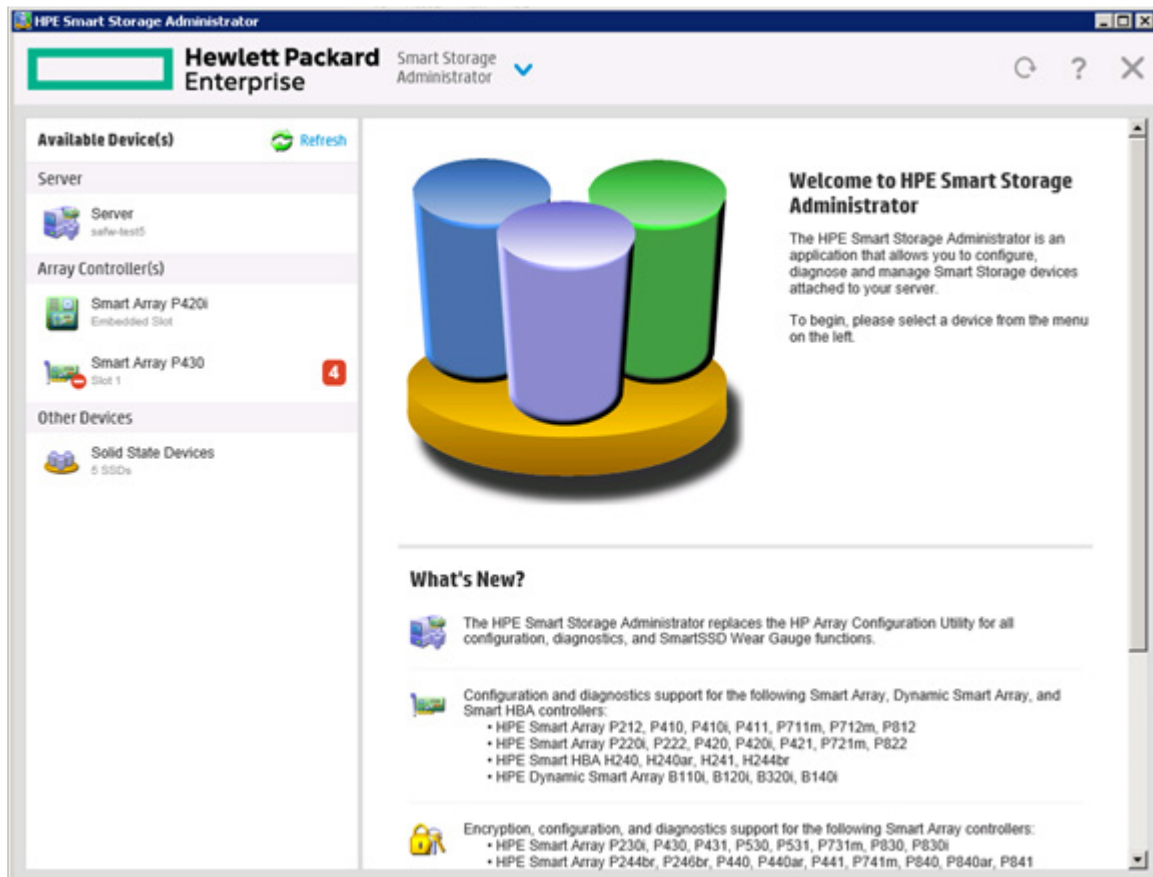
Key	Description
Tab	Cycle through selectable items on a page
Shift + Tab	Cycle backwards through selectable items on a page
F5	Rescan system (equivalent to clicking the Rescan System button)

Key	Description
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 non-action popups*
R	Refresh selected controller*

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

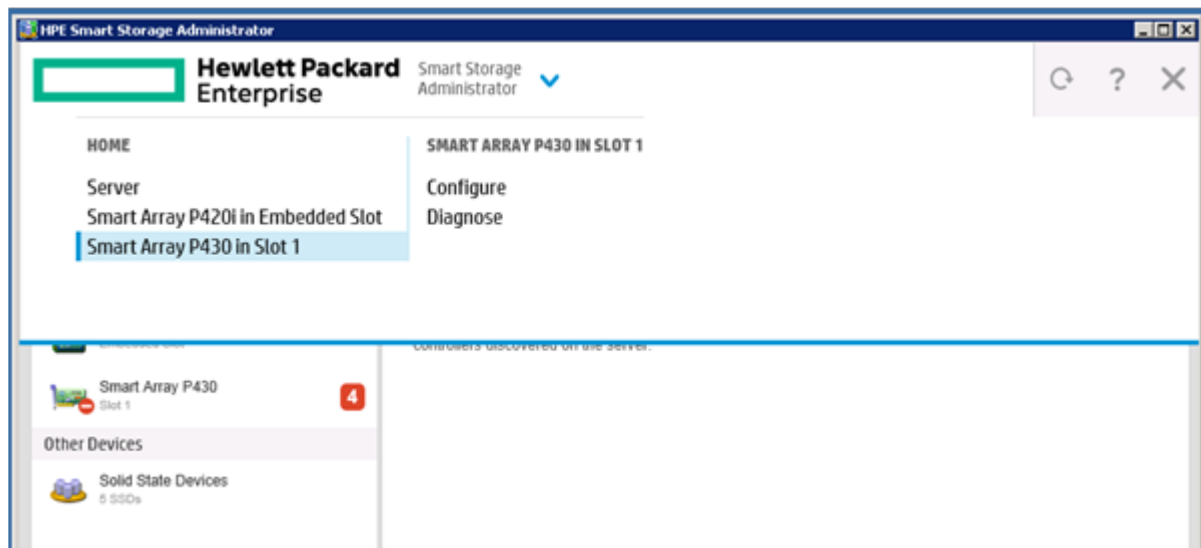
Navigating the GUI

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 (on page 19)" or "Diagnostics screen (on page 21)."

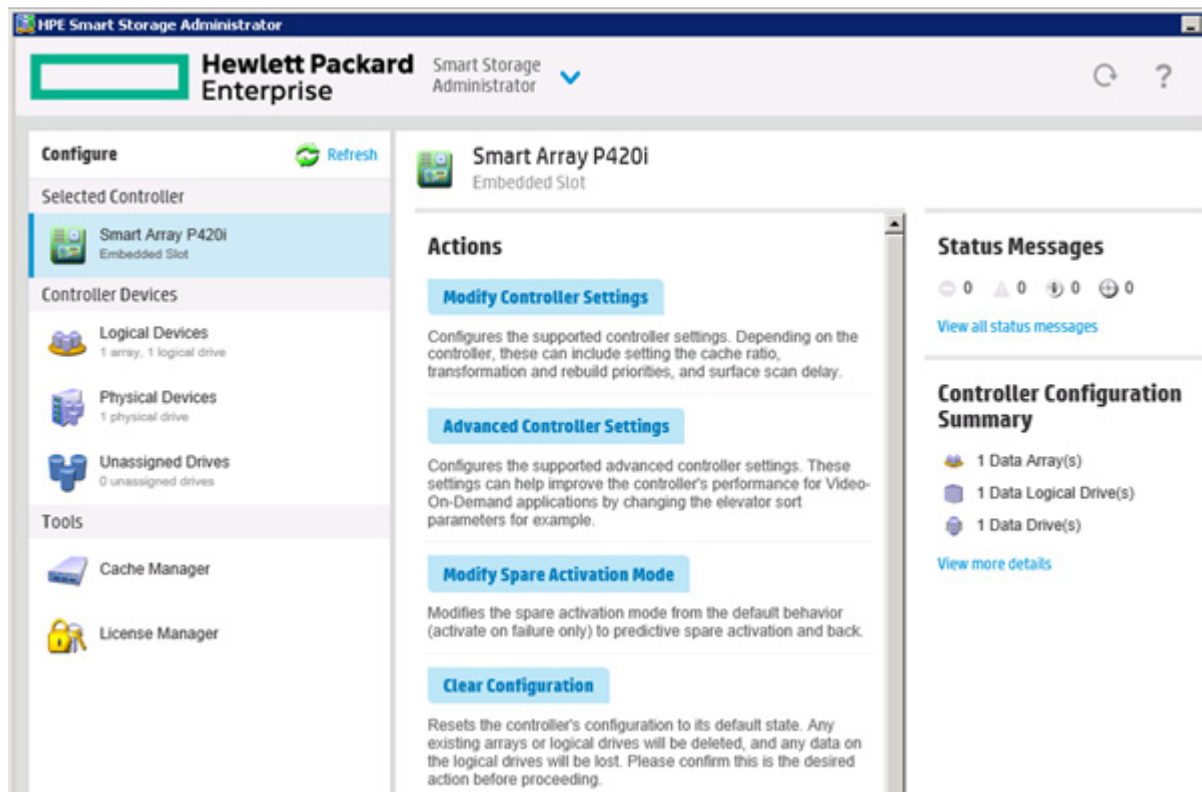


- **Available device(s)** are listed on the left-hand side of the screen. Clicking on a server or array controller displays the available actions, alerts, and summary for that device. You can point to the status alerts to see details on an alert.
- **What's New?** summarizes the changes since Array Configuration Utility became HPE Smart Storage Administrator, and since the previous versions of HPE SSA.
- The **Rescan System** button is near the top right of the screen.
After adding or removing devices, click **Rescan System** 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. (["HPE SSA help"](#) on page 22)"
- The **Exit HPE SSA** button is near the top right of the screen.

Configure screen

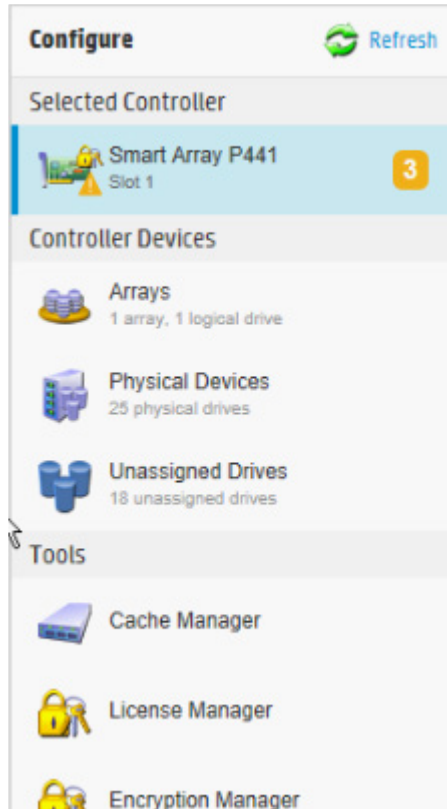
To access this screen, click either a device under **Configuration** in the quick navigation menu, or select an available device from the **Home** screen, and then click **Configure** under the available options.

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:

- **Devices and Tools**—This panel, at left, displays systems, controllers, arrays, physical devices, unassigned drives, and cache, license and 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:
 - Data arrays
 - Data logical drives
 - Data drives
 - Unassigned drives
 - Spare drives

For a list of possible tasks that are available on the **Configure** screen, see "Configuration tasks (on page 23)."

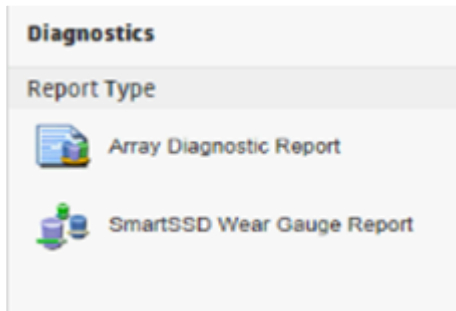
Diagnostics screen

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 (on page 58)."

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 Controls**—An explanation and list of keyboard functions for navigating the GUI
- **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 HPE SSA."

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

Configuration tasks

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

For certain tasks, the controller must have SAAP activated by a registered license key. For more information, see "About SAAP."

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.

Item	Tasks
Controller	Accelerated I/O Path Advanced Controller Settings* ** Array Accelerator Settings Clear Configuration Controller Settings Create Array Disable Standby Controller Enable HBA/RAID/Smart Array Mode operations* Enable SmartCache† Manage Encryption Manage License Keys* Modify Power Modes* More Information Parallel Surface Scan Physical Drive Write Cache Settings Redundancy Settings* 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 Move Drives** Re-Mirror Array** Shrink Array** Spare Management Split Mirrored Array** View Status Alerts Volume Key Rekey

Item	Tasks
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* ** 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 drives	Create Array Erase Drive** More Information

*This task is not available on all controller models.

**If performed with an Smart Array G6 or G7 controller, this task requires a registered SAAP license key or a controller where SAAP functionality is standard. See "About SAAP."

†Only available on a Gen8 or Gen9 controller with a registered SmartCache license key.

Configuring a controller

1. Open HPE SSA.


For more information, see "Using the HPE SSA GUI (on page 16)."

2. Open the **Configure** panel by doing one of the following:
 - o Choose a device and click **Configure** in the quick navigation menu.
 - o 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.

