Hewlett Packard Enterprise

# HPE Smart Storage Administrator GUI User Guide

**Abstract**

This document provides instructions for configuring and managing diagnostics for Smart Storage Administrator GUI users. The primary audience is the system administrator with a good working knowledge of storage hardware and the configuration of logical drives and arrays. Hewlett Packard Enterprise assumes you are qualified in the servicing of computer equipment and trained in recognizing hazards in products with hazardous energy levels.

Part Number: 30-9B6DAFAB-008
Published: August 2026
Edition: 8

# Introduction

Smart Storage Administrator (SSA) is a primary tool for configuring arrays on SmartRAID controllers. It has the following interface formats:

- SSA GUI
- SSA CLI
- SSA Scripting
- SSA Diagnostics Utility CLI (SSADUCLI)

All formats provide support for configuration tasks. Some of the advanced tasks are available in only one format.

SSA is accessible both offline and online:

- Offline environment: There are multiple methods to run SSA. In Offline mode, users can configure or maintain detected and supported SmartRAID (SR) controllers. Some SSA features are only available in the offline environment, such as setting the boot controller and boot volume.
- Online environment: This method requires an administrator to download the SSA executables and install them. You can run SSA online after launching the host Operating System (OS).

This guide provides detailed information about configuration tasks supported on SSA's graphical user interface (GUI).

# Benefits of Using SSA

SSA is an advanced utility that enables you to perform many complex configuration tasks, as well as the ability to run various diagnostic reports.
Using SSA over other configuration utilities provides the following benefits:

- Supports GUI, CLI, and Scripting interfaces
- Supports different languages: English, French, German, Italian, Japanese, Portuguese, Russian, Simplified Chinese, and Spanish
- Applications are executed using the following tools:
  - Any bootable media, such as a software CD
  - Most supported server platform host OS with a web browser
- Supports online and offline environments

# What You Need to Know Before You Begin

### Minimum Requirements

Minimum OS requirements to run any SSA format are:

- Microsoft Windows® 2016 or later
- Red Hat Enterprise Linux 7, 8, 9
- SUSE Linux Enterprise Server 12, 15
- Vmware 7, 8

To run the SSA GUI, you need a monitor with minimum resolution of 1024 x 768 and 16-bit color. The GUI supports the following browsers:

- Mozilla Firefox 9.0 or later
- Microsoft® Edge 88.0.705.50 or later
- Google Chrome ™

**Note:** Ensure the following browser options are enabled:

- Automatically load images and other data types.
  - JavaScript
  - Stylesheets

### Native Support for 64-bit OS

SSA offers a native 64-bit SSA application for supported 64-bit OS, eliminating the need for compatibility libraries. You can choose to install the application that matches the OS installed on the server product.

# Table of Contents


Introduction....2

Benefits of Using SSA....2

What You Need to Know Before You Begin....2

1. Array and Controller Configuration....6

2. Array Configuration Guidelines....7

3. Supported Configurations....8

4. Accessing SSA....9
   - 4.1. Accessing SSA in the Offline Environment....9
   - 4.2. Accessing SSA in the Online Environment....9

5. Familiarize with the SSA GUI....11
   - 5.1. Icons....11
   - 5.2. Keyboard Shortcuts....12
   - 5.3. Navigating the GUI....12

6. Configuration Tasks....17
   - 6.1. Performing a Configuration....18
   - 6.2. Creating an Array....20
   - 6.3. Configuring a Controller....20
   - 6.4. Enabling or Disabling SSD Smart Path....21
   - 6.5. Setting Sanitize Lock....21
   - 6.6. Rapid Parity Initialization (RPI)....22
   - 6.7. Changing the Spare Activation Mode....23
   - 6.8. Changing the Spare Management Mode....24
   - 6.9. Disabling Over Provisioning Optimization....25
   - 6.10. Setting Auto RAID 0....25
   - 6.11. Modifying Parallel Surface Scan....25
   - 6.12. Configuring Controller Cache....26
   - 6.13. SmartRAID SmartCache....26
   - 6.14. Working with Mirrored Arrays....29
   - 6.15. Modifying an Array....31
   - 6.16. Heal Array....41
   - 6.17. Change Array Drive Type....41
   - 6.18. Moving a Logical Drive....42
   - 6.19. Encryption Manager....43
   - 6.20. Power Modes....43
   - 6.21. Viewing Controller Status....45
   - 6.22. Options for Erasing Drives....45
   - 6.23. Erasing a Drive....46
   - 6.24. Managing Flexible Latency Scheduler (FLS)....47
   - 6.25. Configure the Language of the GUI....48
   - 6.26. Maintaining Controller Firmware....49

7. Managing Diagnostics....50
7.1. Performing a Diagnostics Task Using SSA....50
8. Data Encryption....52
8.1. About CBE....52
8.2. About SED....56
8.3. Planning....56
8.4. Configuration....58
8.5. Operations....70
8.6. Maintenance....96
9. Troubleshooting....100
9.1. SSA Diagnostics Utility CLI....100
9.2. 512e Physical Drive Support....103
9.3. Drive Arrays and Fault-Tolerance Methods....103
9.4. Diagnosing Array Problems....113
9.5. Secured Encryption Issues....113
10. Glossary....124
11. Resources....136
12. Revision History....137

# 1. Array and Controller Configuration

After the initial configuration, you can configure arrays and controllers during the initial provisioning of the server or compute module. Configuration tasks are initiated using SSA, which are accessible through Intelligent Provisioning.

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

- When you use SSA offline, you can specify options to enable polling for any drives that are present and build an appropriate array for those drives.
- You can use SSA offline for provisioning.
- You can use the Unified Extensible Firmware Interface (UEFI) System Utilities to create the primary array that is required.

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

# 2. Array Configuration Guidelines

Note 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 NVMe, SAS or all SATA and either all hard drives or all Solid State Drive (SSD)s).
- For the most efficient use of drive space, all drives within an array must approximately have 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 experiences 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.

# 3. Supported Configurations

SSA supports the following configuration tasks. Support for individual tasks varies according to the controller type. It is recommended to have the latest firmware installed to have access to all of the supported features.

- Assign a RAID level to a logical drive
- Assign spare drives to an array
- Configure multiple systems identically
- Copy the configuration of one system to multiple systems
- Create multiple logical drives per array
- Create or delete arrays and logical drives
- Enable or disable a physical drive write cache on physical drives that are configured as part of a logical volume
- Enable or disable a physical drive write cache on unconfigured physical drives
- Enable SSD to be used as caching devices, using SmartCache
- Enable optimized data path to SSDs using SSD SmartPath
- Expand an array
- Extend a logical drive
- Heal an array
- Drive Erase
- Identify devices by causing their LEDs to Flash
- Manage Encryption—Controller Based Encryption (CBE) and Managed SEDs (MSEDs)
- Migrate the RAID level or stripe size
- Move an array
- Move and delete individual Logical Unit Number (LUN)s
- Optimize the controller performance for video
- Re-enable a failed logical drive
- Set the boot controller
- Set the expand priority, migrate priority, and accelerator ratio
- Set the Spare Activation mode
- Set the stripe size
- Set the surface scan delay
- Share a spare drive among several arrays
- Remove a drive from an array
- 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 (Triple), and 10 (Triple) mirrors

# 4. Accessing SSA

This section describes accessing SSA in the off-line and on-line environment.

## 4.1. Accessing SSA in the Offline Environment

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

- Launching within HPE Intelligent Provisioning (IP)
- Launching within Service Pack for Proliant

## 4.2. Accessing SSA in the Online Environment

To access, install, and launch SSA in the online environment, you must download the SSA executables. All three formats have separate executables. SSA Scripting is a standalone application that is distributed with the SSA CLI application. Both SSA and the SSA CLI need sg driver (scsi generic) to access SmartRAID controller in Linux. Red Hat Enterprise Linux 7.1 and later does not load sg driver automatically. You must load the driver and type “modprobe sg” before running SSA or SSA-CLI. To use SSA in the online environment, obtain the executable files from the Hewlett Packard Enterprise website.

To access SSA in the online environment, perform the following steps:

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

### 4.2.1. Launching SSA on a Local Server

This section describes launching the SSA on Microsoft and Linux OS.

**Microsoft OS**

To launch SSA on a Microsoft OS, perform the following steps:

1. Click **Start,** and then select **Programs>Windows System>Smart Storage Administrator**.
   SSA launches in either a browser or application window (v1.50 and later). SSA then scans the system and detects controllers. When controller detection is complete, the controllers are available on the **Controller Devices** menu.
2. Configure a controller. For more information, see Configuring a Controller.
   When configuration is complete, continue with the next step.
3. (Optional) To make newly created logical drives available for data storage, use the OS disk management tools to create partitions and format the drives.

**Linux OS**

Perform the following to launch SSA in Local mode on a Linux OS:

In the command prompt, enter the following: `ssa -local`.

SSA launches in the browser window.

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

# 5. Familiarize with the SSA GUI

The following sections describe the SSA GUI.

## 5.1. Icons

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

**Table 5-1.** Icons

| Image | Description |
|---|---|
|  | Critical |
|  | Warning |
|  | Informational |
|  | Active Task(s) |
|  | Paused/Offline Drive |
|  | Server |
|  | Array Controller |
|  | Array Controller (Embedded) |
|  | Array/Logical Devices or Solid State Devices |
|  | Logical Drive |
|  | Assigned Physical Drive |
|  | Unassigned Physical Drive |
|  | Unassigned Drives |
|  | Spare Drive |
|  | Transient Drive |
|  | Storage Enclosure |
|  | Port or Unassigned Physical Drive Connected to Boot Port |
|  | Tape Drive |
|  | Locked |
|  | License Manager/Encryption Manager |
|  | Cache Manager |

**Table 5-1.** Icons (continued)

| Image | Description |
|---|---|
|  | Array Diagnostic Report |
|  | Smart SSD Wear Gauge Report |
|  | None |

## 5.2. Keyboard Shortcuts

Keyboard functions and shortcuts are used for navigating or performing actions in the GUI. The following table provides keyboard shortcuts and their description.

**Table 5-2.** Keyboard Shortcuts

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

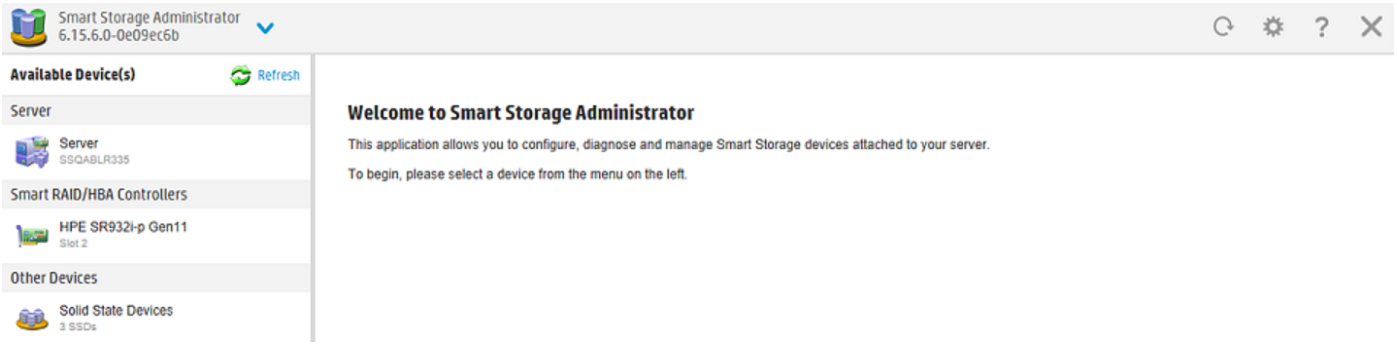
**Note:**

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

## 5.3. Navigating the GUI

When you open SSA, the **Welcome** screen appears. The following figure shows the SSA welcome screen.

**Figure 5-1.** SSA Welcome Screen

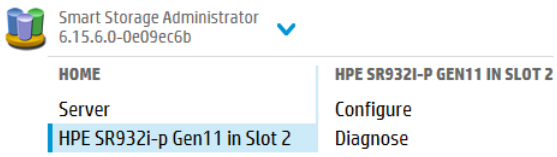



The following elements are visible:

- The SSA quick navigation menu is at the top, left-hand corner of the screen. Click the down arrow to view available devices, and click one of the available devices to view 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 Panel or Diagnostics Screen.

**Figure 5-2.** SSA Welcome Screen—Configuration or Diagnostics



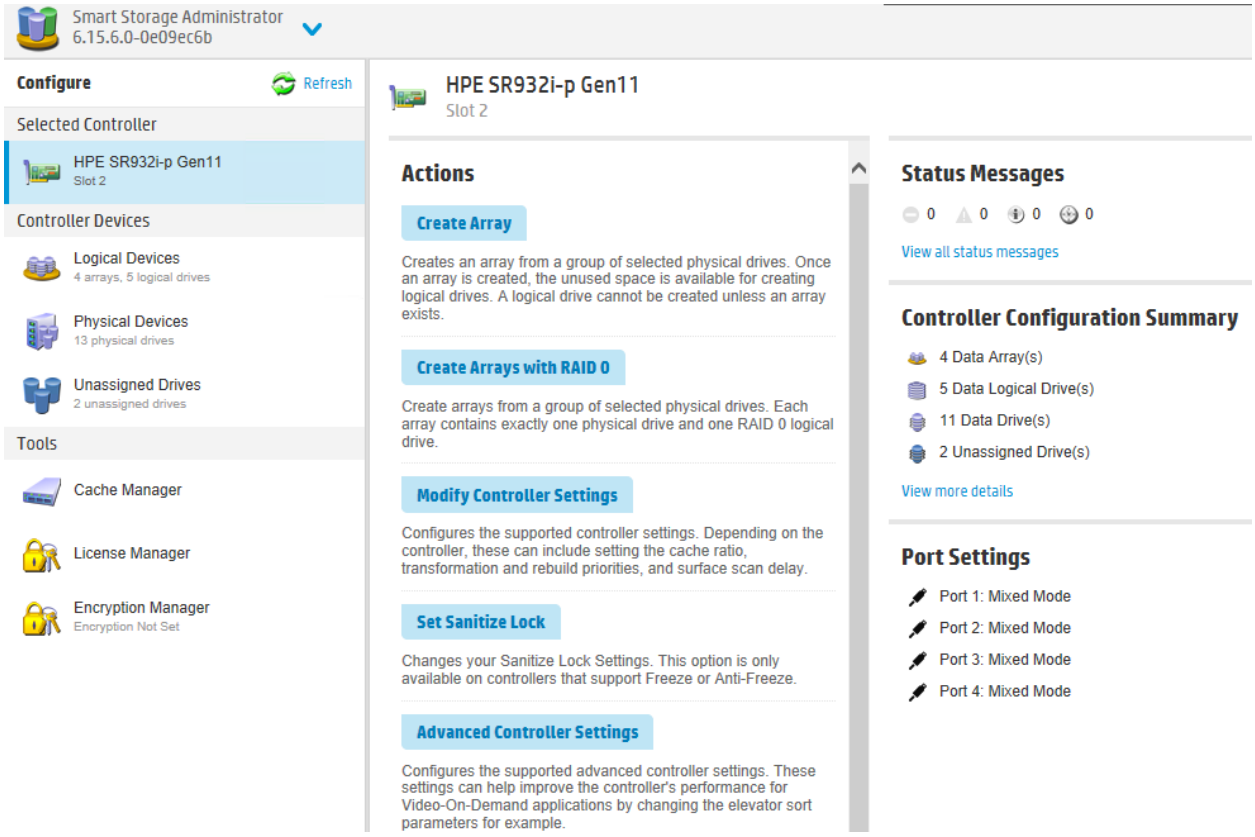

- **Available device(s)** are listed on the left-hand side of the screen. Clicking on a server or array controller displays the available actions, alerts, and summary for that device. You can point to the status alerts to see details on an alert.
- The **Rescan System and Reset Application** button is at the top right corner of the screen. After adding or removing devices, click **Rescan System and Reset Application** to update the list of available devices.
- The **Application Settings** button is used to select the language.
- 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 SSA Help.
- The **Exit** button is near the top right of the screen.

### 5.3.1. Configure Panel

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.

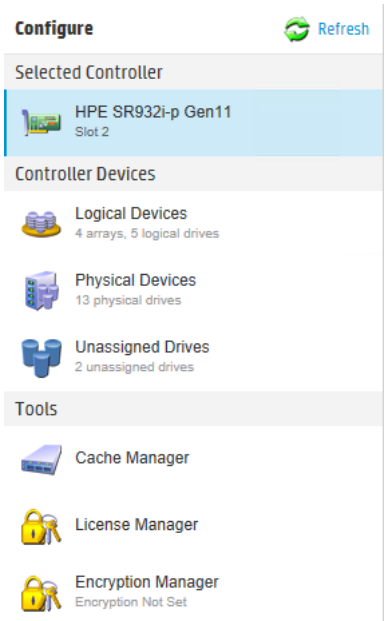
**Figure 5-3.** Configure Panel




When a controller is selected, the following elements appear:

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

**Figure 5-4.** Configure Panel



- **Actions**—Provides the following information and functionality:
  - Tasks that are available for the selected device based on its current status and configuration.
  - Task options and information, after a task is selected.
- **Status Messages**—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**—Provides a summary of the following elements for the selected controller:
  - Data arrays
  - Data logical drives
  - Data drives
  - Unassigned drives
  - Spare drives
  - A **View more details** link that displays more information in a pop-up window
- **Port Settings**—Provides more information about the drives attached to the controller. The values vary according to the controller. For controllers, mixed mode combines RAID and HBA modes, occurs by default, and cannot be disabled.

For more information on the available tasks on the **Configure** dialog box or window, see Configuration Tasks.