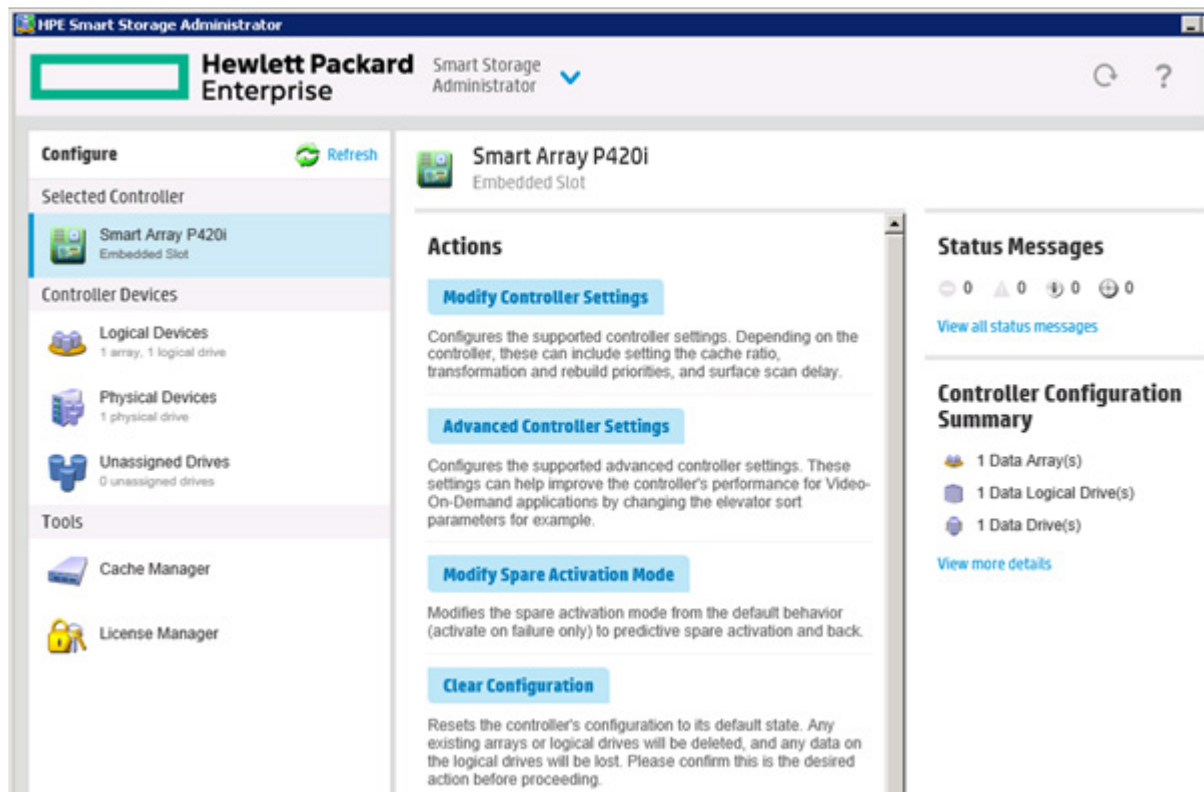


3. Configure the controller. See "Performing a Configuration task (on page 25)."
4. When prompted, save the configuration.
5. Do one of the following:
 - o Configure an additional controller. Repeat steps 3 through 5.
 - o Click **Exit HPE SSA**.

Performing a Configuration task

1. Open HPE SSA.
For more information, see "Using the HPE SSA GUI (on page 16)."
2. Open the **Configure** panel by doing one of the following:
 - o Choose a device and click **Configure** in the quick navigation menu.
 - o 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 (on page 23)."



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**.

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

The command values are `Enable` or `Disable`.

Example:

- `SSDSmartPath= Enable | Disable`

To enable or disable HPE SSD Smart Path:

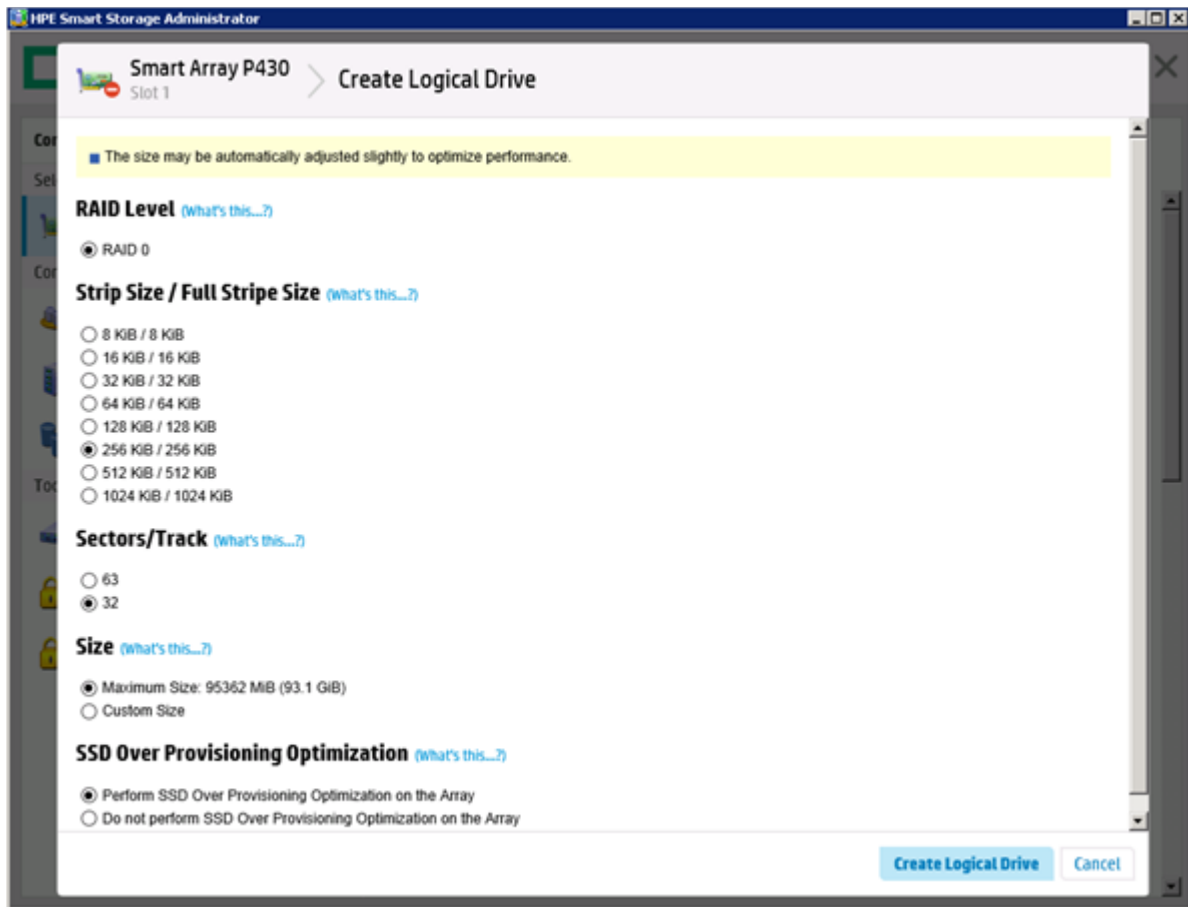
1. Open HPE SSA.

For more information, see "Using the HPE SSA GUI (on page 16)."

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.

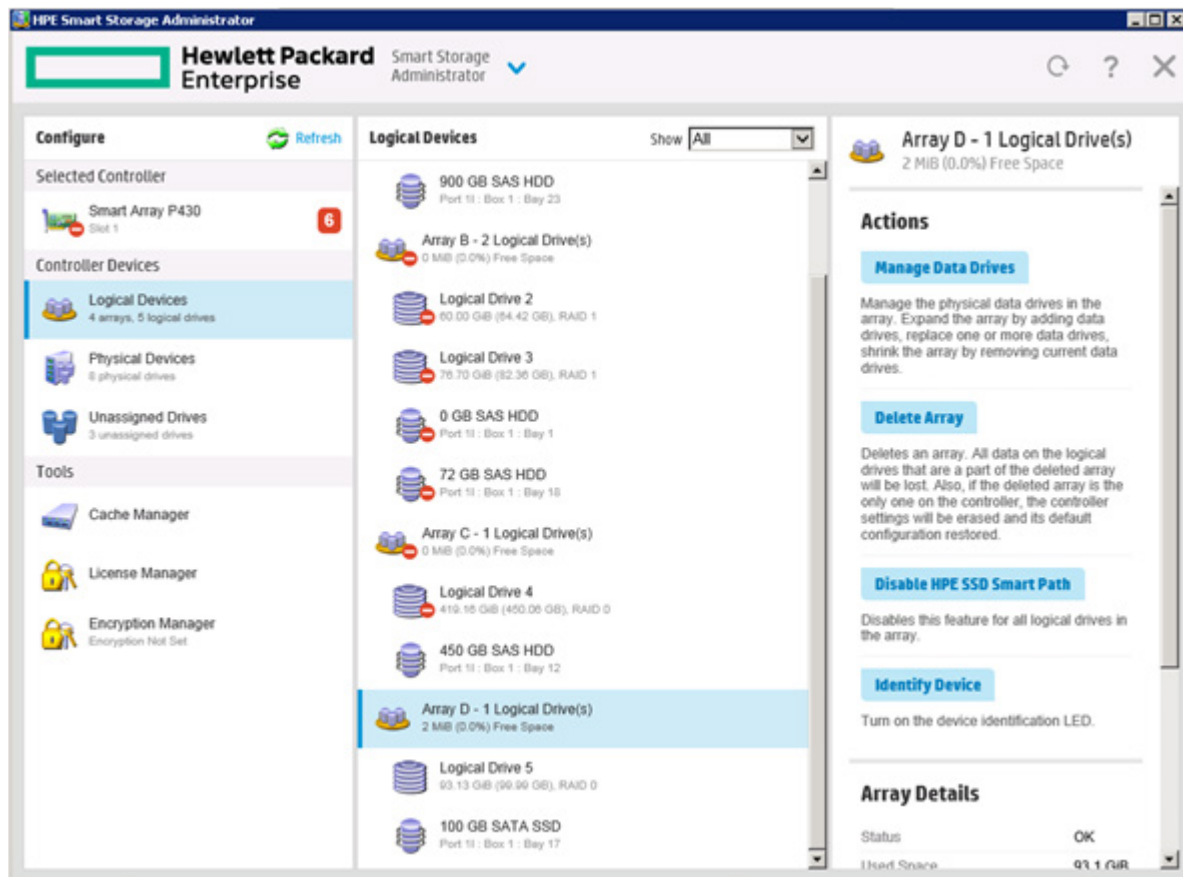
Create a logical drive:

1. Click **Create Array**.
2. Select the physical drives for the new array.
3. Click **Create Array**.
4. Make selections for **RAID Level**, **Strip Size/Full Stripe Size**, **Sectors/Track**, and **Size**.



5. Click **Create Logical Drive**.
6. Click **Finish**.
By default, HPE SSD Smart Path is enabled on SSD drives.
7. Under **Configure>Logical Devices**, choose the array with the logical drives you just created.

- If you want to disable HPE SSA Smart Path, click **Disable HPE SSD Smart Path**.



- Click **Save**.

Rapid Parity Initialization

When you create a logical drive, you must initialize the parity using Rapid Parity Initialization.

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.

To select the method for parity initialization:

- Open HPE SSA.
For more information, see "Using the HPE SSA GUI (on page 16)."
- 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.

The screenshot displays a RAID configuration interface with the following sections and options:

- RAID Level** (What's this...?):
 - ☐ RAID 0
 - ☐ RAID 1+0
 - ☐ RAID 5
 - ☒ RAID 50
- Number of Parity Groups (NPG)** (What's this...?):
 - 2 (selected in a dropdown menu)
- 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**.

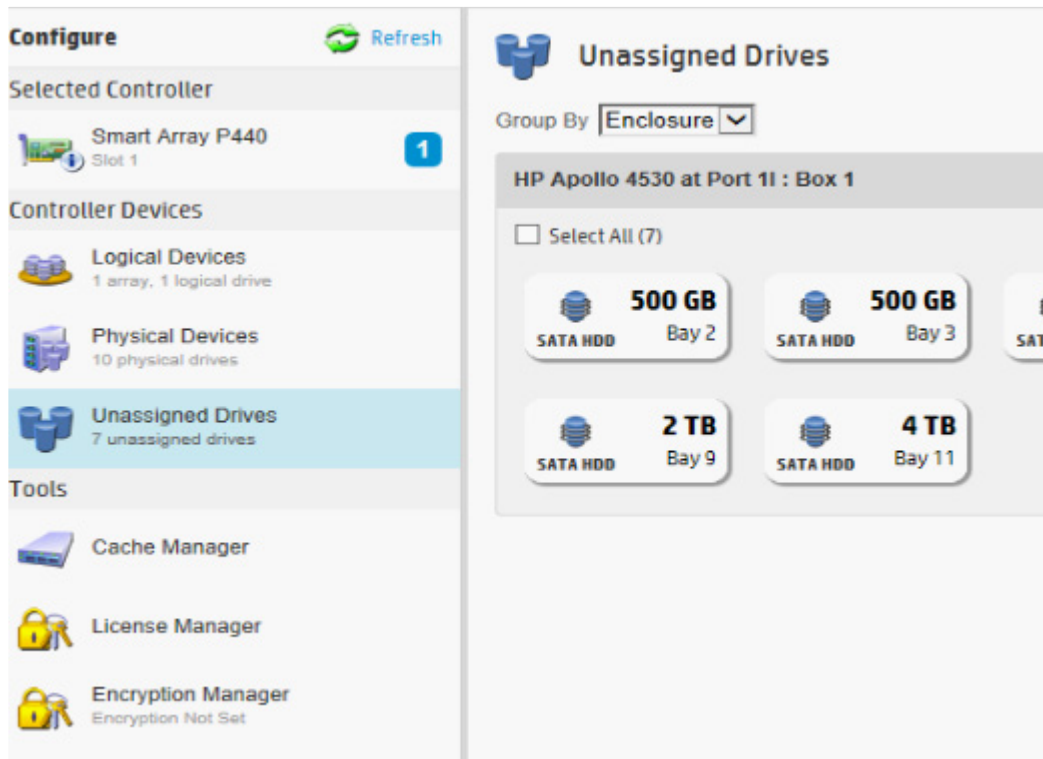
SSD Over Provisioning Optimization

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.

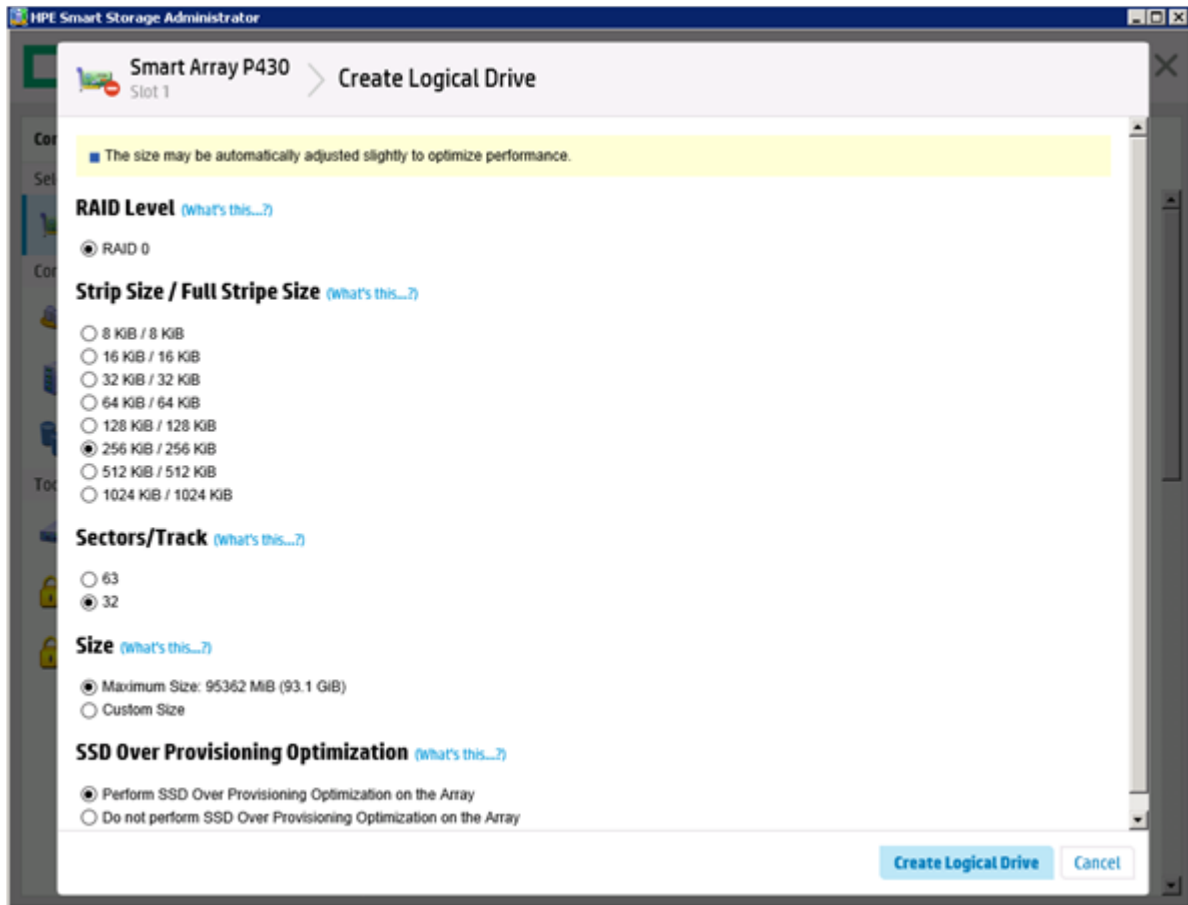
The SSD Over Provisioning Optimization feature allows the user to disable Over Provisioning Optimization in the GUI.

To disable Over Provisioning Optimization:

1. Open HPE SSA.
For more information, see "Using the HPE SSA GUI (on page 16)."
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**.

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 (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.

To change the Spare Activation Mode:

1. Open HPE SSA.
For more information, see "Using the HPE SSA GUI (on page 16)."
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.

To change the Spare Management mode:

1. Open HPE SSA.
For more information, see "Using the HPE SSA GUI (on page 16)."
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**.

The screenshot shows the 'Manage Spare Drives' page for a Smart Array P430. At the top, there's a header with the Smart Array P430 logo and 'Slot 1'. Below the header, a yellow warning box contains two blue square icons and text: 'In a dual domain configuration, mixing single and dual ported SAS drives can lead to a loss of redundancy.' and 'To avoid wasting drive capacity, select physical drives that are same size as the existing drives in the array.' A 'Hide' link is on the right. The main section is titled 'Spare Drive Type (What's this...?)' and has two radio buttons: 'Dedicated Spare Drives. Drives designated as spares can only be used as spare drives.' (selected) and 'Auto Replace Drives. When activated, selected drives will automatically become part of the array.' Below this is 'Add/Remove Spare Drives (What's this...?)' with a 'Group By' dropdown set to 'Enclosure'. There are two sections for 'Internal Drive Cage at Port 11 : Box 1'. The first section has a 'Select All (1)' checkbox and one drive icon: '600 GB SAS HDD Bay 8'. The second section has a 'Select All (2)' checkbox and two drive icons: '600 GB SAS HDD Bay 2' and '600 GB SAS HDD Bay 3'.

8. From the menu, select one of the following spare drive types:
 - o Dedicated Spare Drives
 - o Auto-Replace Drives
9. Select which drives will operate as spare drives in the array.

10. Click **Save**.

Smart Array P431 Slot 5 > Create Logical Drive

■ Logical Drive was successfully created. Please choose one of the actions below.

Array Details

Status	OK
Used Space	205.0 GiB (100.0%)
Total Usable Space	205.0 GiB
Acceleration Mode	Independent. Caching can be enabled or disabled for each individual logical drive

Logical Drives

Logical Drive 2	136.67 GiB (146.75 GiB)
-----------------	-------------------------

Physical Drives

146 GB SAS HDD at Port 2E : Box 1 : Bay 4
72 GB SAS HDD at Port 2E : Box 1 : Bay 34
72 GB SAS HDD at Port 2E : Box 1 : Bay 35

Device Path

Smart Array P431 in Slot 5

Manage Spare Drives Finish

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

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 (on page 128)."

To create an array with RAID 0:

1. Open HPE SSA.
For more information, see "Using the HPE SSA GUI (on page 16)."
2. Select the controller.

3. Click **Create Arrays with RAID 0**.

The screenshot displays the HPE Smart Storage Administrator web interface. The top navigation bar includes the HPE logo, "Hewlett Packard Enterprise", and "Smart Storage Administrator". The left sidebar is titled "Configure" and contains a "Refresh" button. Under "Selected Controller", the "Smart Array P430 Slot 1" is highlighted with a red "6" badge. The "Controller Devices" section lists: "Logical Devices" (6 arrays, 7 logical drives), "Physical Devices" (8 physical drives), and "Unassigned Drives" (1 unassigned drive). The "Tools" section includes "Cache Manager", "License Manager", and "Encryption Manager" (Encryption Not Set). The main content area is titled "Smart Array P430 Slot 1" and features an "Actions" section with four buttons: "Create Array", "Create Arrays with RAID 0", "Modify Controller Settings", and "Advanced Controller Settings". Each button has a corresponding description of its function.

HPE Smart Storage Administrator

Hewlett Packard Enterprise Smart Storage Administrator

Configure Refresh

Selected Controller

Smart Array P430 Slot 1 6

Controller Devices

- Logical Devices**
6 arrays, 7 logical drives
- Physical Devices**
8 physical drives
- Unassigned Drives**
1 unassigned drive

Tools

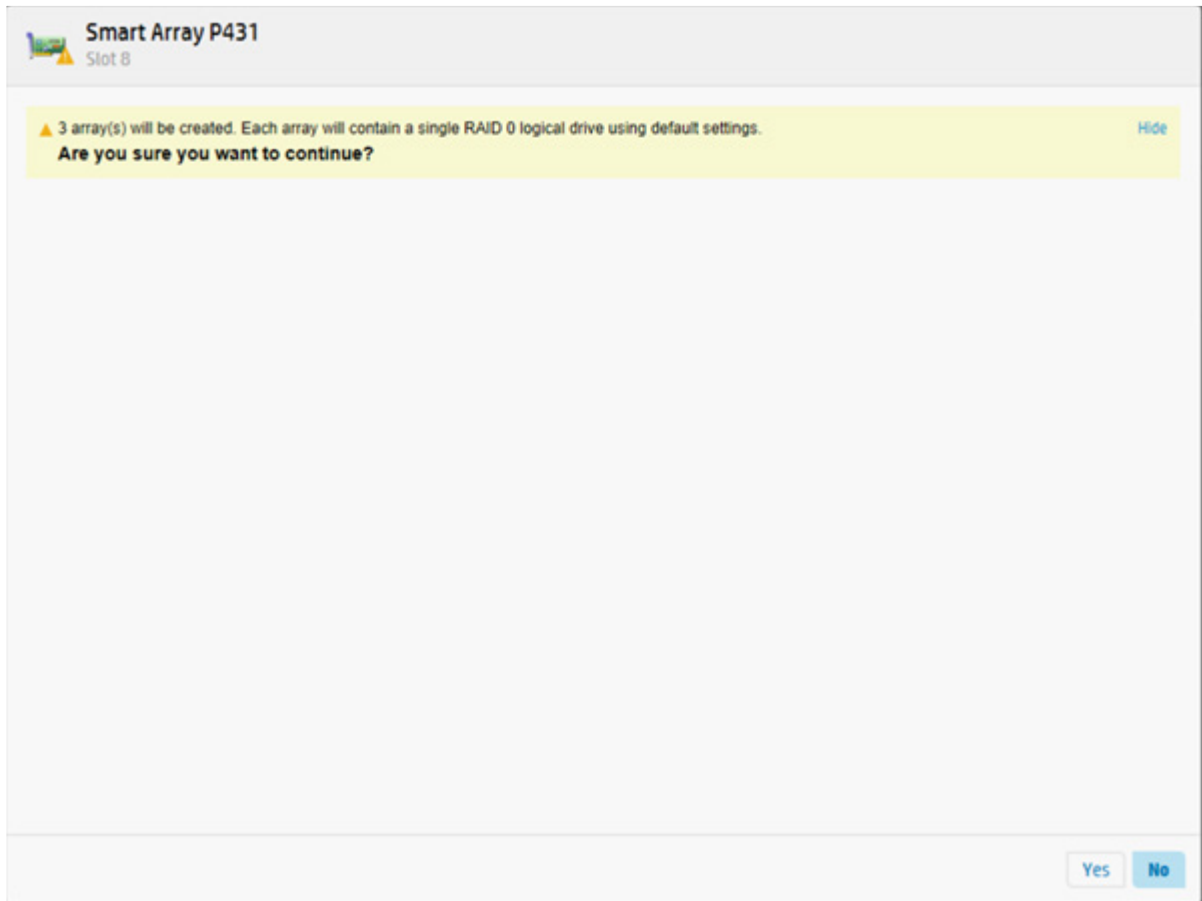
- Cache Manager**
- License Manager**
- Encryption Manager**
Encryption Not Set

Smart Array P430 Slot 1

Actions

- Create Array**
Creates an array from a group of selected physical drives. Once an array is created, the unused space is available for creating logical drives. A logical drive cannot be created unless an array exists.
- Create Arrays with RAID 0**
Create arrays from a group of selected physical drives. Each array contains exactly one physical drive and one RAID 0 logical drive.
- Modify Controller Settings**
Configures the supported controller settings. Depending on the controller, these can include setting the cache ratio, transformation and rebuild priorities, and surface scan delay.
- Advanced Controller Settings**
Configures the supported advanced controller settings. These settings can help improve the controller's performance for Video-On-Demand applications by changing the elevator sort parameters for example.

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.

Smart Array P431
Slot 8

Create Arrays with RAID 0

■ Operation executed successfully.

Controller Hardware Information

Controller	Smart Array P431
Slot	8
Negotiated PCIe Data Rate	PCIe 3.0 x8 (7880 MB/s)
Hardware Revision	A
Controller Serial Number	PCZED0XTA3R07F
Firmware Version	1.49-160
Number of Ports	2 (External only)
Total Cache Size	2048 MiB
Battery/Capacitor Pack Count	1

Controller Settings

Transformation Priority	Medium
Rebuild Priority	High
Surface Scan Analysis Priority	Idle
Surface Scan Delay	3 sec
Physical Drive Write Cache State	Disabled
Spare Activation Mode	Activate on Physical Drive Failure (default)

Settings

Finish

Parallel Surface Scan

1. Open HPE SSA. For more information, see "Using the HPE SSA GUI (on page 16)."
2. Select a controller.

3. Click **Modify Controller Settings**.

Actions

Create Array

Creates an array from a group of selected physical drives. Once an array is created, the unused space is available for creating logical drives. A logical drive cannot be created unless an array exists.

Create Arrays with RAID 0

Create arrays from a group of selected physical drives. Each array contains exactly one physical drive and one RAID 0 logical drive.

Modify Controller Settings

Configures the supported controller settings. Depending on the controller, these can include setting the cache ratio, transformation and rebuild priorities, and surface scan delay.

Advanced Controller Settings

Configures the supported advanced controller settings. These settings can help improve the controller's performance for Video-On-Demand applications by changing the elevator sort parameters for example.

Modify Spare Activation Mode

Modifies the spare activation mode from the default behavior (activate on failure only) to predictive spare activation and back.

4. Under **Current Parallel Surface Scan Count**, use the slider to select the parallel surface scan count.

Smart Array P440 Slot 1 > Modify Controller Settings

Transformation Priority (What's this...?)

High
Medium
Low

Rebuild Priority (What's this...?)

High
Medium High
Medium
Low

Surface Scan Analysis Priority (What's this...?)

Disabled
High
Idle (with delay)

3 Valid Range: 0 (Disabled) - 30 seconds

Current Parallel Surface Scan Count (What's this...?)

1

Physical Drive Write Cache State (What's this...?)