### 5.3.2. Diagnostics Screen

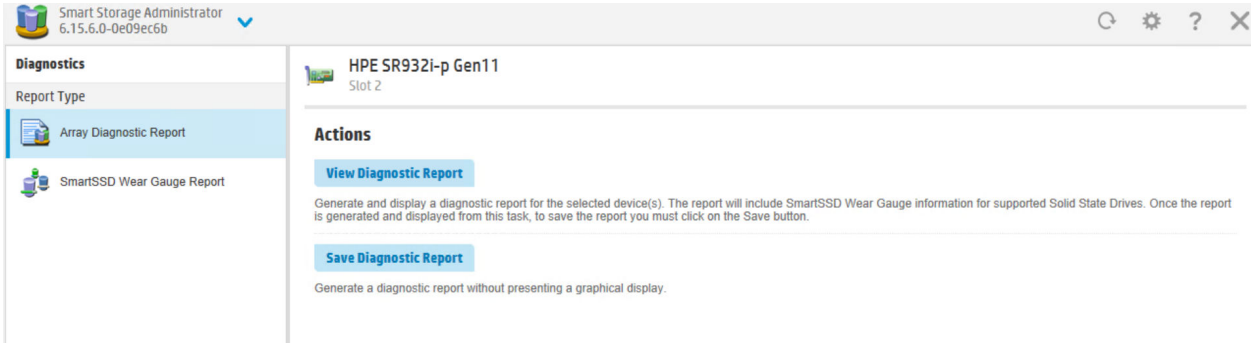
To access this screen, either click a device under **Diagnostics** in the quick navigation menu, or select an available device from the **Home** screen, and then click **Diagnose** under the available options.

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

- Array Diagnostics Report
- SmartSSD Wear Gauge Report

When you select either of the reports, the **View Diagnostic Report** and **Save Diagnostic Report** options are displayed on the **Actions** panel.

**Figure 5-5.** Diagnostics Screen




For more information on the available tasks in the **Diagnostics** window, see Managing Diagnostics.

#### 5.3.2.1. SmartSSD Wear Gauge Report

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

- Total Solid State Drives with Wearout Status

- Total SmartRAID Solid State Drives
- Total Non-SmartRAID Solid State Drives
- Total Solid State Drives

You can generate a report with a graphic representation of the SSD usage and estimated lifetime information or a report without a graphical display, with the option of saving the report.

### 5.3.3. SSA Help

The **Help** button, opens the embedded SSA help file. In addition to providing information about the main screens and tabs, it describes the following:

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

To access help, press the **H** key or click **Help**. When the **Help** window displays, expand the topic **Getting Started with Smart Storage Administrator Overview**.

The glossary in SSA help defines industry standard and related terms as they relate to the SSA application.

# 6. Configuration Tasks

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

When a controller or device is selected, the tasks that appear are a subset of the total number of possible tasks for the selected item. 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.

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

For any compliance issues arising from your operation/usage of encryption within the SmartRAID controller module which violates the above mentioned requirement, you shall bear all the liabilities wholly and solely. The SmartRAID controller vendor is not responsible for any related liabilities.

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

**Table 6-1.** Configuration Tasks

| Item | Tasks |
|---|---|
| Controller | Accelerated I/O Path |
| | Advanced Controller Settings[1] |
| | Array Accelerator Settings |
| | Check Online Firmware Activation Readiness |
| | Clear Configuration |
| | Controller Settings |
| | Create Array |
| | Disable Standby Controller |
| | Enable HBA/RAID/SmartRAID Mode operations [1] |
| | Enable SmartCache |
| | Manage Encryption |
| | Manage License Keys [1] |
| | Modify Power Modes [1] |
| | More Information |
| | Parallel Surface Scan |
| | Physical Drive Write Cache Settings |
| | Redundancy Settings [1] |
| | Set Sanitize Lock |
| | View Status Alerts |

**Table 6-1.** Configuration Tasks (continued)

| Item | Tasks |
|---|---|
| Array | Bypass RAID components using SSA Smart Path<br>Change Array Drive Type<br>Create Array<br>Create Logical Drive<br>Create Split Mirror Backup<br>Convert Plaintext Data to Encrypted Data<br>Delete<br>Expand Array<br>Heal Array<br>Manage Split Mirror Backup<br>More Information<br>Move Drives<br>Re-Mirror Array<br>Shrink Array<br>Spare Management<br>Split Mirrored Array<br>View Status Alerts<br>Volume Key Rekey |
| Logical drive | Create Logical Drive<br>Create SmartCache for Logical Drive<br>Convert Plaintext Data to Encrypted Data<br>Delete<br>Extend Logical Drive<br>Instant Secure Erase<br>Migrate RAID/Stripe Size<br>Modifying Cache Write Policy<br>Move Logical Drive [1]<br>More Information<br>Re-enable Failed Logical Drive<br>View Status Alerts<br>Volume Key Rekey |
| Unused space | Create Logical Drive<br>More Information |
| Physical drive | Erase Drive<br>View Status Alerts<br>Identify Device |
| Unassigned drive | Create Array<br>Erase Drive<br>More Information<br>Identify Device |

**Note:**

1. This task is not available on all controller models.

## 6.1. Performing a Configuration

To perform a configuration, follow these steps:

1. Open SSA.

For more information, see Accessing SSA.

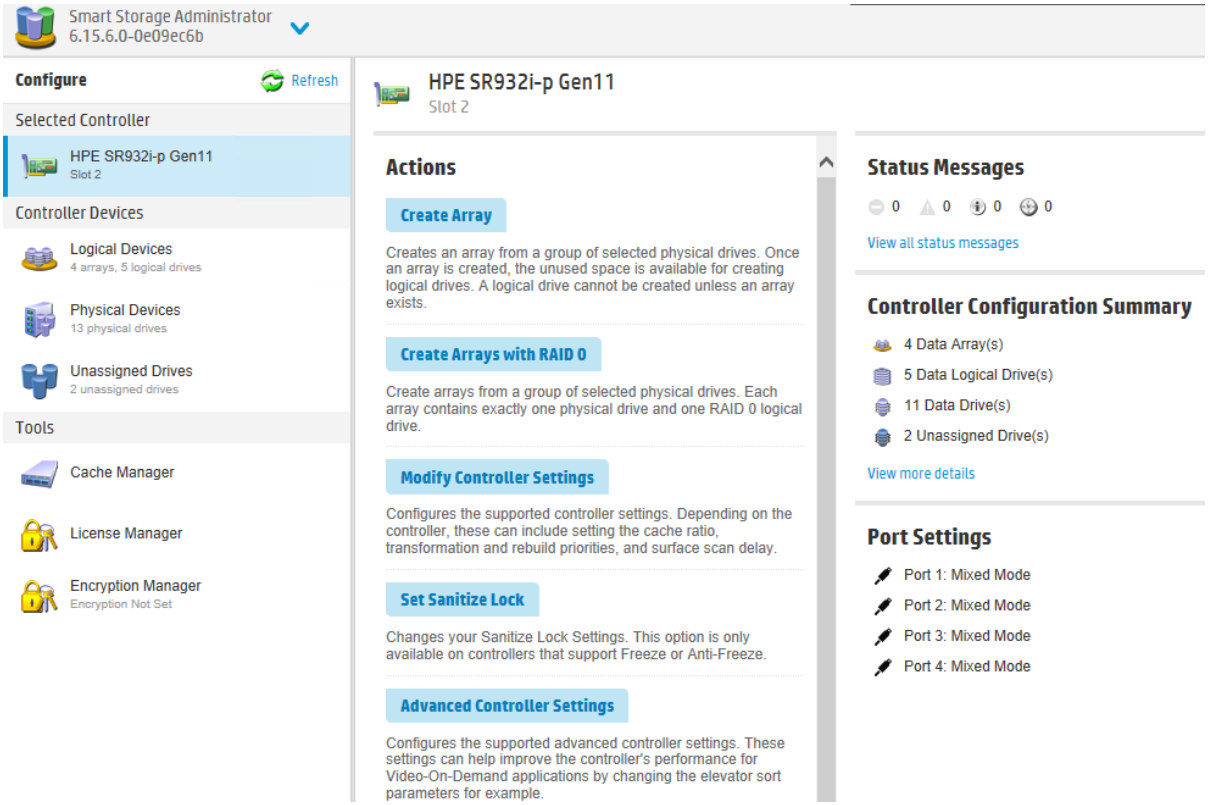
2. Open the **Configure** panel by performing one of the following:
   - Select a device in the **Home** screen and click **Configure** in the quick navigation menu.
   - Click on the drop-down next to the SSA icon. Hover over the device you want to select and click **Configure**.

**Figure 6-1.** Configure from Home Screen




3. Select a device from the **Controller Devices** menu.
   The **Actions**, **Status Messages**, **Controller Configuration Summary**, and **Port Settings** panels appear. The listed tasks are available for this device in its current configuration. For more information, see Configuration Tasks.

**Figure 6-2.** Devices Menu




4. In the **Actions** panel, click on the action you want to perform.
5. Select the settings or configuration options for the device.
6. Use the **Next** and **Back** buttons to navigate through multiple screens for more options.
7. Click **Save** or **OK**.

## 6.2. Creating an Array

To configure RAID and other settings for the array/logical drive, perform the following steps:

1. Open SSA.
2. Open the **Configure** panel by performing one of the following:
   - Select a device in the **Home** screen and click **Configure** in the quick navigation menu.
   - Click on the drop-down next to the SSA icon. Hover over the device you want to select and click **Configure**.
3. Select a controller from the **Controller Devices** menu.

   The **Actions** window appears.
4. Click **Create Array**.
5. Select the physical drives for the new array.
6. Click **Create Array**.
7. Select appropriate **RAID Level**, **Strip Size/Full Stripe Size**, **Sectors/Track**, and **Size** based on your requirements.
8. Click **Create Logical Drive**.
9. Click **Finish**.

By default, SmartPath is enabled on volumes created on SSD drives, and controller cache is not configured for this volume. SmartPath needs to be disabled in order to configure controller cache on the volumes configured on the SSD drives.

## 6.3. Configuring a Controller

To configure a controller, perform the following steps:

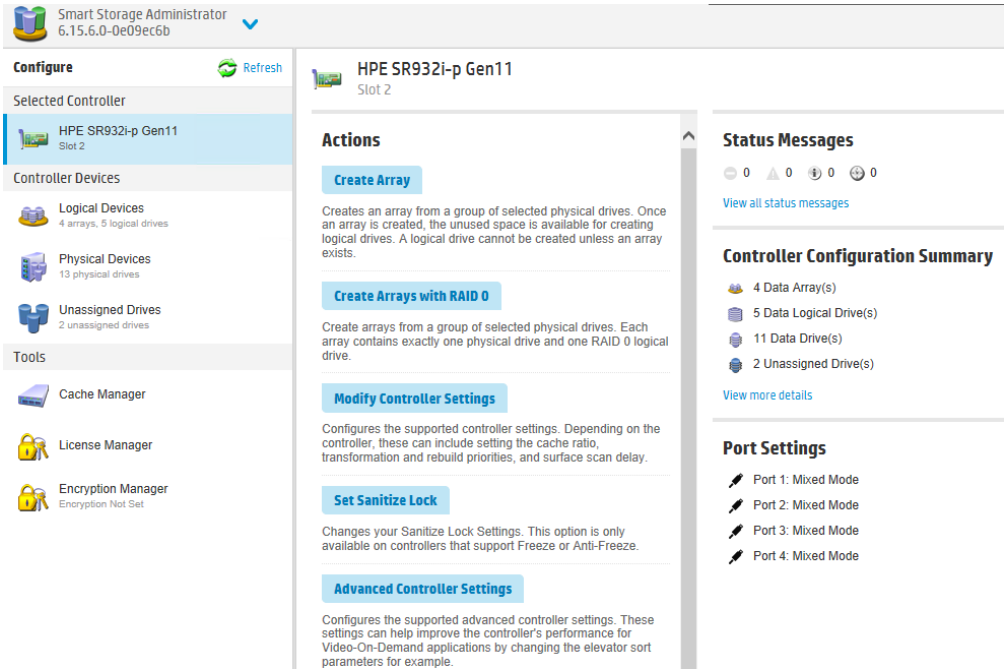
1. Open SSA.
2. Open the **Configure** panel by performing one of the following:
   - Select a device in the **Home** screen and click **Configure** in the quick navigation menu.
   - Click on the drop-down next to the SSA icon. Hover over the device you want to select and click **Configure**.

   The **Configure** panel appears.



**Important:** Screens display different options, depending on the server configuration.

**Figure 6-3.** Configure Panel




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

## 6.4. Enabling or Disabling SSD Smart Path

When you create an array, SSD SmartPath is enabled by default. For more information, see Creating an Array.
To enable or disable SSD SmartPath, perform the following steps:

1. Open SSA.
2. Open the **Configure** panel by performing one of the following:
   - Select a device in the **Home** screen and click **Configure** in the quick navigation menu.
   - Click on the drop-down next to the SSA icon. Hover over the device you want to select and click **Configure**.
3. Select a controller from the **Controller Devices** menu.
   The **Actions** panel appears.
4. Navigate to **Configure** > **Logical Devices** and select the array with logical drives.
5. Click **Disable SSD Smart Path**.
6. Click **Save**.

## 6.5. Setting Sanitize Lock

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

**Note:** This feature is applicable to SATA drives only and is not applicable to the SmartRAID S100i.

To set sanitize lock, perform the following steps:

1. Open SSA.
2. Open the **Configure** panel by performing one of the following:
   - Select a device in the **Home** screen and click **Configure** in the quick navigation menu.
   - Click on the drop-down next to the SSA icon. Hover over the device you want to select and click **Configure**.
3. Select a controller from the **Controller Devices** menu.

   The **Actions** panel appears.
4. In the **Actions** panel, click **Set Sanitize Lock**.
5. In the **Set Sanitize Lock** panel, select one of the following options:
   - **None**: This is the Normal state of the physical disk. Freeze or anti-freeze commands are not sent to any drives.
   - **Freeze**: This setting prevents a drive sanitize operation.
   - **Anti-freeze**: This setting prevents physical disks from being frozen. It enables drive sanitize operations.
6. After you modify the settings, restart the server.

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

## 6.6. Rapid Parity Initialization (RPI)

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

The RPI method works by overwriting both the data and parity blocks in the foreground. The logical drive remains invisible and unavailable to the OS until the parity initialization process is completed. Keeping the logical drive offline eliminates the possibility of any I/O activity, which speeds up the initialization process. This also enables other high-performance initialization techniques, which cannot be used if the volume is available for I/O. After the parity is completed or initialized, the drive is brought online and becomes available to the OS.

This method has the following benefits:

- Speeds up the parity initialization process.
- Ensures that parity volumes use backed-out writes for optimized random write performance.

### 6.6.1. Initializing Rapid Parity

To initialize rapid parity, perform the following steps:

1. Open SSA.
2. Open the **Configure** panel by performing one of the following:
   - Select a device in the **Home** screen and click **Configure** in the quick navigation menu.
   - Click on the drop-down next to the SSA icon. Hover over the device you want to select and click **Configure**.
3. Select **Logical Devices** from the **Devices** menu.

   A list of arrays appear.
4. Select an array

5. Click **Create Logical Drive** in the **Actions** menu.
6. Select appropriate **RAID Level**, **Number of Parity Groups**, **Strip Size/Full Strip Size, Sectors/ Track, Size**, **Parity Initialization Method**, and **Caching** based on your requirements.
   The displayed options differ based on the controller that is used and the selected drive.

**Figure 6-4.** Create Logical Drive




7. Click **Create Logical Drive** to continue.
   A summary page appears.
8. Click **Finish**.

## 6.7. Changing the Spare Activation Mode

The Spare Activation mode enables the controller firmware to activate a spare drive under the following conditions:

- When a data drive reports a predictive failure status.
- When a data drive fails; it is in the Default mode.

In normal operations, the firmware starts rebuilding an additional drive only when a data drive fails. With Predictive Failure Activation mode, rebuilding can begin before the drive fails, reducing the possibility of data loss that might occur if an additional drive fails.
To change Spare Activation mode, perform the following steps:

1. Open SSA.
2. Open the **Configure** panel by performing one of the following:
   - Select a device in the **Home** screen and click **Configure** in the quick navigation menu.
   - Click on the drop-down next to the SSA icon. Hover over the device you want to select and click **Configure**.
3. Select a controller from the **Controller 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**.

## 6.8. Changing the Spare Management Mode

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

- **Dedicated Spare Drives**: 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, you do not need to rebuilt the data drive. 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, perform the following steps:

1. Open SSA.
2. Open the **Configure** panel by performing one of the following:
   - Select a device in the **Home** screen and click **Configure** in the quick navigation menu.
   - Click on the drop-down next to the SSA icon. Hover over the device you want to select and click **Configure**.
3. Select a controller from the **Controller Devices** menu.
4. Click **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 appropriate **Create Plaintext Volume**, **RAID Level**, **Strip Size/Full Strip Size**, **Sectors/Track**, **Size**, and **Caching** based on your requirements.
7. Click **Create Logical Drive** .
8. Click **Manage Spare Drives**.
9. From the menu, select one of the following spare drive types:
   - **Dedicated Spare Drives**
   - **Auto-Replace Drives**
10. Select which drives will operate as spare drives in the array.
11. Click **Save**.

    A confirmation screen appears.

12. Click **Yes** to continue.
13. Click **Manage Spare Drives** to make additional selections, or click **Finish** to complete the task.

## 6.9. Disabling Over Provisioning Optimization

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

The SSD Over Provisioning Optimization feature allows you to disable Over Provisioning Optimization in the GUI. To disable over provisioning optimization, perform the following steps:

1. Open SSA.
2. Click **Unassigned Drives** under **Controller Devices**.
3. To create an array, select a device from the list of available drives. 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**.

## 6.10. Setting Auto RAID 0

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

**⚠ CAUTION** If you select this option for any logical drives, you experience a 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.

For more information about RAID 0, see RAID 0 .
To set auto RAID 0, perform the following steps:

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

## 6.11. Modifying Parallel Surface Scan

To modify parallel surface scan, perform the following steps:

1. Open SSA.
2. Select a controller.
3. Click **Modify Controller Settings**.
4. Under **Current Parallel Surface Scan Count**, use the slider to select the parallel surface scan count.
5. Click **Save Settings**.
   A summary page appears.
6. Click **Finish** to exit.

## 6.12. Configuring Controller Cache

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

To configure controller cache, perform the following steps:

1. Open SSA.
2. Open the **Configure** panel by performing one of the following:
   - Select a device in the **Home** screen and click **Configure** in the quick navigation menu.
   - Click on the drop-down next to the SSA icon. Hover over the device you want to select and click **Configure**.
3. Click **Cache Manager** from the **Tools** menu.
4. Click **Controller Cache** from 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**.

## 6.13. SmartRAID 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 OS without any changes to OS, driver, or applications.

SmartCache is fully enabled when the first SmartCache is created on the controller.

The following features are not enabled unless SmartCache is disabled:

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

SmartCache requires controller cache and a backup power source.

### 6.13.1. Enable SmartRAID SmartCache

Enabling SmartCache with an array accelerates 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 SmartRAID SmartCache, perform the following steps:

1. Open SSA.
2. Open the **Configure** panel by performing one of the following:
   - Select a device in the **Home** screen and click **Configure** in the quick navigation menu.
   - Click on the drop-down next to the SSA icon. Hover over the device you want to select and click **Configure**.
3. Click **Cache Manager** from the **Tools** menu.
4. Click **Enable SmartCache** from the **Actions** menu.
5. Select one or more physical drives from the list of available drives.
6. Click **OK**.
7. A pop-up window appears, indicating certain features are not available when SmartCache is enabled. If you want to continue, click **Yes**.
8. The dialog **Create SmartCache for Logical Drive** appears.
9. Select the following:
   - Logical drive to cache
   - Cache write policy and RAID type
   - Size of the cache. It is recommended to select 10% of the drive size, and it must be a minimum of 16 GiB.
10. Click **Create SmartCache for Logical Drive**.
11. **SmartCache Details**, **SmartCache Statistics**, and **Device Path** appear. Click **Finish** to proceed.
    The SmartCache is created for the logical drive.

### 6.13.2. Caching Write Policies

The following sections describe caching write policies.

#### 6.13.2.1. SmartCache Write-back

SSA contains two different policies for data writes when caching: write-back and write-through. Write-back is a caching method where the 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 comes with the risk of losing data if the cache volume fails.
Write-through is a caching method, where the 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 compared to the write-back policy this method's performance is low.
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 find an appropriate replacement 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 the newest to oldest, in descending order from how they are 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 can 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.

#### 6.13.2.2. Modifying the Physical Drive Write Cache Policy Settings

The procedure varies according to whether your controller supports modifying the policy on configured drives or on configured and unconfigured drives.

##### 6.13.2.2.1. If your Controller Supports Drive Write Caching of Configured Drives Only

The Physical Drive Write Cache State setting appears in the Modify Controller Settings panel if your controller supports drive write caching of configured drives only. To modify physical drive write cache policy setting, perform the following steps:

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

##### 6.13.2.2.2. If your Controller Supports Drive Write Caching of Configured and Unconfigured Drives

The Manage Drive Write Cache Policy option appears on the controller Actions menu if your controller supports drive write caching of configured and unconfigured drives. To modify physical drive write cache policy setting, perform the following steps:

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

#### 6.13.2.3. Modifying the Write Cache Bypass Threshold

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

To modify the write cache bypass threshold, perform the following steps:

1. Open SSA.
2. Open the **Configure** panel by performing one of the following:
   - Select a device in the **Home** screen and click **Configure** in the quick navigation menu.
   - Click on the drop-down next to the SSA icon. Hover over the device you want to select and click **Configure**.
3. Click **Cache Manager** from the **Tools** menu.
4. Click **Controller Cache** from the **Cache Manager** menu.
5. Click **Modify Caching Settings**.
6. Select one or more drives.
7. In the **Write Cache Bypass Threshold** section, use the slider bar to select the value, or type the value into the field.

   A valid value is from 16 KiB to 1024 KiB. If the value that you type is not a multiple of 16, the user interface updates the value to the nearest valid value.
   **Note:** This applies only to non-parity RAID volume types.
8. Click **OK**.

## 6.14. Working with Mirrored Arrays

One of the advanced tasks in the 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.
The following are required to support these procedure:

- The SSA GUI must be run in offline mode. For more information, see Accessing SSA in the Offline Environment.
- Mirrored arrays being split can have RAID 1, RAID 1+0, RAID 1 (Triple), or RAID 10 (Triple) configurations. Arrays with other RAID configurations cannot be split.

### 6.14.1. Creating a Split Mirror Backup

This task splits an array that consists of one or more RAID 1, RAID 1+0, RAID 1 (Triple), or RAID 10 (Triple) logical drives, and then creates two arrays: a primary array and a backup array. To create a split mirror backup, perform the following steps:

1. Run the SSA GUI in offline mode. For more information, see Accessing SSA in the Offline Environment.
2. Open the **Configure** panel by performing one of the following:
   - Select a device in the **Home** screen and click **Configure** in the quick navigation menu.
   - Click on the drop-down next to the SSA icon. Hover over the device you want to select and click **Configure**.
3. Click **Arrays** from the **Controller 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**.

   SSA creates the array according to the following rules:
   - If the original array contained RAID 1 or RAID 1+0 drives, then the primary array contains RAID 0 drives.
   - If the original array contained RAID 1 (Triple) drives, the primary array contains RAID 1 drives.

- If the original array contained RAID 10 (Triple) drives, the primary array contains 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 OS.

8. When 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.

### 6.14.2. Recombining a Split Mirrored Array

To recombine a split mirrored array, perform the following steps:

1. Run the SSA GUI in offline mode. For more information, see Accessing SSA in the Offline Environment.
2. Open the **Configure** panel by performing one of the following:
   - Select a device in the **Home** screen and click **Configure** in the quick navigation menu.
   - Click on the drop-down next to the SSA icon. Hover over the device you want to select and click **Configure**.
3. Click **Arrays** from the **Controller 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 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 might take up to two hours. You can boot the OS during this time, but the logical drive is not fault-tolerant until the rebuild is complete.

### 6.14.3. Splitting a Mirrored Array

To split a mirrored array, perform the following steps:

1. Run the SSA GUI in offline mode. For more information, see Accessing SSA in the Offline Environment.
2. Open the **Configure** panel by performing one of the following:
   - Select a device in the **Home** screen and click **Configure** in the quick navigation menu.
   - Click on the drop-down next to the SSA icon. Hover over the device you want to select and click **Configure**.
3. Click **Arrays** from the **Controller Devices** menu.
4. From the **Arrays** menu, select the appropriate array.
5. Click **Manage Data Drives** from the **Actions** panel.
6. Under **Available Array Action(s)**, click **Mirror Array**.
7. Select a physical drive.
8. Click **OK**.

The mirrored array details are displayed.

9. Click **Finish**.
10. When 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 (Triple) logical drive splits, a RAID 1 logical drive and a RAID 0 logical drive are created.
    - When an array that contains a RAID 10 (Triple) 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. After 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 cannot distinguish between the two arrays when the server is restarted because the arrays are identical.
14. Power-up the server.
15. Restart the OS.

### 6.14.4. Re-mirroring, Rolling back, or Re-activating a Split Mirror Backup

To re-mirror, roll back, or re-activate a split mirror backup, perform the following steps:

1. Run the SSA GUI in offline mode. For more information, see Accessing SSA in the Offline Environment.
2. Open the **Configure** panel by performing one of the following:
    - Select a device in the **Home** screen and click **Configure** in the quick navigation menu.
    - Click on the drop-down next to the SSA icon. Hover over the device you want to select and click **Configure**.
3. Click **Arrays** from the **Controller Devices** menu.
4. Select the appropriate array from the **Arrays** menu.
5. In the **Actions** panel, click **Manage Split Mirror Backup**.
6. Select one of the following actions:
    - Re-mirror the array and save 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. It is recommended 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 OS
    - Activate the backup array.
      This option makes the backup array fully accessible to the OS. SSA removes the designation “backup” from the name of the array.

## 6.15. Modifying an Array

SSA allows you to modify an array by choosing the Manage Data Drives. The following section describes how to manage the data drives.

### 6.15.1. Managing Data Drives

To manage data drives, perform the following steps:

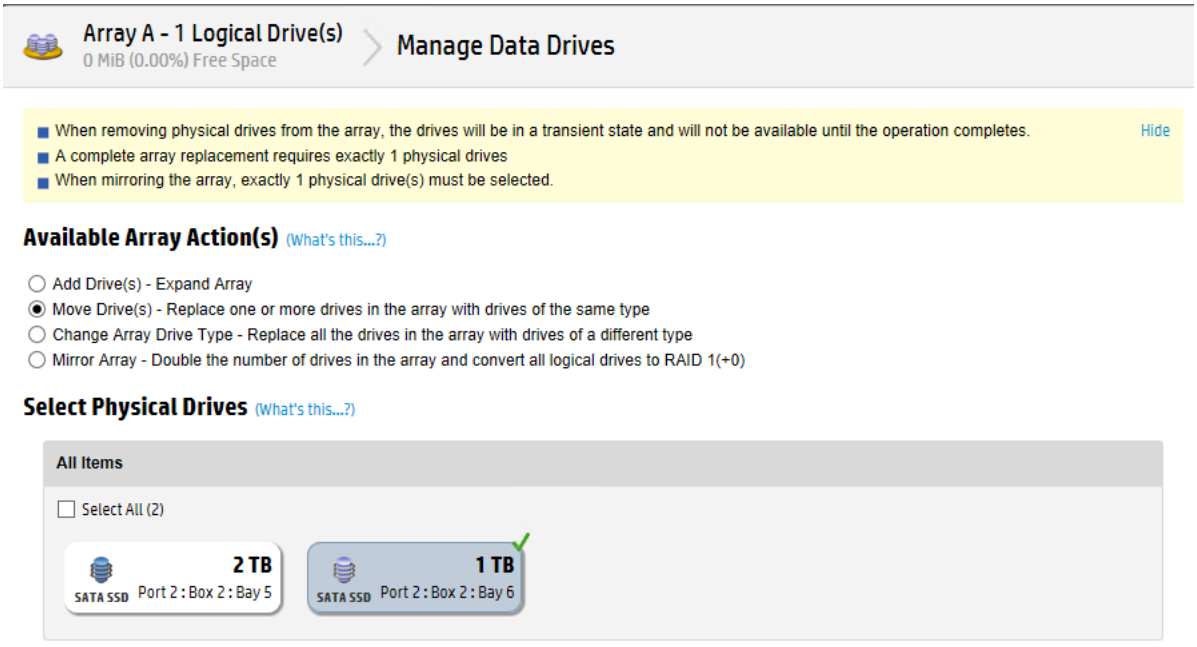
1. Open SSA.
2. Open the **Configure** panel by any of the following ways:
   a. Choose a device and click **Configure** in the quick navigation menu.
   b. Select an available device from the **Home** screen, and then click **Configure** under the available options.
3. Under **Controller Devices**, click **Logical Devices**.
4. Under **Actions** panel, click **Manage Data Drives**.

**Figure 6-5.** Manage Data Drives




5. Under **Available Array Action(s)**, click the required action. Under **Select Physical Drives**, select the required drive. A new window appears.

**Figure 6-6.** Manage Data Drives—Selection Window



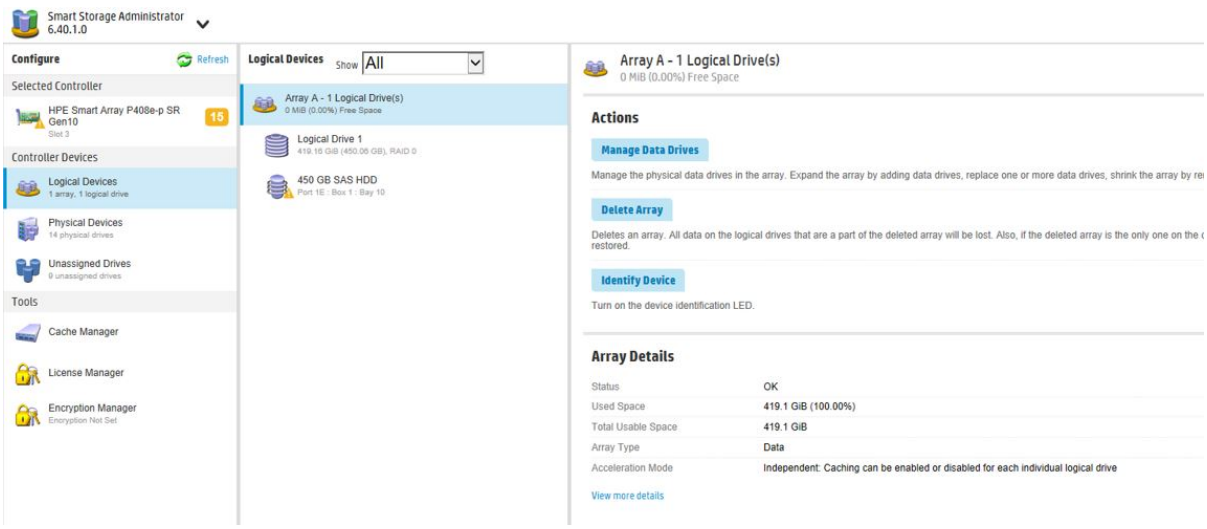

6. Click **Yes**.
7. Click **Finish**.

### 6.15.2. Expand an Array and Extend the Logical Drive

To expand an array and extend the logical drive, perform the following steps:

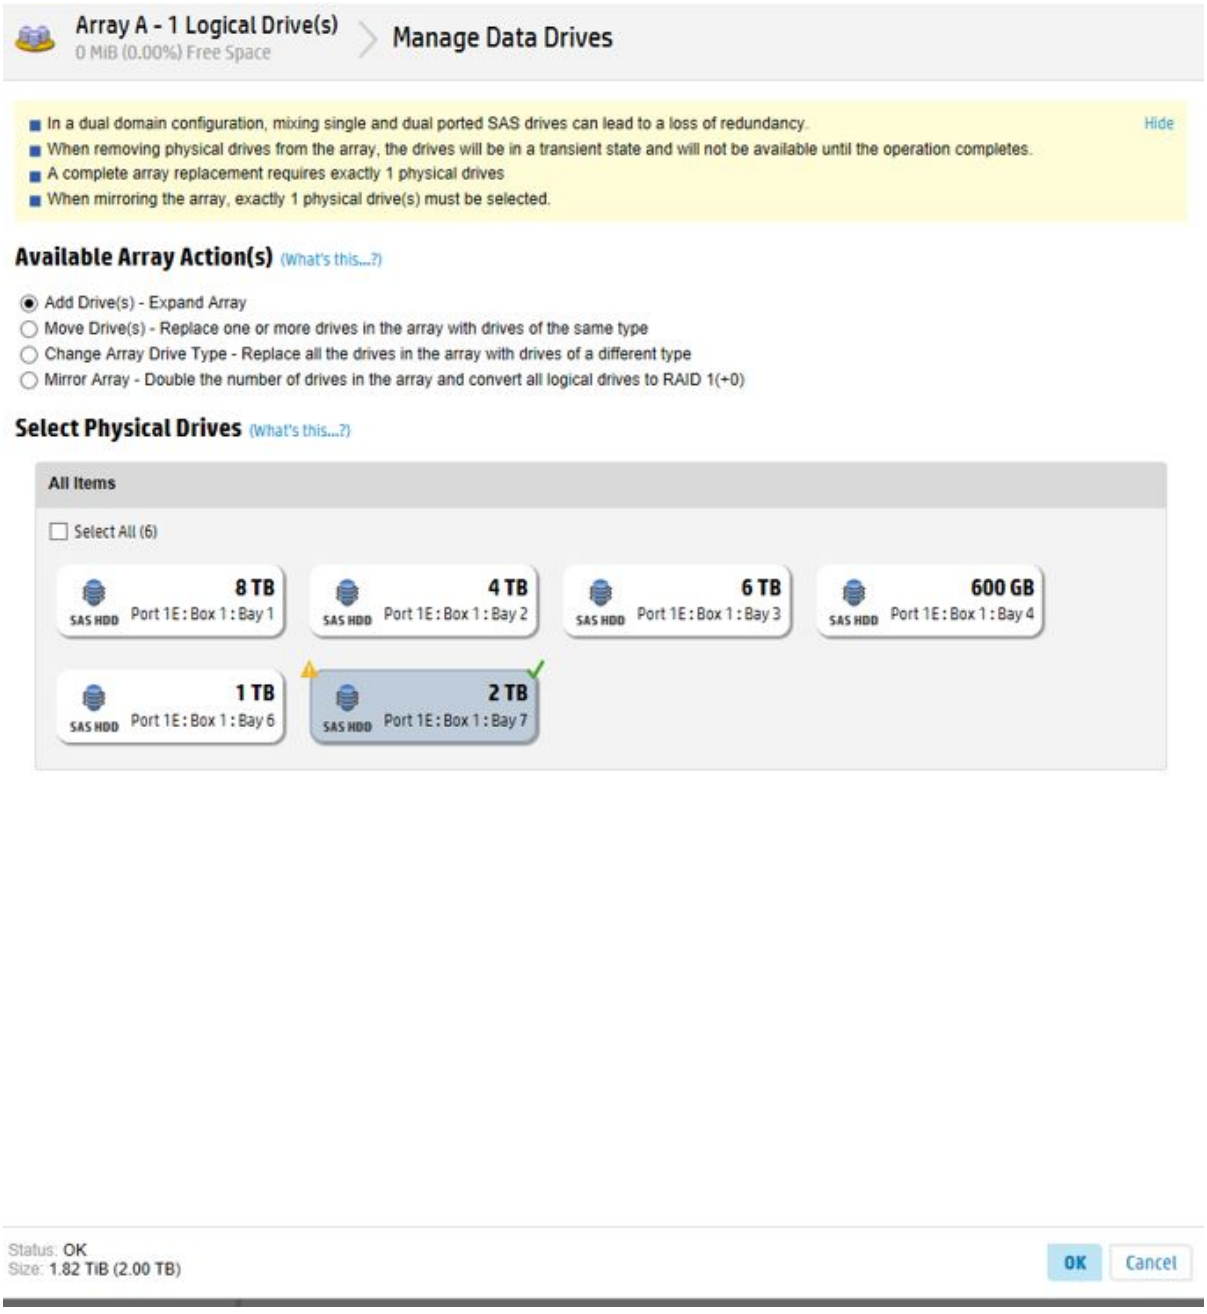
1. Open SSA.
2. Open the **Configure** panel by any of the following ways:
   a. Choose a device and click **Configure** in the quick navigation menu.
   b. Select an available device from the **Home** screen, and then click **Configure** under the available options.
3. Under **Controller Devices**, click **Logical Devices**.
4. Select the required Array in the **Logical Devices** panel. Under the **Actions** panel, click **Manage Data Drives**. A new window appears.