5. Click **Save Settings**.
6. A summary page appears. Click **Finish** to exit.

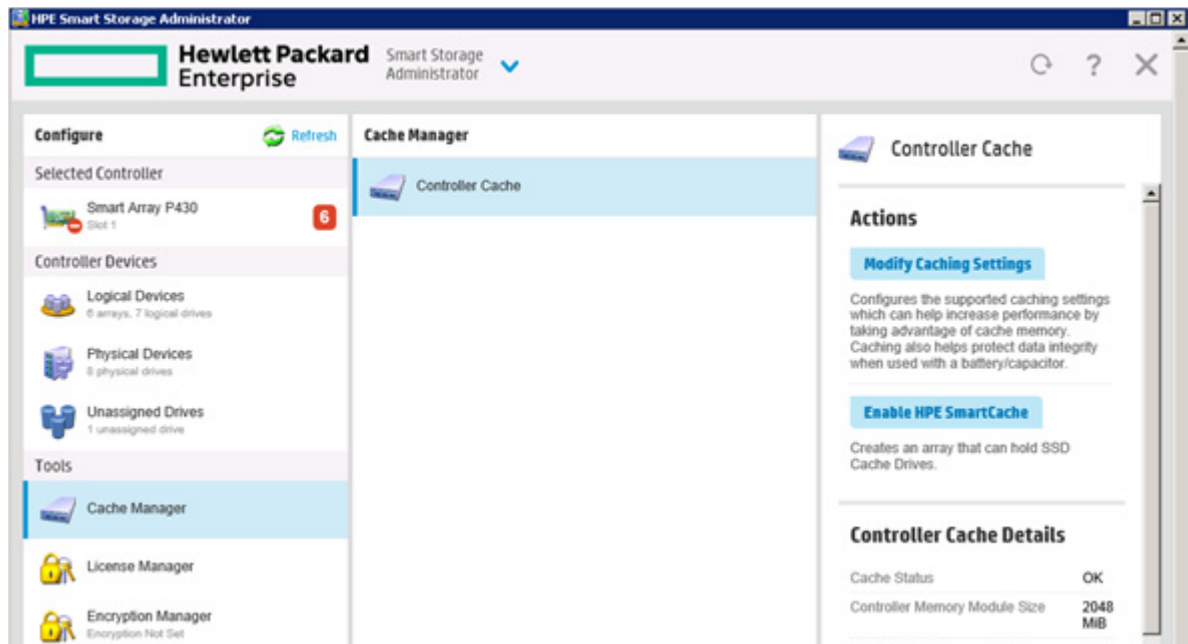
Cache Manager

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.

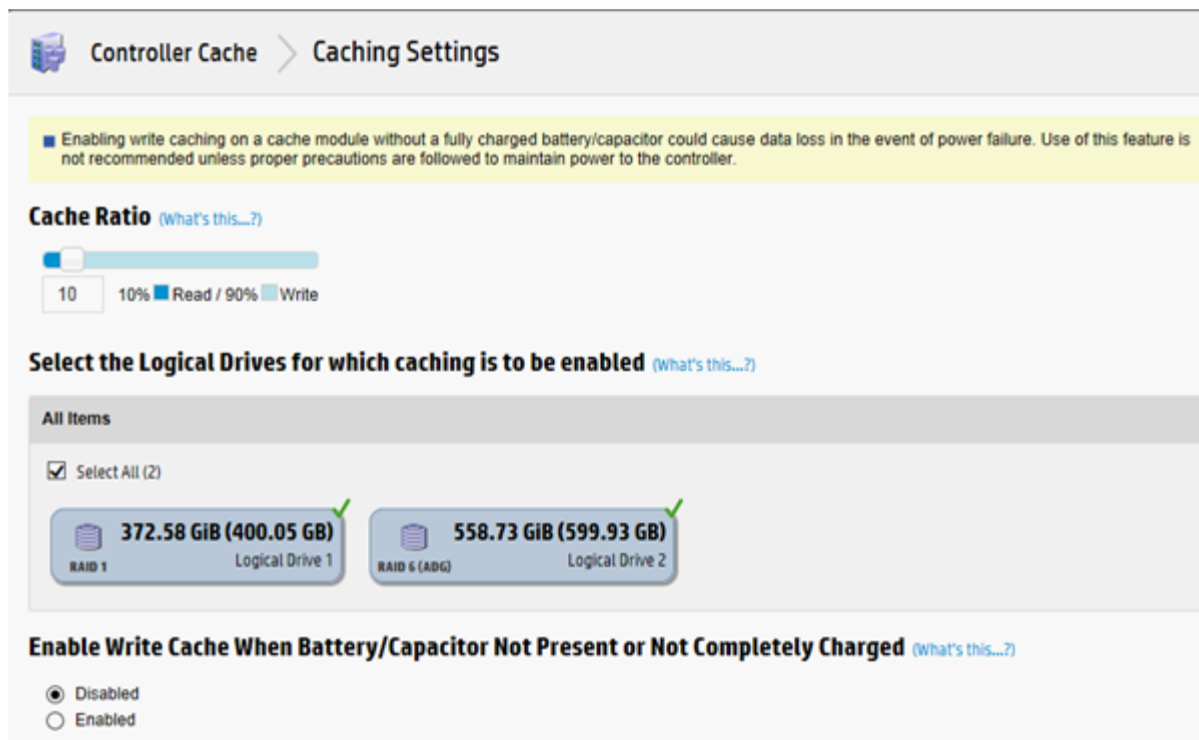
To configure the controller cache:

1. Open HPE SSA.
For more information, see "Using the HPE SSA GUI (on page 16)."
2. Open the **Configure** panel by doing one of the following:
 - o 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**.

About HPE 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 where Smart Array Gen9 controllers are supported, without the need for changes to OS, driver, or applications
- Choice of write-through or write-back cache (Gen9 servers only)

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

To support SmartCache, the Smart Array Controller firmware must meet the following version minimums:

- 6G controller: version 3.42 or later
- 12G controller: all releases

SmartCache requires a SmartCache license. For more information, see the Hewlett Packard Enterprise website (<http://www.hpe.com/servers/smartcache>).

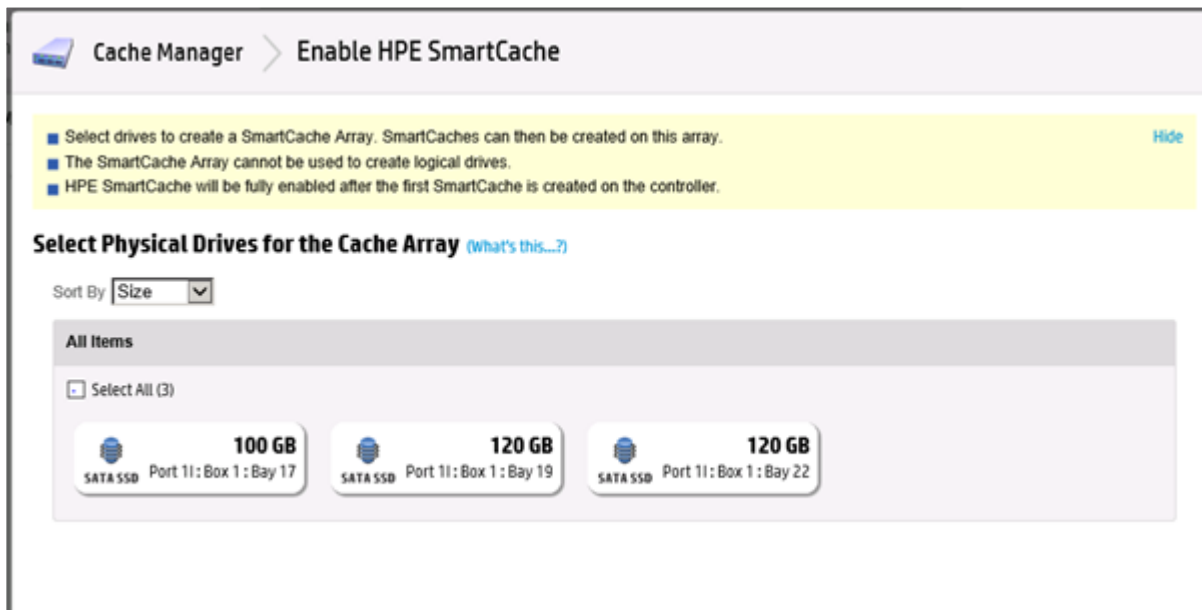
Enable HPE SmartCache

Enabling HPE 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.

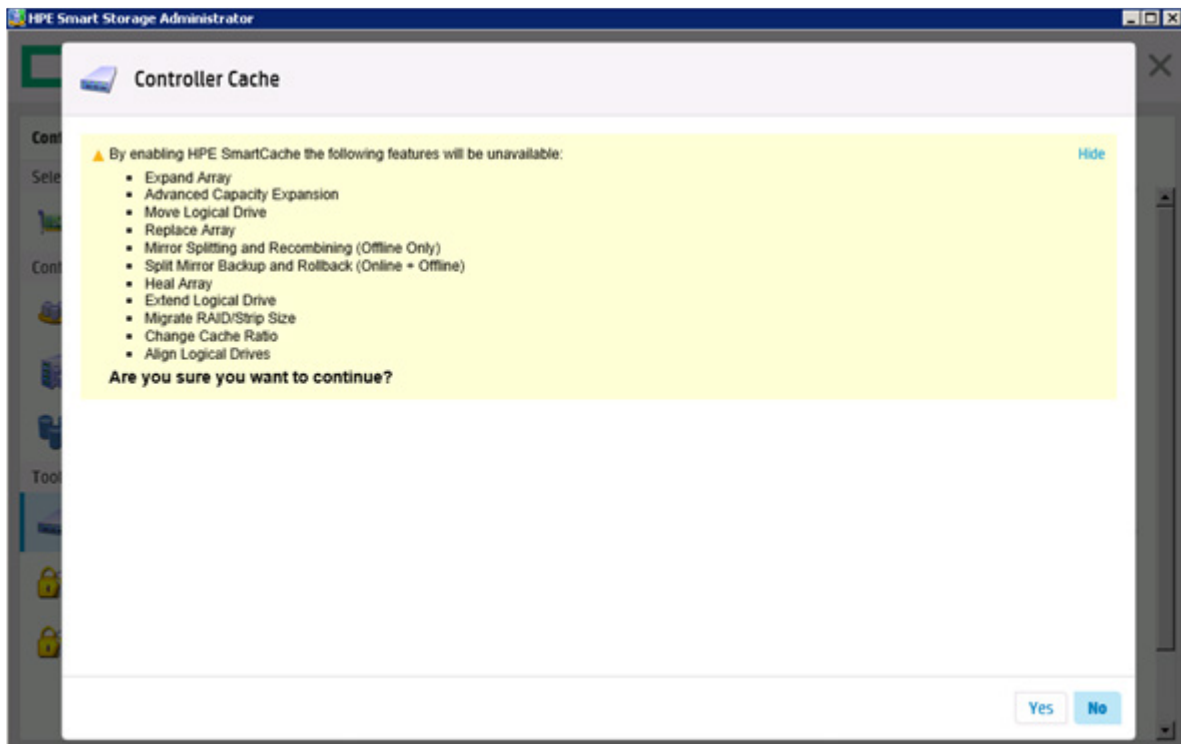
To enable SmartCache:

1. Open HPE SSA.
For more information, see "Using the HPE SSA GUI (on page 16)."
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.

- Click **Enable HPE SmartCache** in the **Actions** menu.



- Select one or more physical drives from the list of available drives.
- Click **OK**.
- A pop-up window appears, indicating certain features will not be available with SmartCache enabled. If you want to continue, click **Yes**.



8. **Create SmartCache for Logical Drive** appears.

Cache Manager > Create SmartCache for Logical Drive

■ The SmartCache will improve the read performance of a logical drive. [Hide](#)
■ If using write-back cache write policy, the HP SmartCache cannot be deleted until its cache write policy is converted to write-through cache write policy.
■ Specifying the write-back cache write policy when using a RAID 0 cache volume could result in data loss in the event of failure of the cache volume.

Logical Drive to Cache [\(What's this...?\)](#)

☒ Logical Drive 1 (136.70 GiB (146.78 GB), RAID 0)

Cache Write Policy and RAID Type [\(What's this...?\)](#)

☐ Write-Back with RAID 0
☐ Write-Back with RAID 1+0
☒ Write-Back with RAID 5
☐ Write-Through with RAID 0
☐ Write-Through with RAID 1+0

Size [\(What's this...?\)](#)

☒ Recommended Size: 16384 MiB (16 GiB)
☐ Maximum Size: 139979 MiB (136.6 GiB)
☐ Custom Size

Create SmartCache for Logical Drive Cancel

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 screenshot shows the 'Cache Manager' interface with the title 'Create SmartCache for Logical Drive'. A yellow message bar at the top states: 'Cache Logical Drive was successfully created. Please choose one of the actions below.' Below this, there are three sections:

SmartCache Details

Status	OK
Cache Write Policy	Write-Through
RAID Level	RAID 0
Drive Number	3
Size	83.83 GiB (90.01 GB)

SmartCache Statistics

Read Cache Hits	0
Read Cache Misses (Total)	0
Read Cache Hit Rate	0
Write Cache Hits	0
Write Cache Misses (Total)	0
Write Cache Hit Rate	0

Device Path

Smart Array P430 in Slot 1
SmartCache Array C - 1 SmartCache(s)

The SmartCache is created for the Logical Drive.