**Figure 6-7.** Manage Data Drives



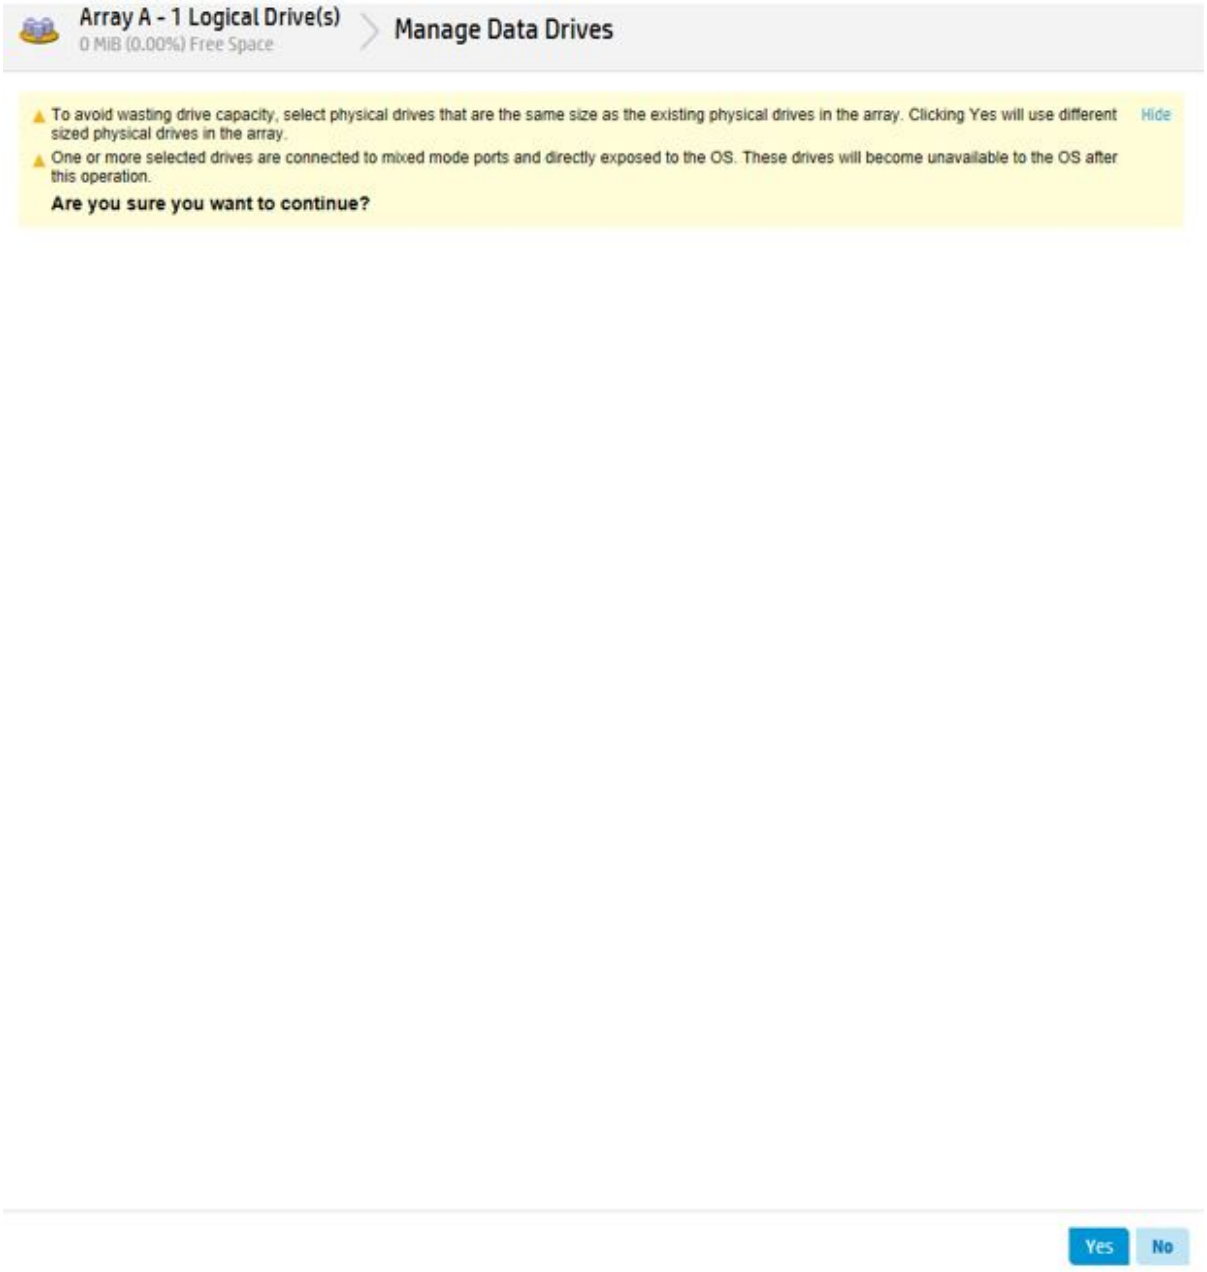
5. Under **Available Array Action(s)**, click **Add Drive(s) – Expand Array**. Under **Select Physical Drives**, select the required drive to add, and then click **OK**. A confirmation screen appears.

**Figure 6-8.** Manage Data Drives—Selection Window



6. Click **Yes** to continue. A summary window of the array appears.

**Figure 6-9.** Manage Data Drives—Confirmation Message



7. Verify the array summary, and then click **Finish**. Wait until the transformation operation is complete.

**Figure 6-10.** Manage Data Drives—Summary

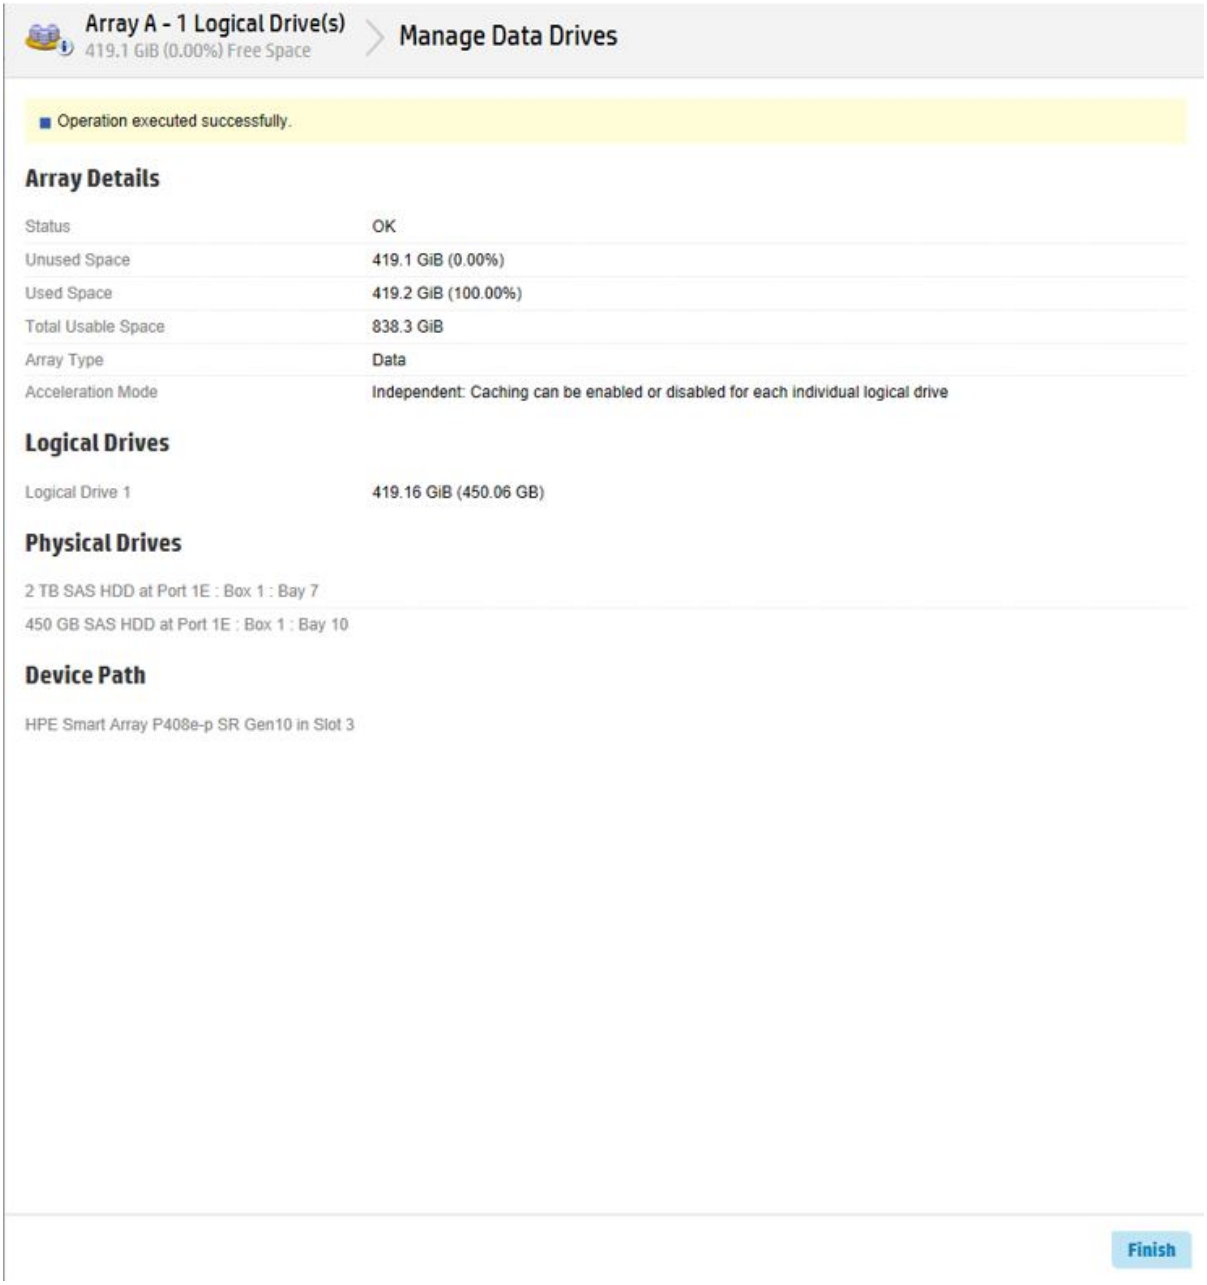
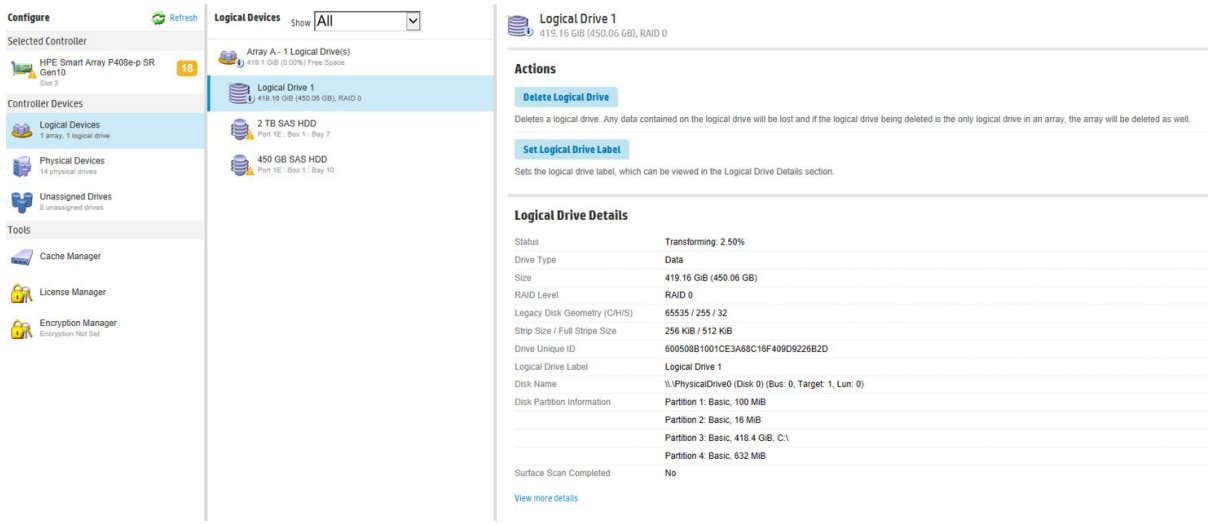
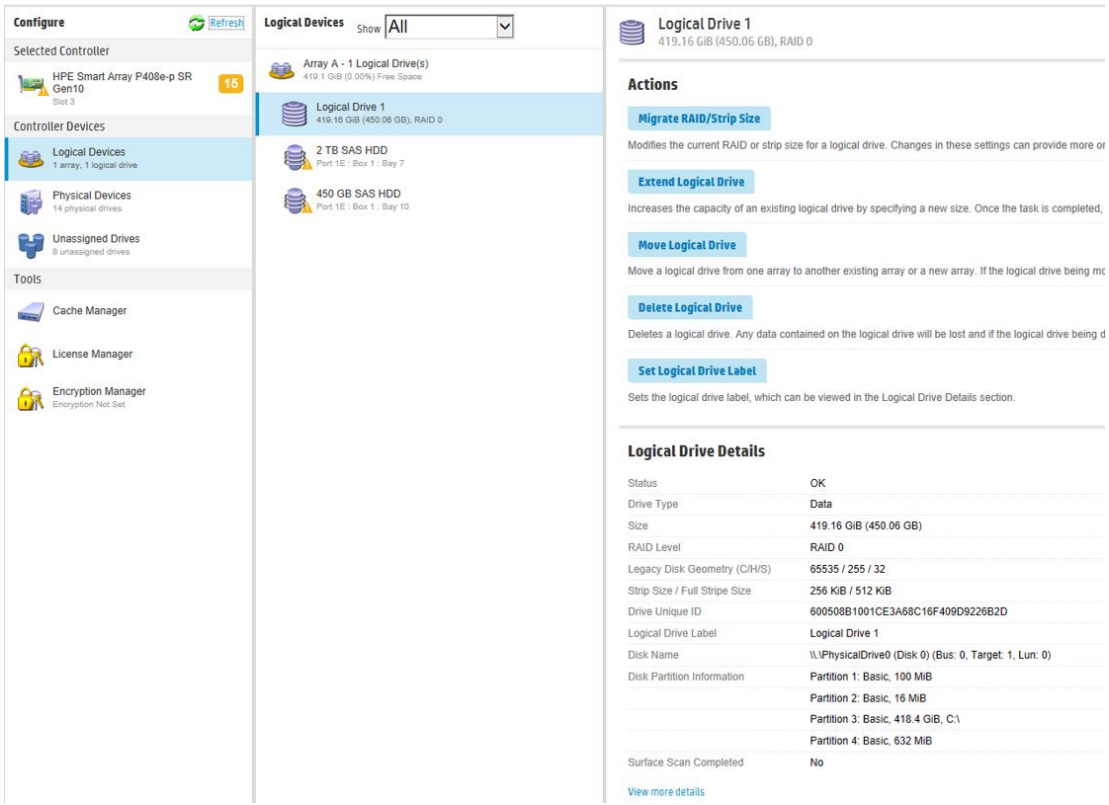