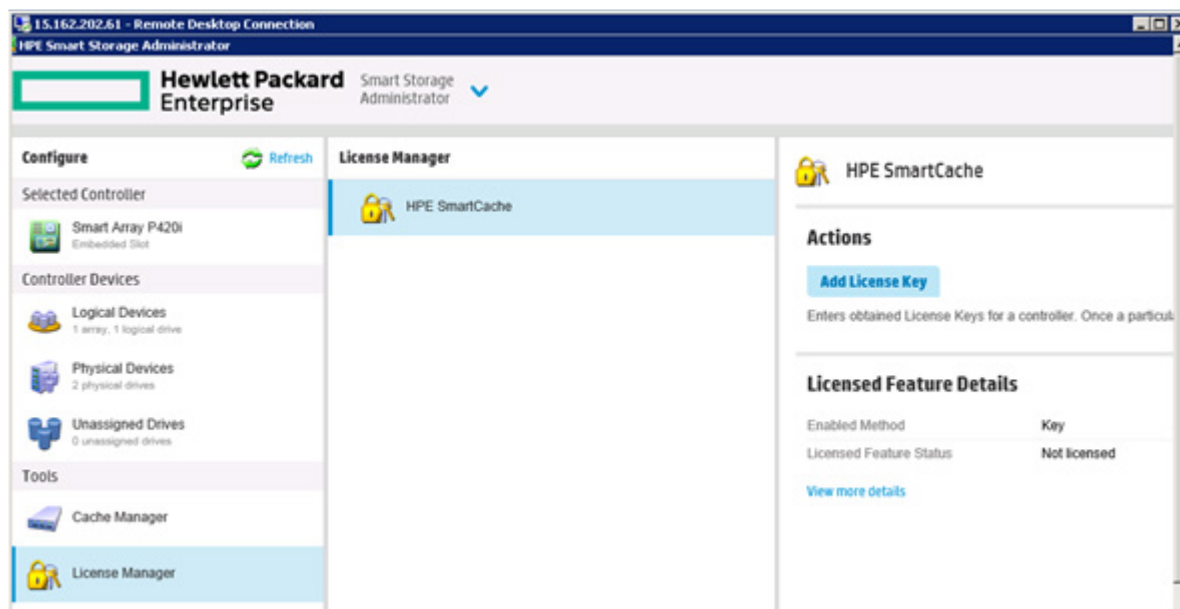
Installing a license key with HPE SmartCache

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

To install a license key:

1. Open HPE SSA.
For more information, see "Using the HPE SSA GUI (on page 16)."
2. Open the **Configure** panel by doing one of the following:
 - o Choose a device and click **Configure** in the quick navigation menu.
 - o Select an available device from the **Home** screen, and then click **Configure** under the available options.
3. In the **Tools** menu, click **License Manager**.

- Under actions, click **Add License Key**.



- Enter the license key number.
- Click **Save**.

Enabling 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.

To modify the logical drive cache policy settings:

1. Open HPE SSA.
For more information, see "Using the HPE SSA GUI (on page 16)."
2. Select the controller.
3. Click **Modify Controller Settings**.
4. Under **Physical Drive Write Cache State**, select one of the following:
 - o Enabled
 - o Disabled

Smart Array P441 Slot 2 > Modify Controller Settings

High
Medium
Low

Rebuild Priority (What's this...?)

☒ High
☐ Medium High
☐ Medium
☐ Low

Surface Scan Analysis Priority (What's this...?)

☐ Disabled
☐ High
☒ Idle (with delay)

3 Valid Range: 0 (Disabled) - 30 seconds

Current Parallel Surface Scan Count (What's this...?)

1

Physical Drive Write Cache State (What's this...?)

☐ Enabled
☒ Disabled

5. Click **Save Settings**.
6. A summary page appears. Click **Finish** to exit.

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 ("Accessing HPE SSA in the offline environment" on page 10).
- Smart Array G6 and G7 controller models must have a valid SAAP license.
- 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 (<http://www.hpe.com/info/enterprise/docs>).

Splitting a mirrored array

1. Run the HPE SSA GUI in offline mode. See "Accessing HPE SSA in the offline environment (on page 10)."
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

1. Run the HPE SSA GUI in offline mode. See "Accessing HPE SSA in the offline environment (on page 10)."
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.

To create a split mirror backup:

1. Run the HPE SSA GUI in offline mode. See "Accessing HPE SSA in the offline environment (on page 10)."
2. Open the **Configure** panel by doing one of the following:
 - o Choose a device and click **Configure** in the quick navigation menu.
 - o 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:

- o If the original array contained RAID 1 or RAID 1+0 drives, then the primary array will contain RAID 0 drives.
 - o If the original array contained RAID 1 (ADM) drives, the primary array will contain RAID 1 drives.
 - o If the original array contained RAID 10 (ADM) drives, the primary array will contain RAID 1+0 drives.
 - o The backup array always contains RAID 0 logical drives.
 - o 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

1. Run the HPE SSA GUI in offline mode. See "Accessing HPE SSA in the offline environment (on page 10)."
2. Open the **Configure** panel by doing one of the following:
 - o Choose a device and click **Configure** in the quick navigation menu.
 - o 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.

Healing an array

The Heal Array operation enables you to enter a command to replace failed physical drives in the array with healthy physical drives. After replacement, the original array and logical drive numbering is unaffected.

To use Heal Array, you must meet the following conditions:

- 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.
- The replacement physical drives and the original drives must be of the same interface type, such as SAS, SATA, and so on.
- A sufficient number of unassigned physical drives of the correct size are available to replace each failed physical drive in the array.

The correct size is defined as a drive as large as the smallest drive on the array, but no larger than the smallest spare.

- If using an ProLiant Smart Array G6 or G7 controller, an SAAP 1.0 license is activated.

When you select Heal Array, and all conditions are met, one of the following actions occurs to the volume in the array:

- If a volume has failed, HPE SSA recreates the volume. This action occurs with a RAID 0 volume.
- If a volume is degraded, HPE SSA rebuilds the volume.

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

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 PX3X and PX4X controllers and Smart HBAs, operating in RAID mode.

For more information about performing a full setup and configuring Secure Encryption, see the *Secure Encryption Installation and User Guide*.

Managing modes

RAID mode or HBA mode

The controller operates in either RAID mode or HBA mode, but not both. Smart Array Controllers allowing HBA mode can be configured to run in RAID mode. In HBA mode, all physical drives are presented directly to the operating system and the hardware RAID engine is disabled.

When HBA mode is enabled, RAID mode is disabled. To enable RAID mode, disable HBA mode.

HBA mode



IMPORTANT: Enabling HBA Mode exposes the physical drives to the operating system. Smart Array configuration is not allowed.

To enable HBA mode:

1. Open HPE SSA. For more information, see "Using the HPE SSA GUI (on page 16)."
2. Select the controller.
3. Click **Configure**.

- Click **Enable HBA mode**.

Smart Array P440ar
Embedded Slot

Actions

Create Array

Creates an array from a group of selected physical drives. Once an array is created, the unused space is available for creating logical drives. A logical drive cannot be created unless an array exists.

Create Arrays with RAID 0

Create arrays from a group of selected physical drives. Each array contains exactly one physical drive and one RAID 0 logical drive.

Enable HBA Mode

Enables HBA mode. A reboot is required for changes to take effect. Any drives configured using a Smart Array will not be available in HBA mode until the configuration data is cleared on the drive.

Modify Power Mode

Modifies the controller's power mode. A reboot may be required after changing power modes to optimize power savings and performance.

Manage Device Identification LEDs

Turn the physical drive identification LED(s) On or Off

Status Messages

0 7 0 0

[View all status messages](#)

Controller Configuration Summary

0 Data Array(s)

7 Unassigned Drive(s)

[View more details](#)

- A prompt appears warning you that entering HBA mode will disable any drives configured using a Smart Array until the configuration is cleared. A reboot is required for HBA mode to be enabled.

Smart Array P440ar
Embedded Slot

Enable HBA Mode

⚠ Enabling HBA mode will disable any drives configured using a Smart Array until the configuration data is cleared. A reboot is required for HBA mode to be enabled. [Hide](#)

Are you sure you want to continue?

[OK](#) [Cancel](#)

If you want to continue, click **OK**.

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

RAID mode

To enable RAID mode:

- Open HPE SSA. For more information, see "Using the HPE SSA GUI (on page 16)."
- Select the controller.
- Click **Configure**.
- Click **Cancel HBA mode**.

Smart Array P440ar
Embedded Slot

Actions

Cancel HBA Mode

Cancels HBA mode. The controller will remain in Smart Array mode. No reboot is required for changes to take effect.

Modify Power Mode

Modifies the controller's power mode. A reboot may be required after changing power modes to optimize power savings and performance.

Manage Device Identification LEDs

Turn the physical drive identification LED(s) On or Off

Status Messages

0 8 0 0

[View all status messages](#)

Controller Configuration Summary

0 Data Array(s)

7 Unassigned Drive(s)

[View more details](#)

5. A prompt appears indicating that the controller will remain in Smart Array mode and no reboot is required. If you want to continue, click **OK**.
6. A summary page appears. Click **Finish** to exit.

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, storage topology, and so forth. 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 very specific applications, but it is not appropriate for the majority of customers. Most applications will suffer significant performance reduction.



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

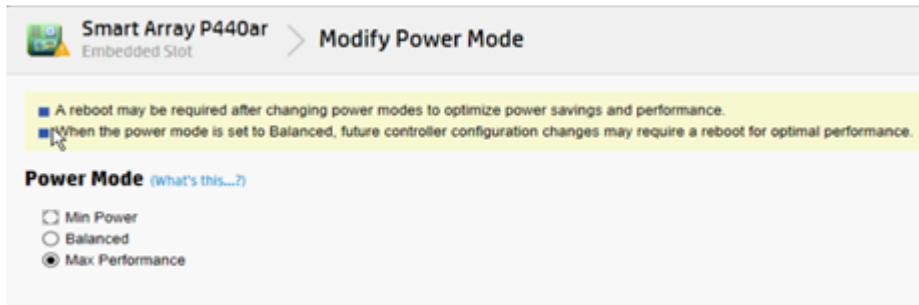


IMPORTANT: When the power mode is set to Balanced, future controller configuration changes may require a reboot for optimal performance.

To modify the power mode:

1. Open HPE SSA. For more information, see "Using the HPE SSA GUI (on page 16)."
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

To view the controller status, including the status of the cache and an attached battery or capacitor, do the following:

1. Open HPE SSA. For more information, see "Using the HPE SSA GUI (on page 16)."
2. Select a controller.
3. Click **Configure**.
4. Under **Controller Configuration Summary**, click **View more details**.
A new window appears.
5. Scroll down to **Controller Status** to view the status of the controller, cache, and battery or capacitor.

Controller Status	
Controller Status	OK
Cache Status	OK
Battery/Capacitor Status	OK

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** – (Gen 8 controllers and HDD only) Sanitize overwrite fills every physical sector of the drive with a pattern. If you select this option, you can specify a sanitize method. If the method is supported or enabled for your HDD, you can specify whether you want the sanitize operation to be restricted (drive is unavailable until successful completion) or unrestricted (drive is recoverable if operation fails). See "Sanitize erase methods (on page 54)" for more information.
- **Sanitize Block Erase** – (Gen 8 controllers and SSD only) Sanitize block erase sets the blocks on the drive to a vendor-specific value, removing all user data. If you select this option, you can specify a sanitize method. If the method is supported or enabled for your SSD, you can specify whether you want the sanitize operation to be restricted (drive is unavailable until successful completion) or unrestricted (drive is recoverable if operation fails). See "Sanitize erase methods (on page 54)" for more information.

- Sanitize Crypto Scramble – (Gen 8 controllers only) Sanitize cryptographic scramble (crypto scramble) changes the internal encryption keys that are used for user data, making the data irretrievable. If you select this option, you can specify a sanitize method. If the method is supported or enabled for your drive, you can specify whether you want the sanitize operation to be restricted (drive is unavailable until successful completion) or unrestricted (drive is recoverable if operation fails). See "Sanitize erase methods (on page 54)" 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 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

1. Open HPE SSA. For more information, see "Using the HPE SSA GUI (on page 16)."
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.

If you are performing a sanitize erase operation (supported on Gen8 controllers only), and the server is rebooted or powered off and back on during the sanitize operation, the drive may not be visible in the HPE SSA user interface until the sanitize operation is complete.

Making the erased drive available for configuration

After completing a drive erase using the one-, two-, and three-pass erase patterns, you must enable the erased drive before it is available for configuration. This step applies to Gen8 controllers.

If you complete a sanitize drive erase on a Gen8 controller, this step is not needed as the drive is automatically enabled following the completion of the erase process.

1. Open HPE SSA.

For more information, see "Using the HPE SSA GUI (on page 16)."
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.

From the list of unassigned drives, select the erased drive and click the drive or drives that you want to erase, and click the **Enable Erased Drive** button at the bottom of the screen.

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 30ms or 10ms. 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 the following controllers:

- Smart Array P244br Controller
- Smart Array P440 Controller
- Smart Array P440ar Controller
- Smart Array P441 Controller
- Smart Array P741m Controller
- Smart Array P840 Controller
- Smart Array P841 Controller
- H240 Smart Host Bus Adapter
- H241 Smart Host Bus Adapter

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.

To enable FLS:

1. Open HPE SSA. For more information, see "Using the HPE SSA GUI (on page 16)."
2. Select the controller.


3. Click **Advanced Controller Settings**.



The screenshot shows a sidebar menu titled "Actions" with five options, each with a descriptive paragraph below it:

- Create Array**: Creates an array from a group of selected physical drives. Once an array is created, the unused space is available for creating logical drives. A logical drive cannot be created unless an array exists.
- Create Arrays with RAID 0**: Create arrays from a group of selected physical drives. Each array contains exactly one physical drive and one RAID 0 logical drive.
- Modify Controller Settings**: Configures the supported controller settings. Depending on the controller, these can include setting the cache ratio, transformation and rebuild priorities, and surface scan delay.
- Advanced Controller Settings**: Configures the supported advanced controller settings. These settings can help improve the controller's performance for Video-On-Demand applications by changing the elevator sort parameters for example.
- Modify Spare Activation Mode**: Modifies the spare activation mode from the default behavior (activate on failure only) to predictive spare activation and back.

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



The screenshot shows the "HDD Flexible Latency Optimization" section with a link "(What's this...?)". Below the title are six radio button options:

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

5. Select one of the following:
 - o Disabled
 - o Low
 - o Middle (100ms)
 - o High
 - o 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:

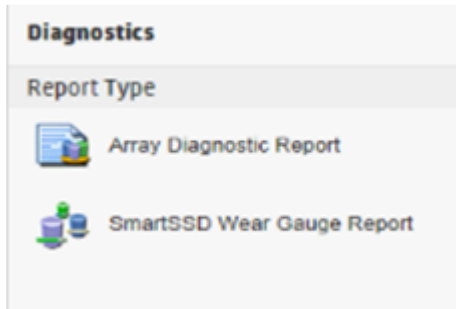
- File system types, contents, or status
- Partition types, sizes, or layout
- Software RAID information
- Operating system device names or mount points

Performing a Diagnostics task

1. Open HPE SSA.

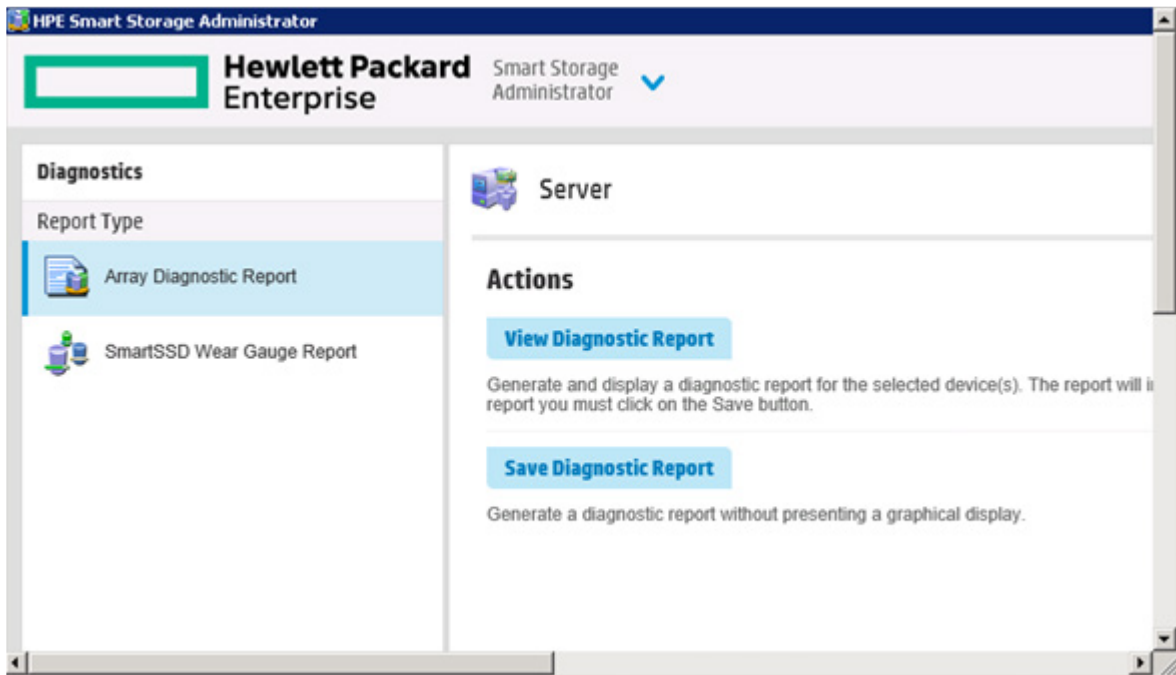
For more information, see "Using the HPE SSA GUI (on page 16)."

2. Open the **Diagnostics** panel by doing one of the following:
 - Choose 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. Choose a report type.



For this example, use the Array Diagnostic Reports selection.

4. Select **Array Diagnostic Report**.
The **Actions** panel for Array Diagnostic Report appears.



5. Click one of the task buttons:
 - If you select **View Diagnostic Report**, the report appears. When you are finished viewing the current report, click **Close** or **Save**.
 - If you select **Save Diagnostic Report**, wait for the report to generate, and then click **Close Report** or **Save Report**.

For more information about reports and viewing them in a browser, see the following sections:

- "Reported information (on page 121)"
- "Identifying and viewing diagnostic report files (on page 123)"
- "Identifying and viewing SmartSSD Wear Gauge report files (on page 124)"

Using the HPE SSA CLI

Access HPE SSA with one of the many methods available:

- Accessing HPE SSA in the offline environment (on page 10)
- 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 ("[Opening the CLI in Command mode](#)" on page 60)
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/hp/hpssa/README.TXT
- HPSSACLI (32bit): /opt/hp/hpssaccli/bld/hpssaccli-1.XX-X.X.XXXX.txt
- HPSSACLI (64bit): /opt/hp/hpssaccli/bld/hpssaccli-1.XX-XX.XXXX_XX.txt

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\hp\hpssaccli\Bin\hpssaccli.exe`
Alternatively, click **Start**, and select **Programs>HP System Tools>HP Smart Storage Administrator CLI>HP Smart Storage Administrator CLI**.
- For Linux, enter the following text:
`# hpssaccli`

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 ("[The <command> variable](#)" on page 61). 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\hp\Hpssacli\Bin\hpssacli.exe help`
- Using Linux:
`[root@localhost root]# hpssacli 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 ("Generating a diagnostic report" on page 67)  
help ("The help command" on page 65)  
rescan ("Rescanning the system" on page 69)  
shorthand ("Keyword abbreviations" on page 62)  
show ("The show command" on page 63)  
version ("Displaying the current versions of the application layers" on page 65)
```

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 ("[Querying a device](#)" on page 62). 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 ("[The help command](#)" on page 65).

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
cacheratio	cr	physicaldrive	pd
chassisname*	ch*	postprompttimeout	ppto
chassisserialnumber	csn	queuedepth	qd
chassisslot	chs	raid1writebuffering	r1wb
configurationmode	cm	rebuildpriority	rp

Keyword	Abbreviation in HPE SSA CLI	Keyword	Abbreviation in HPE SSA CLI
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	ee	surfacescanpriority	sp
expandpriority	ep	tapedrive	td
inconsistencyrepairpolicy	irp	waitforcacheroom	wfcr
licensekey	lk	—	—
logicaldrive	ld	—	—
mnppdelay	mnppd	—	—

*The CLI also uses this keyword and abbreviation for the terms **box name** and **RAID array ID**.

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 ("[The <command> variable](#)" on page 61)
- An HPE SSA CLI keyword or keyword abbreviation ("[Keyword abbreviations](#)" on page 62)
- 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 is only available in the offline environment and only 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 will prevent an error when booting after replacing the boot volume or controller. To assign a designated boot volume, see "Setting the boot volume (on page 66)."

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 only available in the offline 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 (on page 65)."

Scenario: When the bootable volume is assigned Logical Drive 2 and a non-bootable 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 (on page [121](#)).

For VMware ESX 5.x, diagnostic reports are generated differently and require you to run another utility, HPSSADUESXI, 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

For Gen8 controllers

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
- `unrestricted` applies to `crypto`, `block`, or `overwrite` options only

To stop the drive erase process:

Note: `Crypto`, `block`, or `overwrite` drive erase processes cannot be stopped.

```
<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.

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
```

For G6, G7, and Gen9 controllers

Syntax:

```
<target> modify [enableerasedrive]
```

where the target can be any valid physical drive

To begin the drive erase process:

```
<target> modify [erasepattern=zero|random_zero|random_random_zero]
```

where the target can be any valid physical drive

To stop the drive erase process:

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

where the target can be any valid physical drive

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
```

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.

To use this feature on a Smart Array G6 or G7 controller, a registered SAAP license key is required.

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.
size	# min max maxmbr ?	This parameter determines the preferred size of the logical drive. Units are MB.* The default setting is <code>max</code> .

Parameter	Acceptable values	Comments
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 parallel scsi ss_sas ss_sata ?	—

*Use only these units. Do not enter any extra text in the command to specify the units.

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 now enables you to move a single logical drive from one array to another array. This feature is available on Smart Array Gen8 and later controllers.

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 | satalogical | 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

```

SmartCache in HPSSACLI

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
```

```
ssdsmartpath=enable
```

```
controller slot=0 array a modify ssdsmartpath=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 HPSSACLI

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 that 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. Controllers that do not support a battery-backed write cache do not support this process.

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 (on page 128)."

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 do not support a battery-backed write cache 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 (on page 135)."

Consider the following factors before performing a migration:

- For some RAID-level migrations to be possible, you might need to add one or more drives to the array.
- 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 do not support a battery-backed write cache 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 change the ratio only if the controller has a battery-backed cache (because only battery-backed cache can be used for 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: Because physical drive write cache is not battery-backed, you could lose data if a power failure occurs during a write process. To minimize this possibility, use a backup power supply.

Syntax:

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

where <target> is a controller that supports drive write cache.

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]=#
[writeback|writethrough|?]
```

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>
```

Managing modes

HBA mode



IMPORTANT: Enabling HBA Mode exposes the physical drives to the operating system. Smart Array configuration is not allowed.

Syntax:

```
<target> modify [hbamode=on|off|?]
```

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

Example command:

```
controller slot=1 modify hbamode=on
ctrl slot=1 modify hbamode=off
```

RAID mode

Syntax:

```
<target> modify [raidmode=on|off|?]
```

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

Example command:

```
controller slot=1 modify raidmode=on
ctrl slot=1 modify raidmode=off
```

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

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 *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	=> ctrl slot=7 login user=crypto => ctrl slot=7 login user=crypto password=XXXXXX => ctrl slot=7 login user=user
To log out	=> ctrl slot=7 logout

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)	=> ctrl slot=7 enableencryption encryption=on eula=yes masterkey=XXXXXX localkeymanagermode=on mixedvolumes=on password=XXXXXX
Enable encryption (after initial setup)	=> ctrl slot=7 modify encryption=on eula=yes localkeymanagermode=on or =>ctrl slot=7 modify encryption=on eula=yes localkeymanagermode=off masterkey=XXXXXX
Disable encryption	=> ctrl slot=7 modify encryption=off

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

-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:

```
hpssacli 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	=> <code>controller slot=7 setctrlpasswd spassword=XXXXXX</code> => <code>controller slot=7 setctrlpasswd spassword=XXXXXX user=crypto password=XXXXXX</code>
Removing the controller password	=> <code>controller slot=7 removectrlpasswd</code> => <code>controller slot=7 removectrlpasswd user=crypto password=XXXXXX</code>

Setting	Command
Suspending the controller password	=> controller slot=7 modify ctrlpasswd=suspend
Resuming the controller password	=> controller slot=7 modify ctrlpasswdstate=resume => controller slot=7 modify ctrlpasswdstate=resume user=crypto password=XXXXXX

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

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
```

Using HPE SSA scripting

Access HPE SSA with one of the many methods available:

- Accessing HPE SSA in the offline environment (on page [10](#))
- Accessing HPE SSA in the online environment

The HPE SSA Scripting application has two scripting modes:

- Capture mode for capturing a configuration (on page [90](#))
HPE SSA 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 using an Input script (on page [90](#))
HPE SSA reads the array configuration described in a specified script file. See "Creating an HPE SSA script file (on page [90](#)).\" HPE SSA 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:

```
hpssascripting -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 names the default capture file `ssaoutput.ini` and places it in the HPE SSA working directory.

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

The `-e` switch information is used only if HPE SSA must generate an error file. The error filename is optional. By default, HPE SSA names the error file `error.ini` and places it in the HPE SSA working directory.

Using an Input script

To use an Input script to configure or reconfigure a system, first locate a suitable HPE SSA script or see "Creating an HPE SSA script file (on page 90)."

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

```
hpssascripting -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 names the error file `error.ini` and places it in the HPE SSA working directory.

Creating an HPE SSA script file

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

- Modify the sample custom input script (on page 91).
- Create a Capture file for capturing a configuration (on page 90).

You can create a capture file from any server that has HPE SSA 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 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 (on page 91).

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 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.
- HPE SSA Scripting does not distinguish between upper- and lowercase, except for the "CHASSISNAME", "CONNECTIONNAME" and "HOSTMODE" options.
- 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: 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 script files are divided into the following categories:

- Control category (on page 96)
- Controller category (on page 96)
- Array category (on page 101)
- Logical Drive category ("[HPE SSA scripting error messages](#)" on page 115)

Each category has several scripting options, but you do not always need to assign values to every option. HPE SSA 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 Method	These options define the overall behavior of HPE SSA when it processes scripts and creates configurations. Control options can occur only once in a script file and must be listed first.
Controller	Controller CacheState ClearConfigurationWithDataLoss DeleteLicenseKey DPOEnable DriveWriteCache ElevatorSortEnable ExpandPriority HBAMode IRPEnable Latency LicenseKey MNPDelay NoBatteryWriteCache ParallelSurfaceScanCount PowerMode PreferredPathMode QueueDepth RaidArrayId ReadCache RebuildPriority SurfaceScanDelay SurfaceScanDelayExtended SurfaceScanMode WriteCache	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. 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.

Category	Options	Description
Array	Array ArrayR0 Drive DriveType Join OnlineSpare Split	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 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.
Logical Drive	ArrayAccelerator LogicalDrive NumberOfParityGroups PreferredPath RAID Renumbr Repeat ResourceVolumeOwner Sectors ShrinkSize Size SizeBlocks StripSize StripeSize	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 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.

Control category

The Control category has the following options:

- Action mode (on page 96)
- Method mode (on page 96)

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 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 (on page 97)
- CacheState (on page 97)
- ClearConfigurationWithDataLoss (on page 97)

- DeleteLicenseKey ("LicenseKey, DeleteLicenseKey" on page 98)
- DPOEnable ("Video performance options" on page 100)
- DriveWriteCache (on page 98)
- ElevatorSortEnable ("Video performance options" on page 100)
- FLS (on page 98)
- HBAMode (on page 98)
- IRPEnable ("Video performance options" on page 100)
- LicenseKey ("LicenseKey, DeleteLicenseKey" on page 98)
- MNPDelay ("Video performance options" on page 100)
- NoBatteryWriteCache (on page 99)
- Power modes (on page 99)
- PreferredPathMode (on page 99)
- QueueDepth ("Video performance options" on page 100)
- RapidParityInitialization (on page 99)
- ReadCache ("ReadCache, WriteCache" on page 100)
- RebuildPriority ("RebuildPriority, ExpandPriority" on page 100)
- SurfaceScanDelay (on page 100)
- SurfaceScanDelayExtended (on page 100)
- Video performance options (on page 100)

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.

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]`

HBAmode



IMPORTANT: Enabling HBA Mode exposes the physical drives to the operating system. Smart Array configuration is not allowed.

This option enables HBA mode. Command options are `Enable` or `Disable`.

Example: `HBAmode=[Enable | 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 battery is present or when the battery fails. Values are `Enable` or `Disable`. The default is `Disable`.

Some controllers do not support this option.

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]`

PreferredPathMode

The setting that you select for this option determines how the preferred I/O path to a particular logical drive is set for a redundant array controller that is in an active/active configuration.

Not all controllers support this feature, and controllers in an active/standby configuration disregard this option.

- **Auto** is the default setting for new configurations. In this case, the storage system automatically selects the I/O path from the redundant controller to the logical drive and dynamically load balances all paths.
- **Manual** enables you to assign the logical drive to a specific redundant controller. If you select this setting, use the `PreferredPath` (on page 105) command to specify the path.

If you are reconfiguring a controller and do not specify a setting for this option, the existing setting remains unchanged.

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

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 use this feature on a Smart Array G6 or G7 controller, a registered SAAP license key is required.

To optimize the controller performance for video, set values for the following options as indicated:

```
DPOEnable = No
ElevatorSortEnable = Yes
IRPEnable = 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`

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.