**Figure 6-11.** Expand Array—Transforming



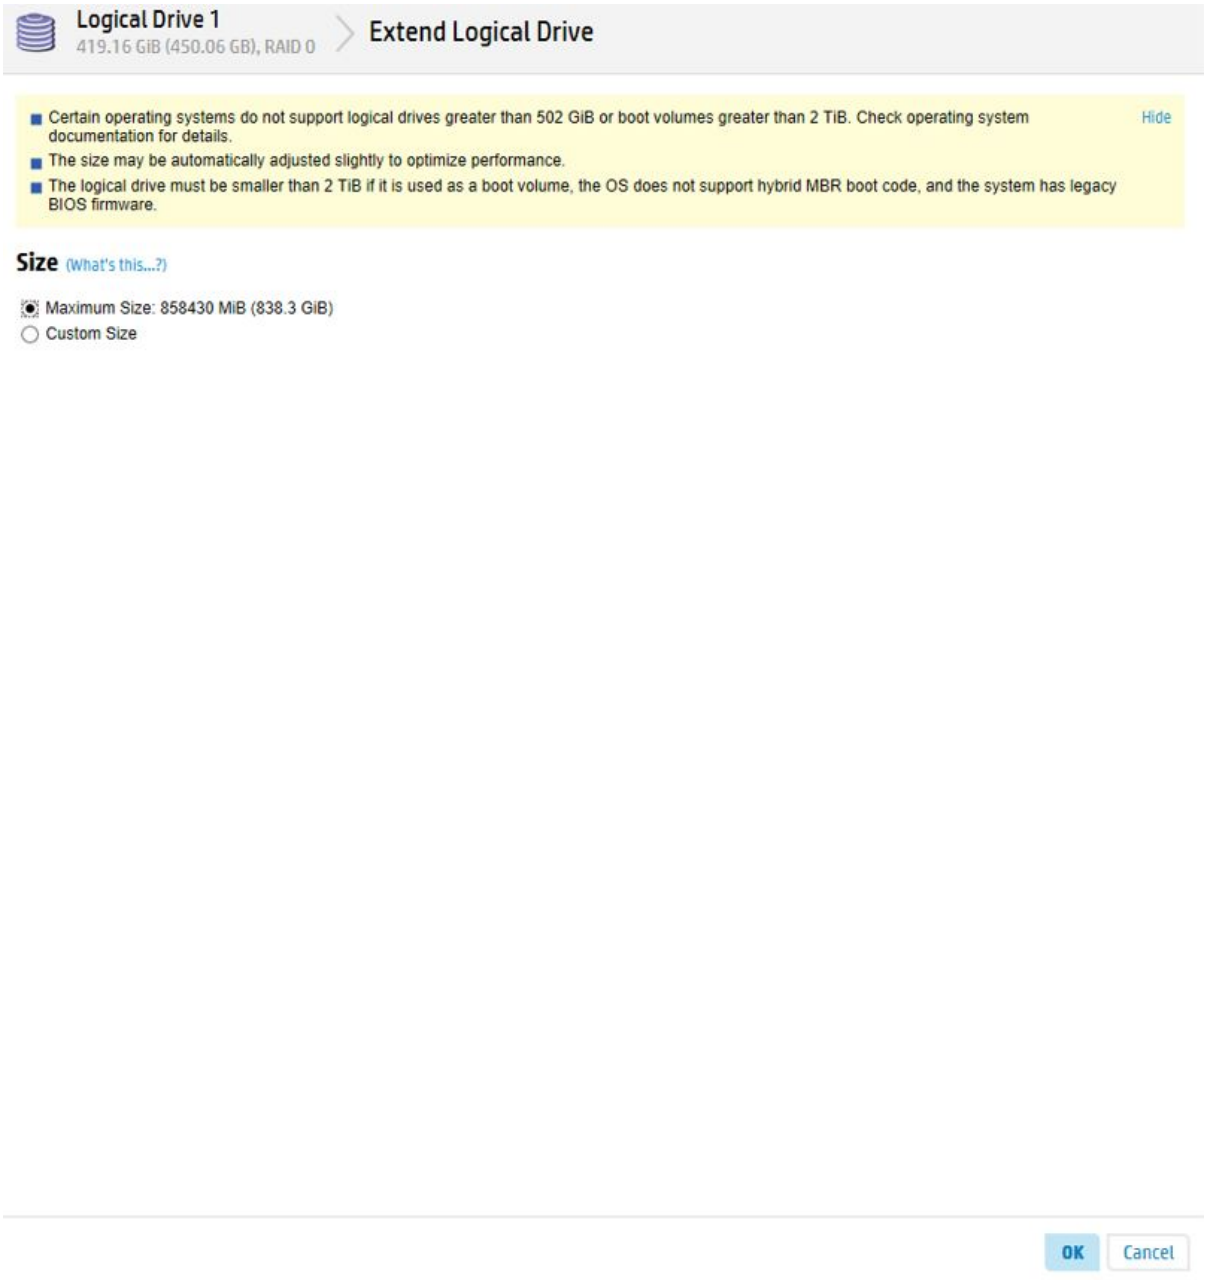

8. Once the transformation is complete, to extend the logical drive, under **Controller Devices**, click **Logical Devices**.
9. Select the required logical drive in the **Logical Devices** panel. In the **Actions** panel, click **Extend Logical Drive**. A new window opens.

**Figure 6-12.** Expand Array—Transformation Complete



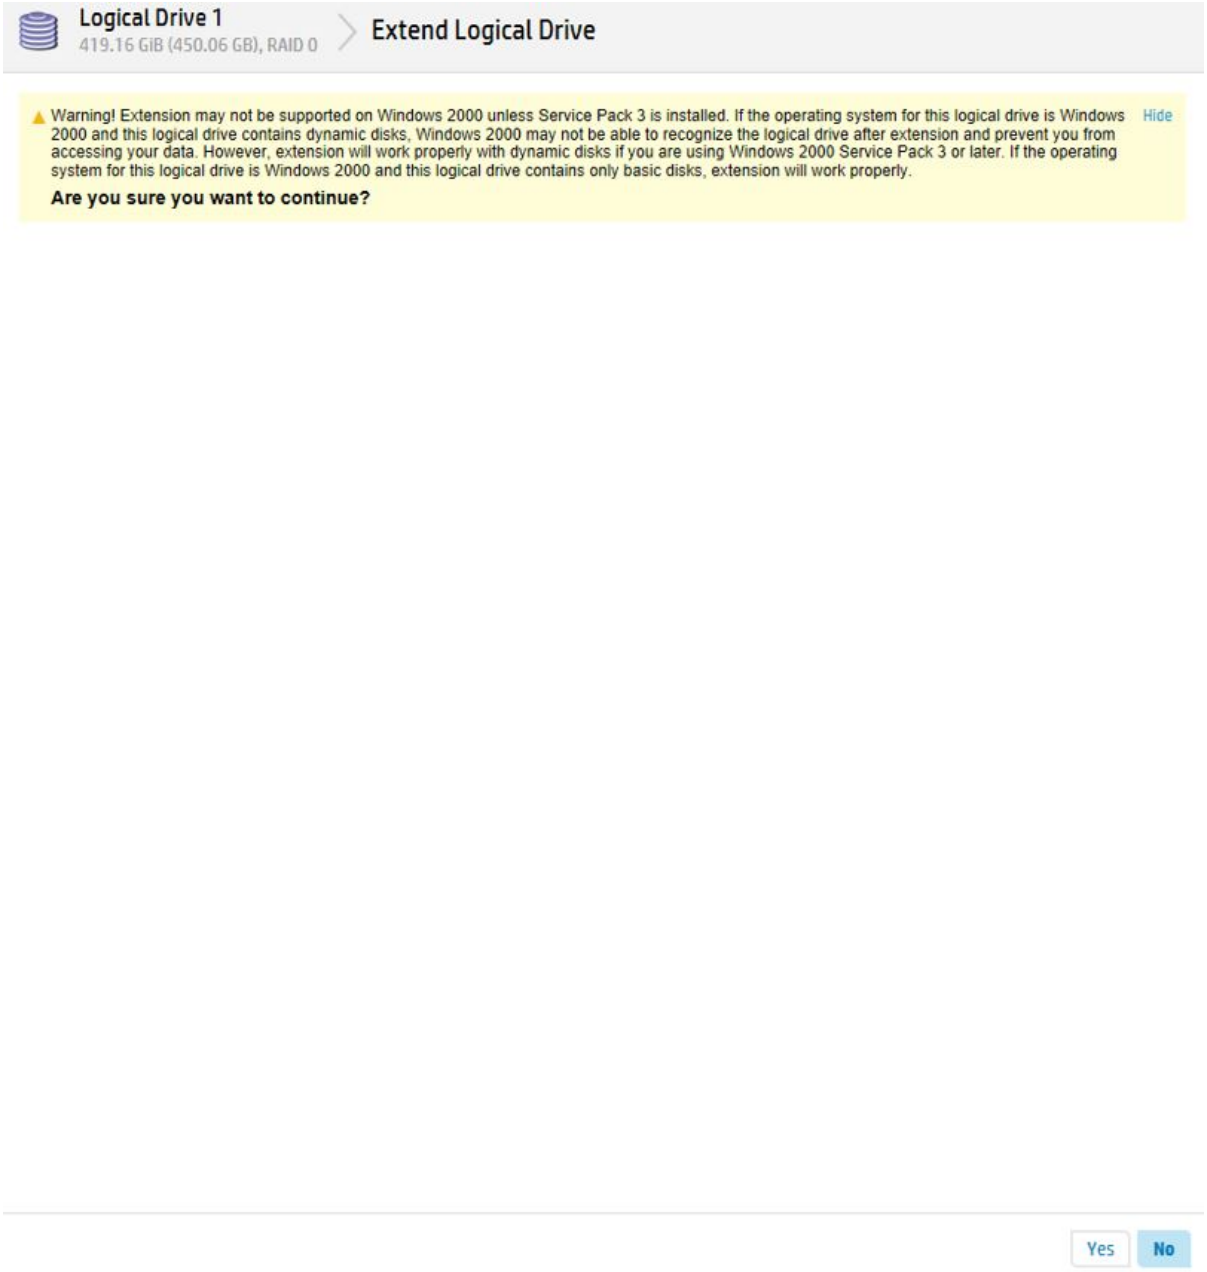
10. Select **Maximum Size** or **Custom Size**, and then click **OK**. A confirmation message appears.

**Figure 6-13.** Extend Logical Drive—Selection Window



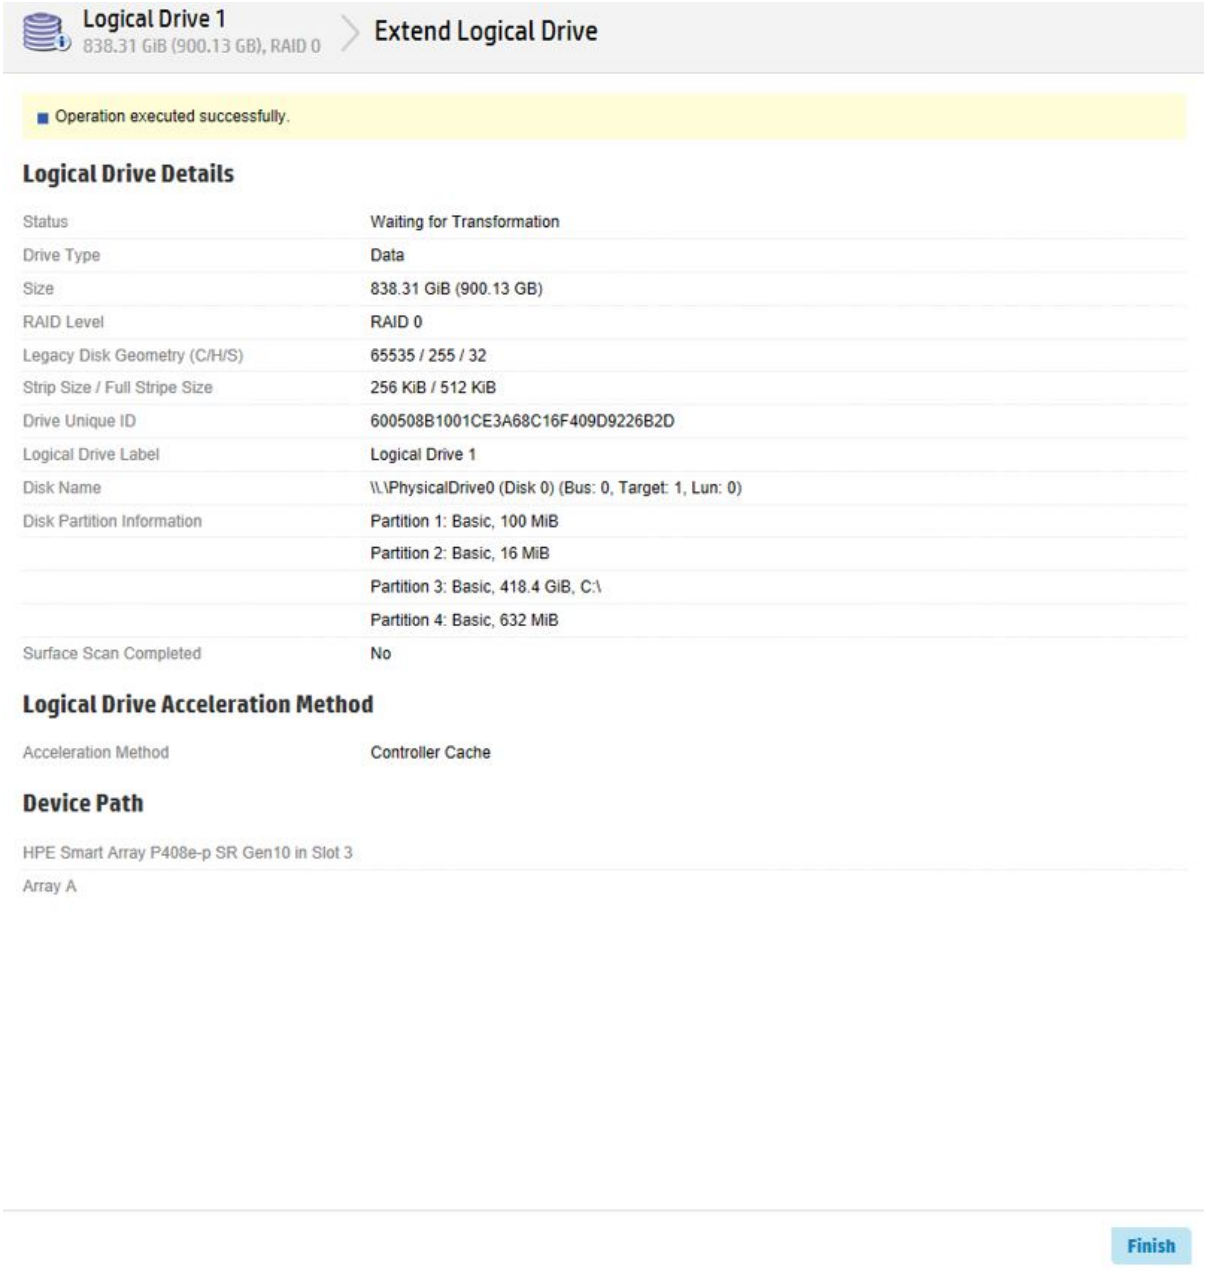
11. Click **Yes** to continue. A summary of the logical drive appears.

**Figure 6-14.** Extend Logical Drive—Confirmation Message



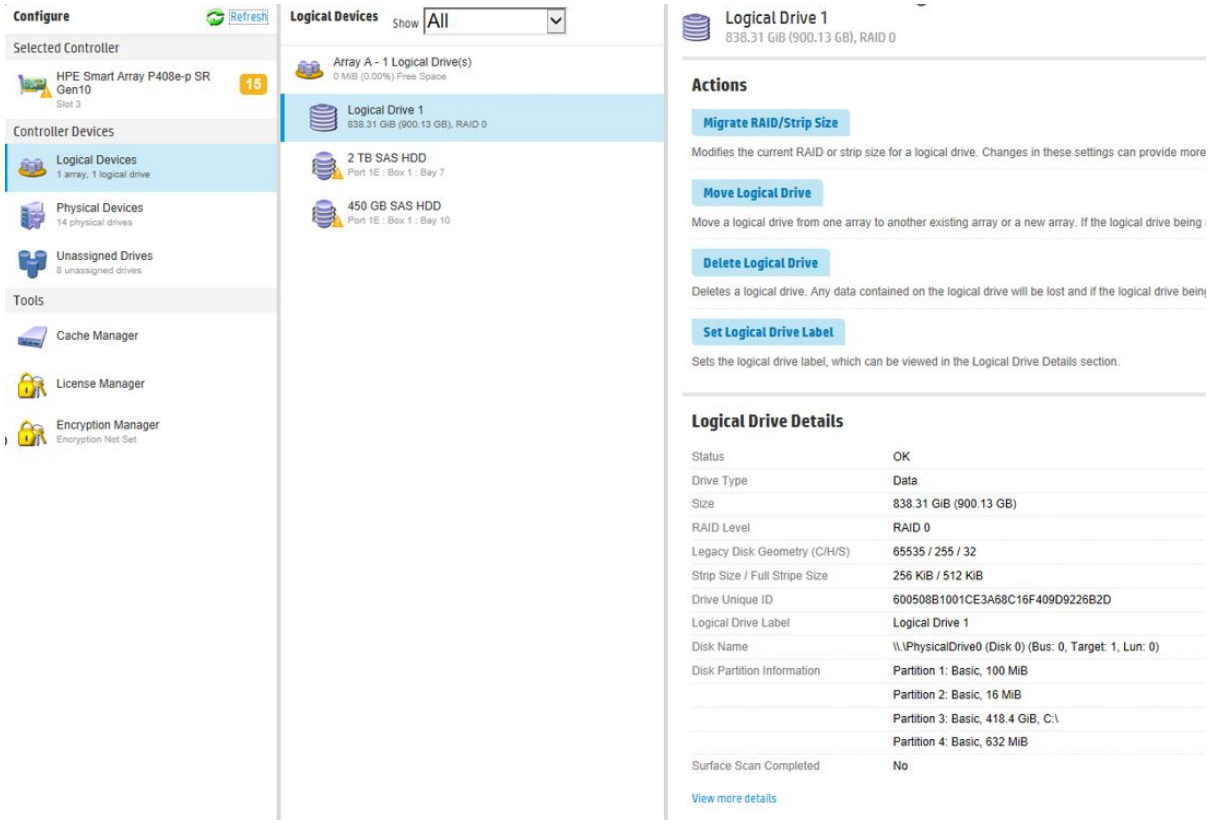
12. Click **Finish** to complete the process.