N - number of scans

Array category

The Array category has the following options:

- Array (on page 101)
- ArrayR0 (on page 101)
- CachingArray (on page 101)
- Drive (on page 102)
- DriveType (on page 102)
- OnlineSpareMode (on page 102)
- OnlineSpare (on page 103)
- SplitMirror (on page 103)

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 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 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 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` (on page 97) 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 uses the first three available drives to build or expand the array that you define in the remainder of the script. HPE SSA 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 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:

- RAID1
- RAID10
- 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` (on page 104)
- `LogicalDrive` (on page 104)
- `CachingLogicalDrive` (on page 104)
- `CachedLogicalDrive` (on page 105)
- `NumberOfParityGroups` (on page 105)
- `PreferredPath` (on page 105)
- `RAID` (on page 105)
- `Renumber` (on page 106)
- `Repeat` (on page 106)
- `SetBootVolumePrimary` (on page 106)
- `SetBootVolumeSecondary` (on page 106)
- `Sectors` (on page 106)
- `Size` (on page 106)
- `SizeBlocks` (on page 107)
- `SmartCache Write-back` (on page 107)
- `SSDOPO` (on page 107)
- `StripSize` (on page 107)

ArrayAccelerator

This option specifies whether the array accelerator is enabled or disabled for the specified logical drive. The default value is Enabled.

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 accepts only the ID number of the next possible logical drive.
- In Reconfigure action mode, HPE SSA also accepts the ID number of any existing logical drive.

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 accepts only the ID number of the next possible logical drive.
- In Reconfigure mode, HPE SSA 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 accepts the ID number of any existing logical drive.

Caching logical drives and cached data drives must be on the same Array Controller.

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).

PreferredPath

If you select the Manual setting for `PreferredPathMode` (on page 99), use the `PreferredPath` command to specify the path for I/O to the logical drive on a redundant controller in active/active mode.

The default setting for this option is **1**. With this setting, the controller in chassis slot 1 is the preferred controller for I/O to the logical drive. If you select **2**, the controller in chassis slot 2 becomes the preferred controller for the logical drive.

To determine the chassis slot numbers, use the `show` command on a controller that supports redundant controllers.

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 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 ("`NumberOfParityGroups`" on page 105).
- 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 either ignores the new setting, or attempts to migrate the logical drive to the specified RAID level (when Method mode is Custom).

HPE SSA 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 1 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

When using a Smart Array G6 or G7 controller, some RAID levels require SAAP. RAID levels supported in a Caching Local Volume are limited by the firmware.

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 is to repeat this logical drive configuration.

Use one of the following values:

- `N`—In Configure mode, HPE SSA creates *N* new logical drives.
- `MAX`—HPE SSA 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 63 if the logical drive is larger than 502 GB. Otherwise, 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 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-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

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, 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 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 (supported starting with Gen8 controllers)

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 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: 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 ("[EncryptionCryptoPasswordSet](#)" on page 109)
- Key management mode ("[EncryptionKeyManager](#)" on page 109)
- Initial Master Encryption Key ("[EncryptionMasterKey](#)" on page 110)
- EULA ("[AcceptEULA](#)" on page 108) (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` (on page 110). 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 ("[EncryptionRecoveryQuestion](#)" on page 110). 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

The following commands have been added to HPE SSA Scripting versions 1.60.0.0 and later to support data encryption features introduced with select Gen8+ controllers.

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.

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:\hpssascripting -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 (on page 111). 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>
	<Action>Configure</Action>
Controller= Slot 1	<Method>Custom</Method>
Array=A	<Controller ID="Slot 1">
Drive= 1I:4:8, 1I:4:7	<Array ID="A">
	<Drive>1I:4:8, 1I:4:7</Drive>
LogicalDrive= 1	<LogicalDrive ID="1">
RAID= 0	<Raid>0</Raid>
Size= 100000	<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.
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 a battery attached	The operation requested requires a battery and one was not attached.
4008	The controller battery status is Failed	The operation requires a battery and the attached battery 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,

Error code	Error message	Comment or clarification
		combining <code>-reset</code> with capture mode.
2818	Invalid Method	The method must be either Custom or Auto.
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.
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

Error code	Error message	Comment or clarification
		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	—
2868	Failure opening error file	—
2877	There are no suitable spares available	HPE SSA 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.

Error code	Error message	Comment or clarification
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 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.
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 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 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.	—

Error code	Error message	Comment or clarification
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.
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	—
3049	Encryption Login Failed	—
3050	Encryption Login command is not available	—
3051	The Encryption command is not available	—

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.
3053	The Firmware Lock/Unlock failed	The controller firmware returned an error when attempting to lock or unlock the firmware.
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.
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.
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.
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.
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	—
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.
3100	SSD Smart Path is not supported	—
3101	The SSD Smart Path command failed	—
3102	Unable to lock controller. Make sure other instances of HPSSA, HPSSACLI, or HPSSASCRIPING are not configuring this controller	—

Troubleshooting

HPE Smart Storage Administrator Diagnostics Utility CLI

About the utility

Formerly known as the Array Diagnostics 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 (on page [121](#))."

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.

Field	Description
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 not displayed when the SSD Utilization is still at 0%.
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

Installing the utility

1. Browse to the HPE Smart Storage Administrator website (<http://www.hpe.com/servers/ssa>).
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\HP System Tools\.

Launching the utility in CLI mode

1. Click **Start>All Programs>HP System Tools>HP Smart Storage Administrator Diagnostics Utility>Read Me**
2. Open a command prompt.
3. Change directory (cd) to the location where `hpssaducli.exe` is installed.
This directory is commonly C:\Program Files\hp\hpssaducli\bin.
4. Do one of the following:
 - o Generate a diagnostic report with the following command:
`hpssaducli -f adu-report.zip`
 - o Generate a SmartSSD Wear Gauge report with the following command:
`hpssaducli -ssd -f ssd-report.zip`

For more options, use the following command:

```
hpssaducli -help
```

Diagnostic report procedures

Viewing the diagnostic report

1. Launch the utility ("[Launching the utility in CLI mode](#)" on page 123).
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.

To view the diagnostic report files in a browser:

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

1. Launch the utility ("[Launching the utility in CLI mode](#)" on page [123](#)).
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

To view the SmartSSD Wear Gauge report files in a browser:

1. Extract the following files to a directory:
 - `SmartSSDWearGaugeReport.json`
 - `SmartSSDWearGaugeReport.htm`
 All files must reside in the same directory.
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 non-optimal logical drive alignment for 512e physical drives.

The following scenarios indicate 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.

As a response, HPE SSA will display a warning indicating the logical drive is not optimally aligned and that performance of the logical drive will not be optimal. Additionally, the array will present a "Re-align Logical Drive" button if the following scenario is met:


- There is enough free space in the array to move the logical drive to be aligned to the native 4K boundary.
- The controller is capable of performing the transformation (requires a cache module with a fully-charged battery or capacitor connected).
- 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 of 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** ("[Fault-tolerance methods](#)" on page 128).

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.

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 and RAID 10 configurations tolerate multiple drive failures if no failed drives are mirrored to one another.
- 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. Those most often used with Smart Array controllers are hardware-based RAID methods.

Two alternative fault-tolerance methods that are sometimes used are also described ("[Alternative fault-tolerance methods](#)" on page 135). However, hardware-based RAID methods provide a much more robust and controlled fault-tolerance environment, so these alternative methods are seldom used.

RAID 0—No fault tolerance

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.



Advantages:

- 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 is needed for fault tolerance).


Disadvantages:

- All data on the logical drive is lost if a physical drive fails.
- Cannot use an online spare.
- Can only preserve data by backing it up to external drives.


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.

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.



In each mirrored pair, the physical drive that is not busy answering other requests answers any read requests that are sent to the array. This behavior is called **load balancing**. 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.

This fault-tolerance method is useful when high performance and data protection are more important than the cost of physical drives.

Advantages:

- This method has the second highest read 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.


Disadvantages:

- This method is expensive, because many drives are needed for fault tolerance.
- Only half of the total drive capacity is usable for data storage.


RAID 1 (ADM) and RAID 10 (ADM)

In RAID 1 (ADM) and RAID 10 (ADM) configurations, data is duplicated to two additional drives.

When the array contains only three physical drives, the fault-tolerance method is known as RAID 1 (ADM).



When the array has more than three physical drives, drives are mirrored in trios, and the fault-tolerance method is known as RAID 10 (ADM).



In each mirrored trio, the physical drives that are not busy answering other requests answer any read requests that are sent to the array. This behavior is called **load balancing**. 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.

This fault-tolerance method is useful when high performance and data protection are more important than the cost of physical drives.

Advantages:


- This method has the highest read performance of any fault-tolerant 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.

Disadvantages:

- This method is expensive, because many drives are needed for fault tolerance.
- Only one-third of the total drive capacity is usable for data storage.

RAID 5—distributed data guarding

In a RAID 5 configuration, data protection is provided by **parity data** (denoted by $P_{x,y}$ in the figure). This parity data is calculated stripe by stripe from the user data that is written to all other blocks within that stripe. The blocks 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 calculated from the remaining parity data and user data on the other drives in the array. This recovered data is usually written to an online spare in a process called a **rebuild**.

This configuration is useful when cost, performance, and data availability are equally important.

Advantages:

- Has high read performance.
- Data is not lost if one physical drive fails.
- More drive capacity is usable than with RAID 1+0—parity information requires only the storage space equivalent to one physical drive.


Disadvantages:

- Has relatively low write performance.
- Data is lost if a second drive fails before data from the first failed drive is rebuilt.

RAID 6 (ADG)—Advanced Data Guarding

NOTE: Not all controllers support RAID 6 (ADG).

RAID 6 (ADG), like RAID 5, generates and stores parity information to protect against data loss caused by drive failure. With RAID 6 (ADG), however, 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.



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 (ADG) is less than it would be if it was configured with RAID 5.

Advantages:


- This method has a high read performance.
- This method allows high data availability—Any two drives can fail without loss of critical data.
- More drive capacity is usable than with RAID 1+0—Parity information requires only the storage space equivalent to two physical drives.

Disadvantages:

The main disadvantage of RAID 6 (ADG) is a relatively low write performance (lower than RAID 5) because of the need for two sets of parity data.

RAID 50

RAID 50 is a nested RAID method in which the constituent hard 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 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, 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.

RAID 50 is particularly useful for large databases, file servers, and application servers.

Advantages:


- 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.

Disadvantages:

- 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 60

RAID 60 is a nested RAID method in which the constituent hard 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.

RAID 60 is particularly useful for data archives and high-availability solutions.

Advantages:

- Higher performance than for RAID 6, especially during writes.
- Better fault tolerance than either RAID 0 or RAID 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.

Disadvantages:

- 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.

Comparing the hardware-based RAID methods

Not all controllers support all RAID levels. To support some RAID levels, Smart Array G6 and G7 controllers require SAAP.

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

Item	RAID 0	RAID 1+0	RAID 5	RAID 6 (ADG)	RAID 1(0) (ADM)
Formula for number of drives usable for data (n = total number of drives in array)	n	$n/2$	$n-1$	$n-2$	$n/3$
Percentage of drive space usable*	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**
Read performance	High	High	High	High	High
Write performance	High	Medium	Low	Low	Medium
Relative cost	Low	High	Medium	Medium	Very high

*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).

**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 (<http://www.hpe.com/servers/smartarray>).

Most important criterion	Also important	Suggested RAID level
Fault tolerance	Cost effectiveness I/O performance	RAID 6 RAID 10 (ADM), RAID 1+0, RAID 50, RAID 60
Cost effectiveness	Fault tolerance I/O performance	RAID 6 RAID 5 (RAID 0 if fault tolerance is not required)
I/O performance	Cost effectiveness Fault tolerance	RAID 5 (RAID 0 if fault tolerance is not required) 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 (<http://www.hpe.com/support/hpesc>). 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 (<http://www.hpe.com/info/InsightDiagnostics>).
- **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 more recent products, array diagnostics is available with HPE SSA v9.0 and later. This utility is available through Intelligent Provisioning on ProLiant Gen8 and later servers, and also on the Hewlett Packard Enterprise website (<http://www.hpe.com/support/hpesc>). For more information about HPE SSA, see "About SSA ("About HPE SSA" on page 5)." For more 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 ("HPE Smart Storage Administrator Diagnostics Utility CLI" on page 121)."

Troubleshooting resources

ProLiant G7 (and earlier) products

The *HPE ProLiant Servers Troubleshooting Guide* provides procedures for resolving common problems and comprehensive courses of action for fault isolation and identification, error message interpretation, issue resolution, and software maintenance on ProLiant servers and server blades. This guide includes problem-specific flowcharts to help you navigate complex troubleshooting processes. To view the guide, select a language:

- English (http://www.hpe.com/support/ProLiant_TSG_en)
- French (http://www.hpe.com/support/ProLiant_TSG_fr)
- Italian (http://www.hpe.com/support/ProLiant_TSG_it)
- Spanish (http://www.hpe.com/support/ProLiant_TSG_sp)
- German (http://www.hpe.com/support/ProLiant_TSG_gr)
- Dutch (http://www.hpe.com/support/ProLiant_TSG_nl)
- Japanese (http://www.hpe.com/support/ProLiant_TSG_jp)

- The *HPE BladeSystem c-Class Enclosure Troubleshooting Guide* provides procedures and solutions for troubleshooting BladeSystem c-Class enclosures. This guide explains how to use the Insight Display to troubleshoot enclosures, and it includes a flowchart to help you navigate the troubleshooting process. To view the guide, see the Hewlett Packard Enterprise website ([http://www.hpe.com/support/BladeSystem Enclosure TSG en](http://www.hpe.com/support/BladeSystem%20Enclosure%20TSG%20en)).

ProLiant Gen8 products

The *HP ProLiant Gen8 Troubleshooting Guide, Volume I: Troubleshooting* provides procedures for resolving common problems and comprehensive courses of action for fault isolation and identification, issue resolution, and software maintenance on ProLiant servers and server blades. To view the guide, select a language:

- English ([http://www.hpe.com/support/ProLiant TSG v1 en](http://www.hpe.com/support/ProLiant%20TSG%20v1%20en))
- French ([http://www.hpe.com/support/ProLiant TSG v1 fr](http://www.hpe.com/support/ProLiant%20TSG%20v1%20fr))
- Spanish ([http://www.hpe.com/support/ProLiant TSG v1 sp](http://www.hpe.com/support/ProLiant%20TSG%20v1%20sp))
- German ([http://www.hpe.com/support/ProLiant TSG v1 gr](http://www.hpe.com/support/ProLiant%20TSG%20v1%20gr))
- Japanese ([http://www.hpe.com/support/ProLiant TSG v1 jp](http://www.hpe.com/support/ProLiant%20TSG%20v1%20jp))
- Simplified Chinese ([http://www.hpe.com/support/ProLiant TSG v1 sc](http://www.hpe.com/support/ProLiant%20TSG%20v1%20sc))

The *HP ProLiant Gen8 Troubleshooting Guide, Volume II: Error Messages* provides a list of error messages and information to assist with interpreting and resolving error messages on ProLiant servers and server blades. To view the guide, select a language:

- English ([http://www.hpe.com/support/ProLiant EMG v1 en](http://www.hpe.com/support/ProLiant%20EMG%20v1%20en))
- French ([http://www.hpe.com/support/ProLiant EMG v1 fr](http://www.hpe.com/support/ProLiant%20EMG%20v1%20fr))
- Spanish ([http://www.hpe.com/support/ProLiant EMG v1 sp](http://www.hpe.com/support/ProLiant%20EMG%20v1%20sp))
- German ([http://www.hpe.com/support/ProLiant EMG v1 gr](http://www.hpe.com/support/ProLiant%20EMG%20v1%20gr))
- Japanese ([http://www.hpe.com/support/ProLiant EMG v1 jp](http://www.hpe.com/support/ProLiant%20EMG%20v1%20jp))
- Simplified Chinese ([http://www.hpe.com/support/ProLiant EMG v1 jp](http://www.hpe.com/support/ProLiant%20EMG%20v1%20jp))

ProLiant Gen9 products

The *HPE ProLiant Gen9 Troubleshooting Guide, Volume II: Error Messages* provides a list of error messages and information to assist with interpreting and resolving error messages on ProLiant servers and server blades. To view the guide, select a language:

- English ([http://www.hpe.com/support/Gen9 EMG en](http://www.hpe.com/support/Gen9%20EMG%20en))
- French ([http://www.hpe.com/support/Gen9 EMG fr](http://www.hpe.com/support/Gen9%20EMG%20fr))
- Spanish ([http://www.hpe.com/support/Gen9 EMG es](http://www.hpe.com/support/Gen9%20EMG%20es))
- German ([http://www.hpe.com/support/Gen9 EMG de](http://www.hpe.com/support/Gen9%20EMG%20de))
- Japanese ([http://www.hpe.com/support/Gen9 EMG ja](http://www.hpe.com/support/Gen9%20EMG%20ja))
- Simplified Chinese ([http://www.hpe.com/support/Gen9 EMG zh cn](http://www.hpe.com/support/Gen9%20EMG%20zh%20cn))

Optional components

HPE Secure Encryption

HPE Secure Encryption is a controller-based, enterprise-class data encryption solution that protects data at rest on any bulk storage attached to the HPE Smart Array controller. The solution is available for both local and remote deployments, and works in conjunction with HPE SmartCache.

HPE 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 has completed FIPS-140-2 (<http://csrc.nist.gov/groups/STM/cmvp/documents/fips140-2/FIPS1402IG.pdf>) validation. 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 HPE 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 the *HPE Secure Encryption Installation and User Guide*.

To purchase a license key, see the HPE Secure Encryption product page on the Hewlett Packard Enterprise website (<http://www.hpe.com/info/enterprise/docs>).

HPE Smart Array Advanced Pack

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 (<http://www.hpe.com/info/enterprise/docs>).

A SAAP license for Smart Array G6 or G7 controllers is no longer available for purchase as it has been discontinued. For Smart Array Gen8 controllers, update the controller firmware to take advantage of the advanced features of SAAP without the need of license. For more information about advanced features, see the SAAP QuickSpecs on the Hewlett Packard Enterprise website (<http://www.hpe.com/info/qs>).

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 ("[LicenseKey](#), [DeleteLicenseKey](#)" on page 98)

Some SAAP features are required for advanced configuration tasks in HPE SSA. For a list of these tasks, see "Support for configuration tasks (on page 6)."

Acronyms and abbreviations

ACU

Array Configuration Utility

ADG

Advanced Data Guarding (also known as RAID 6)

ADM

Advanced Data Mirroring

DHCP

Dynamic Host Configuration Protocol

FLS

Flexible Latency Scheduler

HPE SSA

HPE Smart Storage Administrator

iLO

Integrated Lights-Out

ISO

International Organization for Standardization

POST

Power-On Self Test

PXE

preboot execution environment

RIS

reserve information sector

SAAP

Smart Array Advanced Pack

SAS

serial attached SCSI

SATA

serial ATA

SMART

self-monitoring analysis and reporting technology

SSD

solid-state drive

SSP

Selective Storage Presentation

TFTP

Trivial File Transfer Protocol

WBEM

Web-Based Enterprise Management

WWN

World Wide Name

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 ([**mailto:docsfeedback@hpe.com**](mailto: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.

Index

5

512e physical drive support 124

A

action mode, HPE SSA Scripting 96
ACU scripting 101
ADG (advanced data guarding) 132
advanced configuration tasks, support for 5, 6
advanced data guarding (ADG) 132
array accelerator, enabling or disabling 82
array accelerator, enabling or disabling, ACU CLI 82
array concepts 125
array configuration utilities, features 5
array configuration, copying 90
array, healing 49
array, moving 77, 102
array, replacing 49
array, shrinking 25, 77, 102
ArrayAccelerator 104

B

boot controller, setting 65
boot volume, setting 66, 106

C

Cache Manager 39
cache ratio, setting 81
CacheState 97
capture script 92
choosing a RAID level 135
ClearConfigurationWithDataLoss 97
clearing a configuration 89, 109
CLI (Command Line Interface) 60
CLI abbreviations 62
CLI syntax 61, 65
CLI, accessing 10, 123
Command Line Interface (CLI) 60
comparison of RAID features 134
comparison of the utilities 5
configuration tasks, advanced, support for 6
configuration tasks, performing 25, 65, 86, 90
configuration tasks, standard, support for 6
configuration utilities, comparison of 5
configuration utilities, description of 8
configuration utility, choosing 8
configuring a controller 96
control category 96

controller category, scripting 96
controller duplexing 135
controller password 86, 108
controller state 53
controller, specifying, HPE SSA scripting 97
criteria for selecting a RAID level 135
crypto officer 109, 110

D

data protection methods 128, 135
data striping 125, 128
default settings 91
DeleteLicenseKey 98
deleting a device 67
device information, obtaining 63
devices, discovering 18
devices, identifying 67
diagnostic report 58, 67, 123
diagnostic tools 136
diagnostics 21, 22, 58, 67
disabling drive cache 81
disabling the array accelerator 82, 104
distributed data guarding 131
DPOEnable 98, 100
drive array concepts 125
drive cache, enabling or disabling 81, 98
drive information 73
drive mirroring 128, 130
drive support 124
drive type, HPE SSA Scripting 102
DriveType 102
DriveWriteCache 98
duplexing 135

E

effects of a hard drive failure 127
ElevatorSortEnable 98, 100
enabling drive cache 81
enabling the array accelerator 82, 104
enclosure information 73
encryption 50, 93, 108, 138
encryption key manager 88, 109
encryption keys 87, 109, 110
encryption script 93, 94, 108, 111
encryption settings 85, 108, 109, 110
erase a drive 55
erase, sanitize 54
error messages 115
EULA (end user license agreement) 89, 108
example CLI script, creating logical drive 71

- exit on error 82
- expand priority, setting 80
- expanding an array 102
- expanding an array, HPE SSA Scripting 102
- ExpandPriority 100
- Express Local Encryption 85, 110
- extending logical drive capacity 106

F

- F5 prompt 10
- failed logical drive, re-enabling 81
- fault-tolerance methods 128
- features of the utilities 5
- firmware lock 88
- Flexible Latency Scheduler 98

G

- GUI, accessing 10

H

- HBA mode 50
- healing an array 49
- help command 65
- help resources 22
- HPE Secure Encryption 138
- HPE Smart Storage Administrator (HPE SSA) 10, 121
- HPE Smart Storage Administrator Diagnostics Utility CLI 60, 84, 121, 136
- HPE SmartCache 41, 44, 45, 74
- HPE SSA (HPE Smart Storage Administrator) 5

I

- I/O path, modifying 99
- identifying devices 67
- information about devices, obtaining 63
- Input script, using 90, 91
- installing the utility 121, 123
- instant secure erase 88, 111
- Intelligent Provisioning 10
- IRPEnable 98, 100
- ISO image 11

K

- keystroke combinations 17
- keyword abbreviations 62

L

- languages 5
- LEDs, activating 67
- LicenseKey 69, 98
- local application, using HPE SSA as 10, 14
- logical drive capacity extension 78, 106

- logical drive, caching 104
- logical drive, creating 70
- logical drive, description of 125
- logical drive, failed, re-enabling 81
- logical drive, migrating, HPE SSA Scripting 105, 107
- logical drive, moving, HPE SSA CLI 72
- LogicalDrive 70, 104

M

- MaxBoot setting 106
- method mode, HPE SSA Scripting 96
- methods for opening the HPE SSA GUI 16
- migrating stripe size or RAID level, HPE SSA Scripting 105, 107
- minimum requirements 8
- mirrored arrays 46
- mirrored drives 128, 130
- MNPDelay 99, 100
- mounting virtual media 11
- moving an array 77, 102

N

- navigating the interface 18
- nested RAID 133, 134
- NoBatteryWriteCache 99
- Novell NetWare, configuration tool for 8
- NumberOfParityGroups 105

O

- offline deployment 10
- online spare, HPE SSA Scripting 103
- OnlineSpare 102
- opening the HPE SSA GUI 16
- operating system support 9
- options, scripting, list of 95
- overview of configuration utilities 8

P

- parallel surface scan 37, 81, 101
- parameters, in CLI, obtaining values of 62
- parity groups 105, 133, 134
- password recovery, administrator 85
- physical drive write cache, enabling or disabling 81
- plaintext volume 87, 108
- POST error messages 121, 136
- power modes 52, 99
- PreferredPath 105
- PreferredPathMode 99
- PXE-based deployment 12
- PXELinux configuration 12

Q

- querying a device 62
- QueueDepth 99, 100

R

- RAID 34, 78
- RAID 0 volume, creating 101
- RAID configuration 34, 78
- RAID level migration 105
- RAID level selection criteria 135
- RAID levels 128, 130, 131, 133, 134
- RAID levels, comparison of features 134
- RAID mode 50, 84, 98
- RAID, software-based 135
- RAID-level migration, HPE SSA Scripting 105
- Rapid Parity Initialization 75, 99
- ReadCache 100
- read-write ratio, setting 81, 104
- RebuildPriority 79, 100
- recombining a split mirrored array 47
- redundant controller, modifying I/O path of 99, 105
- redundant controller, preferred path settings for 99, 105
- re-enabling a failed logical drive 81
- registering a license key 44
- re-mirroring a split mirror backup 48
- remote server, configuring 15
- Renumber 106
- Repeat 106
- repeat configuration 106
- replacing an array 49, 78
- rescan 18, 69
- running the HPE SSA GUI 16

S

- sample script 91, 92
- sanitize/erase 54
- screen description 18, 19, 21, 22
- scripting modes 89
- scripting options 95
- scripting syntax 90
- Sectors 106
- Secure Encryption 50, 84, 111
- selecting a RAID level 135
- shorthand in CLI 62
- show (CLI command) 63
- shrinking an array 25, 77
- Size 106
- SizeBlocks 107
- SmartCache 41, 44, 45
- SmartCache Writeback 45, 107
- SmartSSD Wear Gauge report 58, 124
- software-based RAID 135
- spare activation mode 31, 76
- spare drives 75
- spare drives, HPE SSA Scripting 103
- spare management 32, 76
- Split 82
- Split Mirror 82, 103
- split mirror backup 48

- split operation, performing 82
- splitting a mirrored array 47
- SSD OPO 75, 107
- standard configuration tasks, support for 6
- stripe size migration, HPE SSA Scripting 107
- striping data 125, 128
- StripSize 107
- supported operating systems 13
- supported tasks 6
- surface scan delay 80, 100
- surface scan mode 80
- SurfaceScanDelay 100
- SurfaceScanDelayExtended 100
- SurfaceScanMode 100
- syntax, CLI 61, 65

T

- target device, setting 66
- tasks 23, 25, 58
- tasks supported in each utility 6
- troubleshooting 121, 136

U

- unlocked volumes 89
- utilities, description of 8

V

- version command 65
- video performance optimization 69, 100
- volatile keys 87, 111

W

- warning messages 115
- WriteCache 100

X

- XML DTD 114
- XML output file 111
- XML support 111