**Figure 6-15.** Extend Logical Drive—Summary Window



13. Verify the logical volume size after the completion of the extend operation under **Logical Drive Details**.

**Figure 6-16.** Logical Drive Details




## 6.16. Heal Array

The Heal Array operation allows you to replace failed physical drives in the array with healthy physical drives. The original array and logical drive numbering are not affected after the replacement. Note the following conditions and restrictions for the Heal Array operation:

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

## 6.17. Change Array Drive Type

Some controllers may not support this option or might require a license key to enable the feature. SSA allows 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 I/O 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

## 6.18. Moving a Logical Drive

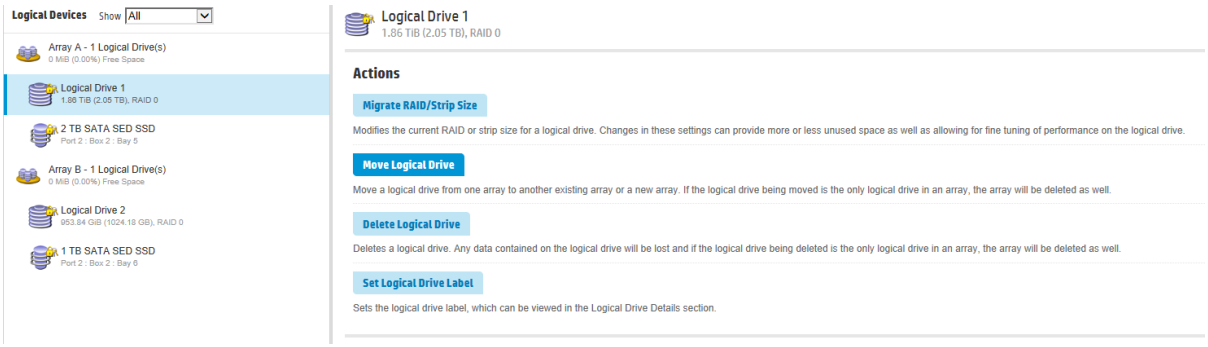
SSA allows you to move a single logical drive from one array to another array. You can choose the following destinations:

- Move logical drive to a new array
- Move logical drive to an existing array

To move a logical drive to a new array, perform the following steps:

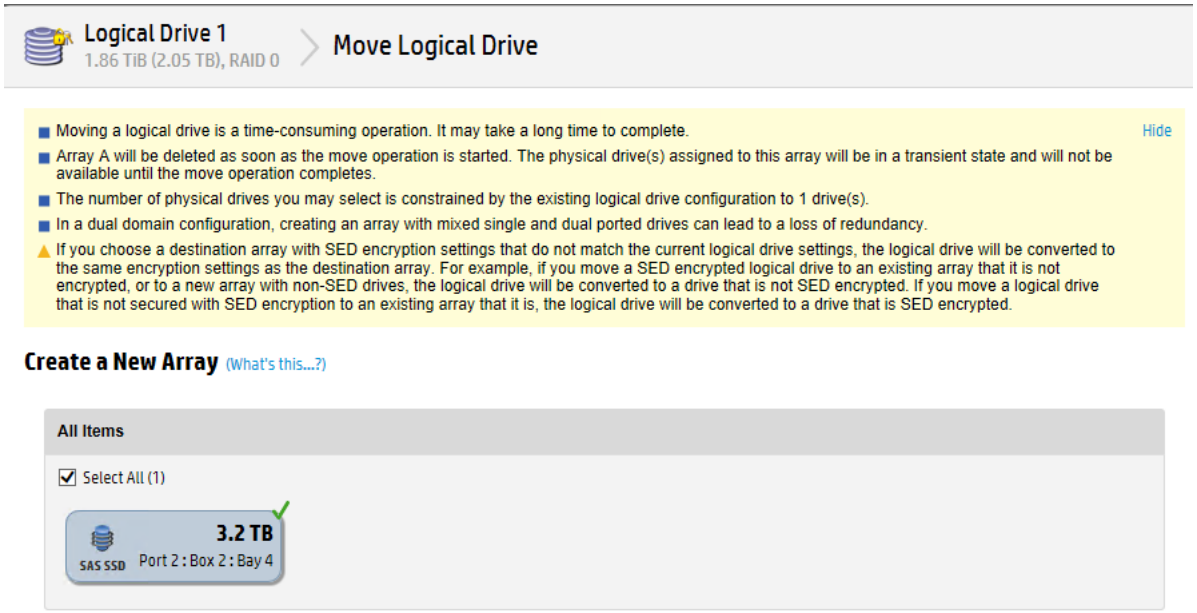
1. Open SSA.
2. Open the **Configure** panel by any of the following ways:
   a. Choose a device and click **Configure** in the quick navigation menu.
   b. Select an available device from the **Home** screen, and then click **Configure** under the available options.
3. Under **Controller Devices**, click **Logical Devices**.
4. Under **Action** panel, click **Move Logical Drive**.

**Figure 6-17.** Move Logical Drive




5. Under **Create a New Array**, select the logical drive. A new window appears.

**Figure 6-18.** Move Logical Drive—Select the Drive




6. Click **Yes**.
7. Click **Finish**.

## 6.19. Encryption Manager

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

For any compliance issues arising from your operation/usage of encryption within the SmartRAID controller module which violates the above mentioned requirement, you shall bear all the liabilities wholly and solely. The SmartRAID controller vendor is not responsible for any related liabilities.

For data encryption, you need to enable Controller Based Encryption (CBE) or Self-Encrypting Drive (SED) in Encryption Manager. For more information on configuring CBE or SED in Encryption Manager, see the Configuration section in the Data Encryption chapter.

## 6.20. Power Modes

The following three power modes are available:

- Maximum performance
- Minimum power
- Balanced

**Maximum Performance (Default)**

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

**Balanced**

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

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

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

**Minimum Power**

When settings are selected without regard to system performance, maximum power savings is achieved. It is recommended that to use this setting for specific applications, but it is not appropriate for most customers. Most applications suffer significant performance reduction.

**Important:**

- A reboot might be required after switching power modes to optimize savings and performance.
- When the Power mode is set to Balanced, future controller configuration changes may require a reboot for optimal performance.

### 6.20.1. Modifying Power Modes

To modify the power modes, perform the following steps:

1. Open SSA. For more information, see Accessing SSA.
2. Select the controller.
3. Click **Configure**.
4. Click **Manage Power Settings**.
5. Select a power mode:
   - Min Power
   - Balanced
   - Max Performance

**Figure 6-19.** Power Modes Window




6. **Optional:** Enable the **Survival Mode** only if required.
   **Note:** Enabling the Survival mode may decrease the controller's performance.
7. Click **OK**.
8. A summary page appears. Click **Finish** to exit.

## 6.21. Viewing Controller Status

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

To view controller status, perform the following steps:

1. Open SSA. For more information, see Accessing SSA.
2. Select a controller.
3. Click **Configure**.
4. Under **Controller Configuration Summary**, click **View more details.**
   A new window appears.
5. To view the status of the controller, cache, and energy pack, scroll down to **Controller Status**.

## 6.22. 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.

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

- Sanitize Overwrite – (HDD only)— It fills every physical sector of the drive with a pattern. If you enable this option, you can specify a sanitize method. If you select the restricted option, the drive is unavailable until the sanitize operation is completed successfully. If you select the unrestricted option, the drive is recoverable if the operation fails. For more information, see Sanitize Erase Methods.
- Sanitize Block Erase – (SSD only)— It sets the blocks on the drive to a vendor-specific value, removing all user data. If you enable this option, you can specify a sanitize method. If you select the restricted option, the drive is unavailable until the sanitize operation is completed successfully. If you select the unrestricted option, the drive is recoverable if the operation is failed. For more information, see Sanitize Erase Methods.
- Three-pass erase —This 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 —This 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 —This 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 continues to sanitize after a hot-plug or server reboot.
- During the sanitize erase operation, the drive is unusable until after the process is completed.
- 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 cancels the erase process.
- A physical drive remains offline after the erase process has completed if the one-, two-, or three-pass erase pattern is used. To bring the drive online and make it available for configuration, select the drive and click the **Enable Erased Drive** button.
- If the drive reboots during the one-, two-, or three-pass erase, the erase process may not persist following the reboot.

### 6.22.1. Sanitize Erase Methods

When you select a Sanitize Erase method, 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 is 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 us. This method may take several hours.
- Unrestricted: Using the unrestricted sanitize method means that the drive is 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.

The following are some considerations for the preceding sanitize erase methods:

- 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 takes to complete the sanitize erase.

  **Note:** 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.

## 6.23. Erasing a Drive

To erase a drive, perform the following steps:

1. Open SSA. For more information, see Accessing SSA.
2. Open the **Configure** panel by choosing a device and clicking **Configure** in the **Actions** panel.
3. In the **Configure** panel, click **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 method 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 with the drive being erased message.
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.

## 6.24. Managing Flexible Latency Scheduler (FLS)

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

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

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

In general, high settings result in lower maximum latencies for hard drive logical volumes, but lower throughput for patterns that have improved latency.
FLS is only supported by hardware RAID-based controllers. The following table lists the FLS settings.

**Table 6-2.** FLS Settings

| Setting | Description |
|---|---|
| Disable | Controller allows the drives to optimize for throughput, resulting in higher maximum latencies for some workloads. |
| High, Very High | Controller attempts to compel disks to complete operations in a timely manner, but still allow for some drive optimization. Doing so lowers the effective throughput for some workloads. |
| Low, Medium | Controller compels 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 might result in a loss of throughput for some request patterns.

### 6.24.1. Enabling FLS

To enable FLS, perform the following steps:

1. Open SSA. For more information, see Accessing SSA.
2. Select the controller.
3. Open the **Configure** panel by performing one of the following:
   - Select a device in the **Home** screen and click **Configure** in the quick navigation menu.

- Click on the drop-down next to the SSA icon. Hover over the device you want to select and click **Configure**.

4. Click **Advanced Controller Settings** in the **Configure** panel.
5. Scroll down to the option HDD Flexible Latency Optimization.
6. Select one of the following:
   - Disabled
   - Low
   - Middle (100 ms)
   - High
   - Very High (30 ms)
   - Very High (10 ms)
7. Click **OK**.
8. A summary page appears. Click **Finish** to exit.

## 6.25. Configure the Language of the GUI

The following sections describe the process to configure the GUI language.

### 6.25.1. Configuring the GUI Language (Windows)

To configure the GUI language, perform the following step:

1. In the top panel, click on .
   The list of languages appear.
2. Select your preferred language from the list.

### 6.25.2. Configuring the GUI Language (Linux)

To configure the GUI language (Linux), perform the following steps:

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

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

2. Specify the language using the appropriate two-letter language code:

```
ssa -local -lang <languageCode>
```

   Alternately, set the language using the settings  button in the top panel.

## 6.26. Maintaining Controller Firmware

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

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

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

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

### 6.26.1. Checking Online Firmware Activation Readiness

To check online firmware activation readiness, perform the following steps:

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

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

   If your controller is not enabled, a list of reasons is displayed.

# 7. Managing Diagnostics

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

For either of the tasks, you can save the report. In online and offline environments, 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 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

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
- OS device names or mount points

## 7.1. Performing a Diagnostics Task Using SSA

To perform a diagnostics task using SSA, perform the following steps:

1. Open SSA.
2. Open the **Diagnostics** panel by performing one of the following:
   - Select a device in the **Home** screen and click **Diagnose** in the quick navigation menu.
   - Click on the drop-down next to the SSA icon. Hover over the device you want to select and click **Diagnose**.
3. Select a report type.

For this example, use the Array Diagnostic Report selection.

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

**Related Links**

Reported Information
Identifying and Viewing Diagnostic Report Files
Identifying and Viewing SmartSSD Wear Gauge Report Files

# 8. Data Encryption

SSA features an Encryption Manager that protects the data at rest on bulk storage hard drives and SSDs attached to a SmartRAID controller. SmartRAID Secure Encryption supports the following encryption modes:

- Controller Based Encryption—CBE mode
- Self-Encrypting Drive—SED mode

CBE is an enterprise-class data encryption solution that protects data at rest on any SAS/SATA/NVMe drive configured as a member of a RAID volume. This solution is available for both local and remote deployments and works in conjunction with SmartRAID SmartCache. CBE supports both Local and Remote mode of operation.

SED based encryption provides the capability to manage all the SEDs connected to the controller by using an active pin called Master Key. The controller keeps Master Key in NVRAM. When the controller boots, the Master Key is read from NVRAM and applied to all the attached controller owned SEDs. SED based encryption supports both enterprise class and Opal encryption solutions. SED based encryption is available only for Local mode deployment and works in conjunction with SmartRAID SmartCache.

## 8.1. About CBE

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

CBE provides encryption for data at rest as an important component for complying with sensitive data protection requirements including PCI-DSS, HIPAA/HITECH, Sarbanes/Oxley, and state privacy laws. CBE secures any data deemed sensitive and requiring extra levels of protection through the application of XTS-AES 256-bit data encryption. Many companies under the government regulations require that sensitive privacy data be secured and uncompromised using NIST-approved algorithms and methodologies for local key management. Secure Encryption is validated for FIPS-140-2 Level 2 for the SmartRAID Px3x controllers and is validated for FIPS 140-2 Level 1 for the SmartRAID Px4x controllers and SmartRAID Gen10 P-Class RAID controllers. For more information about the controllers that are validated, see the *Cryptographic Module Validation Program (CMVP)* on the National Institute of Standards and Technology website.

CBE operates in Remote Key Management mode, or Remote mode, through the use of a separate, clustered, appliance-based server called the Enterprise Secure Key Manager (ESKM). The ESKM manages all encryption keys throughout the data center. When utilizing the ESKM, the communication path between the ESKM and the SmartRAID Controller is established through the iLO interface. The controller communicates with the ESKM as new keys are generated and old keys are retired. The ESKM acts as a key vault where all keys are managed through a web browser interface. For more information about the ESKM, see Enterprise Secure Key Manager. For more information about iLO connectivity, see iLO.

The following additional components are required for operating Secure Encryption in Remote mode:

- Integrated Lights Out (iLO) Advanced License, per ProLiant server
- ESKM

CBE can also operate without an attached key management solution through Local Key Management mode, or Local mode.

### 8.1.1. Benefits

This section describes the benefits of CBE.

- **Broad Encryption Coverage**
  - Encrypts the data on both the attached bulk storage and the cache memory of SmartRAID Controllers
  - Supports any hard drive or SSD for ProLiant Gen8 or later servers or the Supported Storage Enclosures
- **High Availability and Scalability**
  - Scales to meet individual data privacy requirements:
    - Server counts up to 25,000
    - Millions of drives
    - Millions of encryption keys
  - The ESKM supports High Availability Clustering.
- **Simplified Deployment and Management**
  SSA configures the cryptographic features of the Secure Encryption and manages the controller and other direct-attached storage devices.

### 8.1.2. Solution Components

This section describes the solution components for CBE.

#### 8.1.2.1. SmartRAID Controller

CBE is supported on the following:

- HPE Gen10 Plus and Gen11 controllers
- SmartRAID Gen10 E-class (For example, E208i-a) and P-class (For example, P408i-a) controllers
- SmartRAID PX3X and PX4X controllers
- HX4X SmartHBAs operating in RAID mode

#### 8.1.2.2. Encryption Features

Most CBE features and security settings are available through SSA. Additional features for Remote mode deployments are available through ESKM and iLO.
The following table lists the encrytion features and security settings available through ESKM.

**Table 8-1.** Encryption Features and Security Settings

| Feature | Description | Notes |
|---|---|---|
| Automatic Key Management | SmartRAID Controllers automatically creates, saves, and deletes the Encryption Keys, without the need for user intervention or management when logical drives are created or deleted. | — |
| Compliance | CBE is designed to meet NIST-approved standards. The ESKM has completed FIPS 140-2 Level-2 validation (certificate #1922). CBE helps enterprises to comply with the data privacy and protection requirements associated with the U.S. Health Insurance Portability and Accountability Act (HIPAA) and the Sarbanes-Oxley Acts. | — |
| Controller Key Cache | SmartRAID Controllers optionally stores all keys required at boot time inside the controller, enabling the server to survive a variety of network outages. | Remote mode only |
| Controller Password | Protects the server in the event of theft by applying a secondary password upon boot to lock down the controller. | For more information, see Set or Change the Controller Password. |
| Dynamic Encryption | Enables smooth transitions between Local and Remote modes, the conversion of plaintext data to encrypted data, and rekey services for both the data and key wraps. | — |
| Encryption Keys | Data is protected using a series of keys that provide layered protection at the volume and drive levels. The solution utilizes XTS-AES 256-bit encryption. | — |

**Table 8-1.** Encryption Features and Security Settings (continued)

| Feature | Description | Notes |
|---|---|---|
| Enterprise Secure Key Manager | The ESKM unifies and automates an organization's encryption controls by securely creating, protecting, serving, controlling, and auditing access to encryption keys. | Remote mode only. For more information, see Enterprise Secure Key Manager. |
| ESKM Key Search | Individual Drive Encryption Keys are visible by serial number identification on the ESKM to enable unique tracking and management from a central location. The ESKM supports query by serial number, server name, bay number, PCI slot, and date. | Remote mode only. For more information, see the HPE Enterprise Secure Key Manager - Document List in the Resources. |
| Firmware Lock | Prevents controller firmware from being updated unintentionally or by unauthorized personnel. | For more information, see Enabling/Disabling the Firmware Lock. |
| Hardware-based Encryption | Utilizes the SmartRAID Controller hardware to accelerate all cryptographic algorithms when securing data and keys. | For more information about SmartRAID Controllers, see the website. |
| iLO | iLO Management is a comprehensive set of embedded management features supporting the complete life cycle of the server, from initial deployment, through ongoing management, to service alerting and remote support. iLO is provided on all ProLiant Gen8 and later servers.<br>iLO 4 Advanced License editions v1.40 or later, connect and auto-register with the ESKM. iLO provides key exchange support between the SmartRAID Controller and the ESKM to enable preboot support for OS disk encryption. All key management transactions receive audit support. | Remote mode only. For more information, see iLO. |
| Instant Volume Erase | Provides ability to instantly and cryptographically erase logical volumes without having to delete the volume first. | — |
| Key Rotation Support | Supports the rekeying of all keys utilized by the controller to enable a robust key rotation strategy. | — |
| Local Key Management Mode | Focused on single server deployments where there is one Master Encryption Key per controller that is managed by the user. In Local mode, all volumes still have their own unique key for data encryption. | For more information, see Local Key Management Mode. |
| One-way Encryption | As a security feature, data volumes cannot be converted back to plaintext after the volume is encrypted. Restoration of data is required to revert to plaintext. | — |
| Pre-deployment Support | Supports the ability to preconfigure all cryptographic security settings while in a server, then store the powered-off controller for later use while retaining the settings securely. | — |
| Remote Key Management Mode | Designed for enterprise-wide deployments with the SmartRAID Controller. It requires ESKM to manage all keys related to encryption deployments. All keys are managed automatically between the SmartRAID Controller, iLO, and the ESKM. | For more information, see Remote Key Management Mode. |
| Security Reset Function | The feature clears all secrets, keys, and passwords from the controller, and places the encryption configuration of the controller in a factory new state. | For more information, see Clearing the Encryption Configuration. |
| Smart Storage Administrator | Smart Storage Administrator v1.60.xx.0 and later provides the configuration and management of the cryptographic features of Secured Encryption associated with SmartRAID Controllers. | For more information, see Introduction. |
| Two Encryption Roles | Secured Encryption supports two roles for managing encryption services: a Crypto Officer role and a User role. | — |
| Volume Level Encryption | Provides flexibility in allowing the user to selectively encrypt at the volume or logical drive level regardless of RAID level. | — |

#### 8.1.2.3. SmartRAID SmartCache Secure Encryption

SmartRAID SmartCache can also be used in conjunction with Secure Encryption. SmartRAID SmartCache enables SSDs to be used as caching devices for hard drive media. Data is also accessed

from the SSDs instead of hard drives. Data stored on the SmartCache drive utilizes the same encryption methods and keys as the originating volume where the data is permanently stored, extending protection to the SmartCache drives.

SmartCache provides the following features:

- Accelerates application performance
- Provides lower latency for transactions in applications
- Supports all OS, without the need for change

#### 8.1.2.4. iLO

iLO Management is a set of embedded management features that support the complete life cycle of the server, from initial deployment, to ongoing management, and to service alerting and remote support.

The iLO subsystem is a standard component of servers that simplifies initial server setup, server health monitoring, power and thermal optimization, remote server administration, and key exchanges between the ESKM and the SmartRAID Controller. The iLO subsystem includes an intelligent microprocessor, secure memory, and a dedicated network interface. This design makes iLO independent of the host server and its OS. This system provides client credentials, registration to the key management database, the key management, the encryption activation, and the audit support for the devices within the platform.

For the complete implementation of Secure Encryption with the ESKM, iLO Advanced License editions are required to connect and auto-register with the ESKM. iLO provides key exchange support between the SmartRAID Controller and the ESKM to enable pre-boot support for OS disk encryption. Audit support is provided for all for the key management transactions.

#### 8.1.2.5. Enterprise Secure Key Manager

ESKM acts as a secure, reliable repository for keys used by Secure Encryption. In Remote Key Management mode, iLO connects to the ESKM using the username/password and digital certificate authentication to securely store and retrieve keys. Each iLO must be registered as an ESKM user by an administrator, or Crypto Officer, of the ESKM for access to be granted. If a user is registered and has the necessary permissions, the ESKM accepts requests and provides keys to the client. As a standard practice, communication with the ESKM is configured for SSL to ensure the security of the connection and authorized access to keys.

The ESKM keys and users are organized into different groups depending on the policies set by an administrator. These groups determine whether a particular user retrieves a particular key, and supports both key sharing and separation for multi-tenant and hosted service provider environments.

**Characteristics**

- Used only in Remote mode, requiring a network connection
- Supports high-availability clustering of ESKM nodes for automatic replication and failover
- Provides key services to iLO clients using the username and the password, certificate authentication, or both
- Communicates using SSL encryption to ensure the security of the connection and authorized access to keys
- Provides reliable, secure access to business-critical encryption keys
- Supports audit and compliance requirements, including PCI-DSS and HIPAA/HITECH
- Provides scalability for multiple data centers, thousands of clients, and millions of keys
- Uses a FIPS-140-2 Level 2 validated secure appliance, which supports the latest NIST cryptographic guidance. Following are the ESKM vendors partnered with HPE:

- Thales: cpl.thalesgroup.com/partners/hpe
- Utimaco: utimaco.com/partners/hpe

#### 8.1.2.5.1. ESKM and Key Management

The SmartRAID controller manages keys by separating them into the following categories:

- Keys stored off-controller on the ESKM
- Keys stored on the drive media
- Keys stored on the controller

The separation of keys ensures the safety of the data residing on the drives, the portability of the drives, and the ability to manage keys in a centralized manner. The controller uses the ESKM to back up a segment of its keys using an encryption method that protects the keys from exposure in plaintext.

## 8.2. About SED

A SED encrypts data through disk-based encryption with a Media Encryption Key (MEK). The MEK is known only to the SED and cannot be recovered through forensic analysis. Smart controllers enable the use of SEDs as logical drives or physical drives.

The controller is responsible for managing and delivering the credentials required by the SED for enabling the disk-based encryption. SAS, SATA, and NVME drives that are compliant to the Opal 2.0 and Enterprise 1.01 industry standards are supported.

## 8.3. Planning

This section describes the encryption setup guidelines, security settings at remote sites, encrypted backups, security domains, and deployment scenarios.

### 8.3.1. Encryption Setup Guidelines for CBE

This section describes the CBE setup guidelines for SmartRAID. When setting up CBE, consider the information described in the following table.

**Table 8-2.** Encryption Setup Guidelines

| Configuration | | Description |
|---|---|---|
| Encryption mode | • Local Key Management Mode<br>• Remote Key Management Mode | Choose Local Key Management mode when:<br>• Data is stored at a site without network access<br>• In a small deployment center or lab<br>• Manual key management is available<br>Choose Remote Key Management mode when:<br>• Using a large number of servers<br>• A network is available between the ESKM and a server<br>• Automatic key management is preferred, including backups and redundancy configurations |
| Plaintext volumes | • Allow<br>• Disallow (default) | Allow future plaintext logical drives when:<br>• Drive migration might occur to a non-encrypting controller.<br>• Data is not privacy-sensitive. For more information, see Enabling/Disabling Plaintext Volumes |
| Key naming conventions | Master Encryption Keys are customizable. | Create a specific naming convention when managing multiple keys and multiple servers. |

#### 8.3.1.1. Recommended Security Settings at Remote Sites

For added security, it is recommended to use the following configuration when operating Secure Encryption at remote sites outside the main data center.

- Firmware lock enabled. For more information, see Enabling/Disabling the Firmware Lock.
- Controller password enabled. For more information, see Set or Change the Controller Password.
- Plaintext volumes disabled. For more information, see Enabling/Disabling Plaintext Volumes.
- Local Key Cache disabled. For more information, see Enabling/Disabling Local Key Cache.
  **Note:** Applies to Remote Key Management mode only.

#### 8.3.1.2. Encrypted Backups

At system startup, all encrypted data-at-rest becomes accessible to the host system in an unencrypted form through the controller and the appropriate keys. This method of startup allows the system to boot into an OS installed on an encrypted volume. As a result, encrypted backups are not available, and all data appears unencrypted when accessed from the host system and placed on tape. Secure Encryption does not impact software or hardware, which utilizes an independent encryption feature.

#### 8.3.1.3. Security Domains

A security domain is a blueprint for separating out different groups of servers or key management escrows where access to a set of keys is inhibited by the structure of the various domains. The best mechanisms for establishing separate security domains are either through using a separate ESKM or by using groups within the ESKM. Unique groups provide a software mechanism for each server to divide their key sets from one server to another. Groups are created on the ESKM and are assigned to a server through the iLO Key Manager page. For more information, see Remote Key Management Mode.

#### 8.3.1.4. Deployment Scenarios

This section describes the deployment scenarios for Remote and Local Key Management modes.

##### 8.3.1.4.1. Remote and Local Key Management Requirements

This section describes the remote and local key management requirements. Use the following table to determine appropriate encryption mode.

**Table 8-3.** Remote and Local Key Management Requirements

| Mode parameters | Local Key Management Mode | Remote Key Management Mode |
|---|---|---|
| Number of servers | <99 (recommended) | 100 or more |
| ESKM available | No | Yes |
| iLO Advanced License available | No | Yes |
| Requirement to escrow keys | No | Yes |
| Manual tracking of keys | Yes | No |

### 8.3.2. Encryption Setup Guidelines for SED

This section describes the encryption setup guidelines for SED. When setting up SED based encryption, consider the information described in the following table.

**Table 8-4.** Encryption Setup Guidelines

| Configuration | Description |
|---|---|
| SED Management | Shows SED Based Encryption status and allows users to disable support |
| Key Management Mode | Supports both Local and Remote Key Manager modes. In the Local mode, you must provide the master key at the time of setup. In the Remote mode, the master key is not required. A system reboot is required to apply the changes. Option to perform rekey, and option to configure controller password are only available in UEFI HII. |

**Table 8-4.** Encryption Setup Guidelines (continued)

| Configuration | Description |
|---|---|
| Master Key | When the controller is managing SEDs, a Master Key is created during the initial setup. The Master Key is required to secure the SEDs and unlock the user data on managed SEDs.<br>**Local Key Management**<br>The Master Key is stored locally in the controller NVRAM. Optionally, a Master Key Identifier can also be entered at the time of Master Key creation.<br>**Remote Key Management**<br>The Master Key is generated and stored by the key management server external to the controller. The controller communicates with the server to retrieve the Master Key.<br>**Note:** If the Master Key is lost, the data on the encrypted volumes cannot be retrieved. |
| Master Key Identifier | **Local Key Management**<br>Master key identifier is a hint that helps you to remember the master key. You may use the "Use Default" button to create a default key identifier.<br>**Remote Key Management**<br>UEFI HII utilizes the auto generated Master Key and Master Key ID from the key management server. |
| Controller Password | The controller password is an optional setting while configuring the controller managed SED encryption.<br>**Local Key Management**<br>The controller password is intended to provide an extra level of security for local SED management and guards against theft of the server, adapter, and the SEDs. The adapter will not unlock any SED until the controller password input is provided in the configuration utility.<br>**Remote Key Management**<br>Controller password in the Remote Key Management mode serves as a backup option to unlock controller and encrypted devices in the case when the key management server becomes unavailable. Controller password option is only provided in UEFI HII utility. If you configure the controller password, an encrypted version of the Remote Master Key is stored in the controller NVRAM. If the controller is not able to connect to the remote key server, you can use the controller password to retrieve and decrypt the Remote Master key from controller NVRAM. |
| SED Ownership | • In order for SED Based Encryption to manage SED drives, the controller must take the ownership of the drive. The drive must be in original factory state before ownership can be taken.<br>• When a SED is owned by the controller, it can be reverted to Original Factory State. This process destroys all user data stored in the drive.<br>• Use Import Foreign SED operation to take ownership of a SED if it is foreign owned. You will enter the SED's foreign master key. A reset foreign master key may be required when the foreign drive was removed during a rekey process. |

## 8.4. Configuration

For data encryption, you need to enable Controller Based Encryption (CBE) or Self-Encrypting Drive (SED) in Encryption Manager. For more information, see CBE Setup and SED Based Encryption Setup.

### 8.4.1. CBE Setup

CBE is an enterprise-class data encryption solution that protects data at rest on any SAS/SATA/NVMe drive configured as a member of a RAID volume. The solution is available for both local and remote deployments. CBE is configured using the SSA. The following sections describe configuring the Local and the Remote Key Management modes in the CBE setup.

#### 8.4.1.1. Local Key Management Mode

Local Key Management mode, or Local mode, is a solution designed for small to medium-size data centers using few encrypting controllers. The solution utilizes a paraphrase password, or Master Encryption Key name, to set the security on the controller and enable encryption. The Master Encryption Key must be tracked independently of the controllers in case the controller needs replacement or drive migration is required among controllers with different passwords. In Local mode, the Master Key name is considered as a cryptographic secret and must be protected as such. Key creation and management are maintained at the local controller level without using a key manager.

**Characteristics**

Following are the characteristics of Local Key Management mode:

- Requires physical paraphrase password management, such as writing and storing Master Key information in a notebook or computer file.
- Utilizes one paraphrase password-derived 256-bit key to encrypt a unique, per-volume XTS-AES 256-bit data encryption key.

##### 8.4.1.1.1. Opening Encryption Manager

To open Encryption Manager, perform the following steps:

1. Start SSA.
2. Select a Secure Encryption-compatible controller.
3. Click **Configure**.
4. Under **Tools**, click **Encryption Manager**.

##### 8.4.1.1.2. Configuring the Controller to Enable CBE (Local Key Encryption)

To configure Secure Encryption using command-line or scripting methods, see *Smart Storage Administrator CLI Guide*.

**Important:** It is recommended that you keep a record of the Master Encryption Keys when encryption is configured in Local mode. The local Master Encryption Key is not displayed by any available tool or firmware because it is considered a cryptographic secret by FIPS 140-2. Secure Encryption design follows the NIST architecture requirements and does not allow us to assist in the recovery of a lost Master Encryption Key.

To configure the controller to operate in Local Key Management mode, perform the following steps:

1. Open Encryption Manager. For more information, see Opening Encryption Manager
2. Click **Controller-Based Encryption (CBE) Setup**.

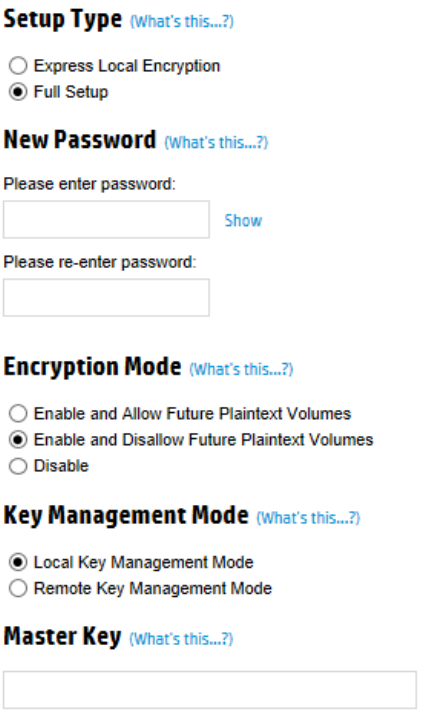
**Figure 8-1.** Encryption Manager—CBE



The following CBE Setup screen appears.

**Note:** "Setup Type" is unavailable in "Perform Initial Setup" if encrypted logical drives are attached to a replacement controller.

**Figure 8-2.** CBE Setup Selection Screen




3. Complete the following:
   - Under **Setup Type**, select **Full Setup**.
   - Enter your password in the **New Password** section.
   - Under **Encryption Mode** section, select either:
     - **Enable and Allow Future Plaintext Volumes**: Allowing future plaintext volumes still requires authentication by Crypto Officer or the user before a plaintext volume is created.
     - **Enable and Disallow Future Plaintext Volumes**: This option prevents the creation of new plaintext volumes on the controller. This setting is changed later by the Crypto Officer. Selecting this option does not prevent the migration of a set of drives with existing plaintext volumes to the controller.
   - Under **Key Management Mode**, select **Local Key Management Mode**.
   - Type the **Master Key** name in the field provided. The Master Encryption Key name must be between 10 and 64 characters.
4. Click **OK**.
5. A warning appears, prompting the user to record the Master Encryption Key. Click **Yes** to continue.
6. If you have read and agreed to the terms of the EULA, select the check box and click **Accept**. A summary screen appears.
7. Click **Finish**.

**Important:** It is recommended to set up a password recovery question and answer after the initial configuration. If the Crypto Officer password is lost and a recovery question and answer have not been set, you need to erase and reconfigure all Secure Encryption settings to reset the Crypto Officer password. For more information, see Set or Change the Password Recovery Question.

#### 8.4.1.1.3. Express Local Encryption

This section describes the Express Local Encryption.

**About Express Local Encryption**

**Important:** Express Local Encryption configures Secure Encryption in Local Key Management mode. Once configured, Crypto Officer password is not available.

**Important:** Express Local Encryption uses a randomly-generated Master Encryption Key. Features requiring the input of a Master Encryption Key, such as migrating volumes to a new controller is not available while Express Local Encryption is enabled.

Express Local Encryption configures the controller with predetermined encryption settings and a randomly-generated Master Encryption Key. Once configured, you cannot change the encryption settings without clearing the encryption configuration.

Express Local Encryption enables the following:

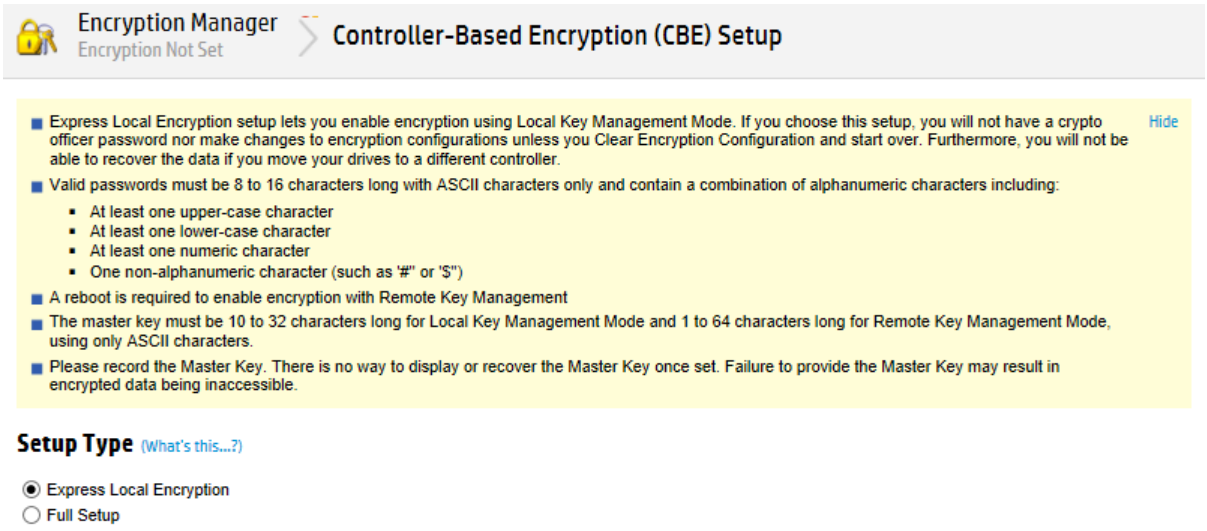
- Controller encryption
- Local Key Manager mode
- Random crypto password, not recoverable
- Random master key name, not recoverable
- Future plaintext volumes not allowed

**Setting up Express Local Encryption**

To set up express local encryption, perform the following steps:

1. Open Encryption Manager. For more information, see Opening Encryption Manager .
2. Click **Perform Initial Setup**, a new window appears.

**Figure 8-3.** CBE Initial Setup Screen




3. Under **Setup Type**, select **Express Local Encryption**. Once selected, all other encryption setup options disappear. Click **OK** to continue.
4. Click **Yes** to continue when the warning message is displayed.
5. If you have read and agreed to the terms of the EULA, select the check box and click **Accept**.
6. The **Encryption Manager** screen appears with updated **Settings**, **Accounts**, and **Utilities** options.

### 8.4.1.2. Remote Key Management Mode

**Important:**

- ESKM must already be installed and configured to operate Secure Encryption in Remote mode. For more information, see Configuring the ESKM.

In Remote Key Management mode, keys are imported and exported between the controller and the ESKM, which provides a redundant, secure store with continuous access to the keys. To enable key exchanges between the SmartRAID Controller and the ESKM, a network connection is required both during pre-OS boot time and during OS operations. Because the controller does not have direct network access, iLO provides the necessary network access to facilitate key exchanges between the controller and the ESKM. iLO has both network presence and is constantly running on AUX power regardless of the server state. The keys that are exchanged between iLO, ESKM, and the controller are all secured.

A valid Secure Encryption license for each server to be encrypted is required. This license must be purchased, but it does not need to be input into SSA.

#### Characteristics

Following are the characteristics of the Remote Key Management mode:

- High volume key storage
- Keys are kept in separate storage from servers to protect against physical removal
- Requires network availability and a remote key management system

#### 8.4.1.2.1. Configuring Remote Key Management Mode

**Important:** Secure Encryption and other encryption client products must be coordinated for a successful installation and configuration. It is recommended to refer to each product's user guide to ensure proper installation and encryption protection.

To configure Secure Encryption to operate in Remote mode, perform the following steps:

1. Configure the ESKM. For more information about installation, configuration, and operation of the ESKM, see the HPE Enterprise Secure Key Manager - Document List in the Resources.
2. Connect iLO to the ESKM.
3. Install SSA.
4. Configure the SmartRAID Controller.

**Configuring the ESKM**

To configure the ESKM, perform the following steps.
**Note:** For more information about step details, see the HPE Enterprise Secure Key Manager - Document List in the Resources.

1. Log in to the ESKM.
2. Create initial user accounts.
3. Create a group.
4. Assign the user account for hosting Master Encryption Keys to the group created in step 3.
5. Create a Master Encryption Key to be used by the controller. You must set the owner of the key to the user account created to host the Master Encryption Key created in step 2.
6. Place the Master Encryption Key in the group created in step 3.

## Configuring iLO

iLO manages key exchanges between the ESKM and the SmartRAID Controller. iLO initially uses user credentials with administrative privileges created on the ESKM to automatically register and create a private, unique, and MAC address-based username account for all key exchanges. The administrative account is termed the deployment user account. All iLO accounts can be viewed in the ESKM under Users And Groups and take the form iLO-MAC Address. The iLO-specific account is placed in the group indicated in the group field on the **iLO Key Manager** page. If the group does not exist, iLO creates one and places the account in that group along with all future keys generated.

**Prerequisites**

- The ESKM must be configured with a deployment user. For more information, see Configuring the ESKM.
- iLO must be installed and operating properly with the appropriate iLO-supporting license.

For more information on installing and configuring iLO, including scripting and command line methods, see the Hewlett Packard Enterprise website.

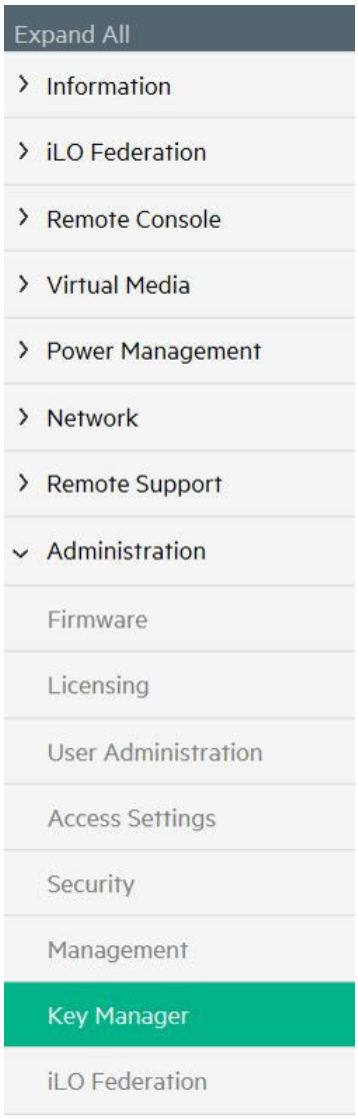
## Connecting iLO to the ESKM

If you intend to use a second ESKM for a redundant key repository, complete the fields under **Secondary Key Server** and select the **Enable Enterprise Secure Key Manager Redundancy** check box. It is recommended that there should be a redundant pair of ESKM devices in a cluster configuration.

To connect iLO to the ESKM, perform the following steps:

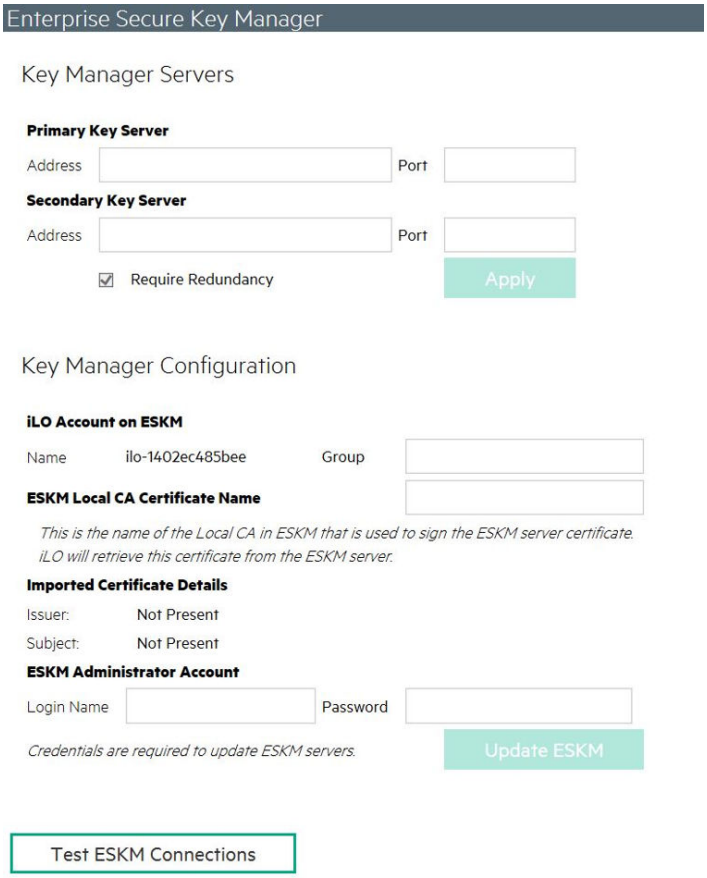
1. Log in to iLO using the server credentials.
2. From the left panel, expand the **Administration** menu and select **Key Manager**.

**Figure 8-4.** Administration Menu—Key Manager



The **Enterprise Secure Key Manager** configuration window appears.

**Figure 8-5.** Enterprise Secure Key Manager




3. Under **Key Manager Servers**, complete the following:
   a) **Primary Key Server**
      - Type the primary IP address of the ESKM in the **Address** box.
      - Type the primary port number of the ESKM in the **Port** box. This port number must match the value on the ESKM, located under **KMS Server Settings** on the **Device** tab. SSL must be enabled on the ESKM as well.

   b) Optional: **Secondary Key Server**
      - Type the secondary IP address of the ESKM in the **Address** box
      - Type the secondary port number of the ESKM in the **Port** box.

   c) Optional: Select the **Require Redundancy** check box. This option enables iLO to verify that encryption keys are copied to all configured key servers. For configurations with a primary and secondary key server. It is recommended to enable this option.
4. Click **Apply**. A confirmation message appears.
5. Under **Key Manager Configuration**, type the group name created previously in the ESKM in the **Group** field.
6. Under **ESKM Administrator Account**, complete the following fields using the deployment username and password created earlier on the ESKM.
   a) **Login Name**: Type the deployment account username.
   b) **Password**: Type the deployment account password.

7. Click **Update ESKM**. A confirmation screen appears, indicating that the configuration is saved and connected successfully.

#### 8.4.1.2.2. Configuring the Controller to Enable CBE (Remote Key Encryption)

To configure the controller to operate in Remote Key Management mode, perform the following steps:

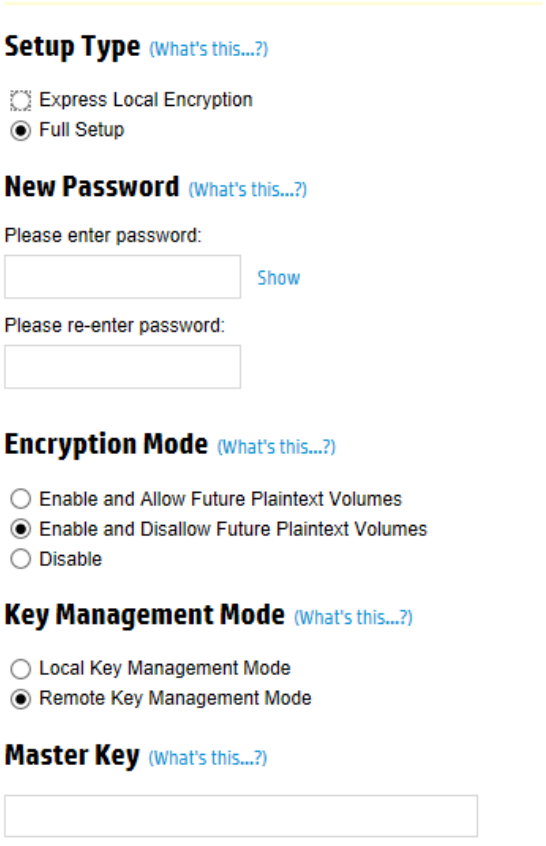
1. Open Encryption Manager. For more information, see Opening Encryption Manager .
2. Click **Controller-Based Encryption (CBE) Setup**.

**Figure 8-6.** Encryption Manager—CBE




The following CBE Setup screen appears.

**Figure 8-7.** Encryption Manager—CBE Setup Screen




3. Complete the following:

- Under **Setup Type**, click **Full Setup**.
- In the boxes under **New Password**, type a suitable password.
- Under **Encryption Mode**, select either:
  - **Enable and Allow Future Plaintext Volumes**: Allowing future plaintext volumes still requires authentication by the Crypto Officer or the User before a plaintext volume is created.
  - **Enable and Disallow Future Plaintext Volumes**: This option prevents the creation of new plaintext volumes on the controller. This setting is changed later by the Crypto Officer. Selecting this option does not prevent the migration of a set of drives with existing plaintext volumes to the controller.
- Under **Key Management Mode**, select **Remote Key Management Mode**.
- Type the **Master Key** name in the field provided. The Master Encryption Key name must be between 10 and 64 characters.

4. Click **OK**.
5. A warning appears, prompting the user to record the Master Encryption Key. Click **Yes** to continue.
6. If you have read and agreed to the terms of the EULA, select the check box and click **Accept**.
7. A summary screen appears, click **Finish**.

**Important:** It is recommended to set up a password recovery question and answer after the initial configuration. If the Crypto Officer password is lost and a recovery question and answer have not been set, you have to erase and reconfigure all the Secure Encryption settings in order to reset the Crypto Officer password. For more information, see Set or Change the Password Recovery Question.

#### 8.4.1.3. Changing from Local Key Management Mode to Remote Key Management Mode

To change the Key Management mode from Local Key Management mode to Remote Key Management mode, perform the following steps.

**Note:**
You do not need to back up or restore your data when performing this task.

1. Connect iLO to the ESKM. For more information, see Connecting iLO to the ESKM.
2. Create the Master Key on the ESKM.
3. Start SSA.
4. Under **Array controllers**, select the controller.
5. Under **Actions**, click **Configure**.
6. From the side menu, click **Encryption Manager**.
7. Log in to Encryption manager.
8. Change the Key Management mode from Local Key Management mode to Remote Key Management mode.
9. In the **Master Key** box, type the master key.

### 8.4.2. SED Based Encryption Setup

This section describes configuring the Local Key Management mode in the SED setup.

### 8.4.2.1. Configuring the Controller to Enable SED (Local Key Management Mode)

To configure the controller to operate in Local Key Management mode, perform the following steps:

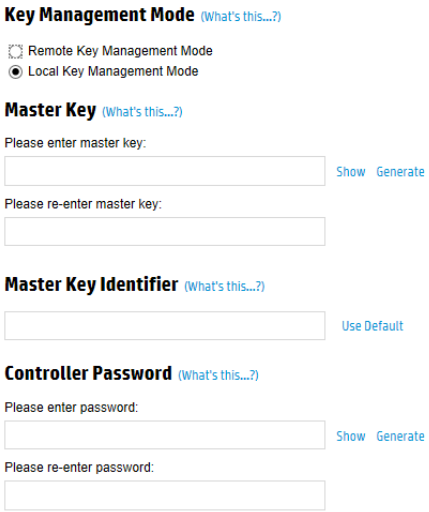
1. Open Encryption Manager.
2. Click **Self-Encrypting Drive (SED) Setup**.

**Figure 8-8.** Encryption Manager—SED




The following screen appears.

**Figure 8-9.** Encryption Manager—SED Setup




3. Complete the following:
   - In the boxes under the **Master Key**, type a suitable value. The Master Encryption Key name must be between 8-32 characters.
   - In the **Master Key Identifier** box, type a hint that helps you remember Master Key. You can use **Use Default** option.
   - In the boxes under **Controller Password**, type a suitable password.
   - Under **Key Management Mode**, it is **Local Key Management Mode**.
4. Click **OK**.
5. A warning appears, click **Yes** to continue.
6. The **Encryption Manager** screen appears with updated **Settings**.

### 8.4.2.2. Configuring the Controller to Enable SED (Remote Key Management Mode)

To configure the controller to operate in the Remote Key Management mode, perform the following steps:

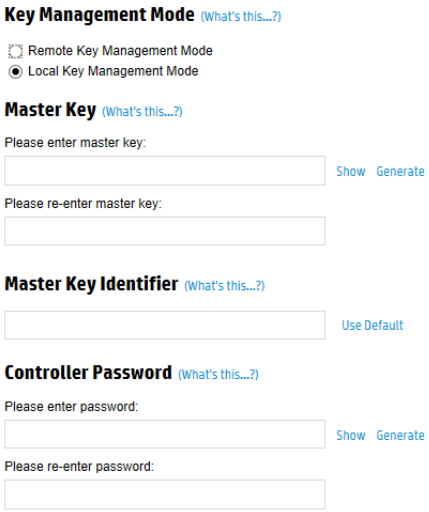
1. Open Encryption Manager.
2. Click **Self-Encrypting Drive (SED) Setup**.

**Figure 8-10.** Encryption Manager—SED



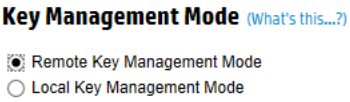

The following screen appears.

**Figure 8-11.** Encryption Manager—SED Setup




3. Under **Key Management Mode**, select **Remote Key Management Mode**.

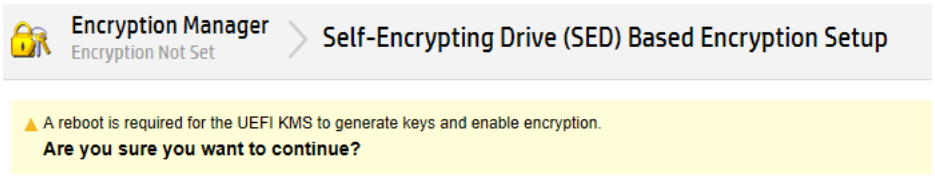
**Figure 8-12.** SED Based Encryption Setup—Key Management Mode




**Note:** The **Master Key**, **Master Key Identifier**, and **Controller Password** fields are disabled, use UEFI to set or change their values in the Remote Key Management mode.

4. Click **OK**.
5. A warning appears, prompting a reboot for the UEFI KMS to generate keys and enable encryption, click **Yes** to continue.

**Figure 8-13.** SED Based Encryption Setup—Warning Message




6. The **Encryption Manager** screen appears with updated **Settings**.

## 8.5. Operations

This section describes the operations to access encryption manager, manage passwords, work with keys, create a plaintext volume, convert plain text volumes to encrypted volumes, change key management modes, and enable or disable plain text volumes, the firm lock, and local key cache.

### 8.5.1. CBE Operations

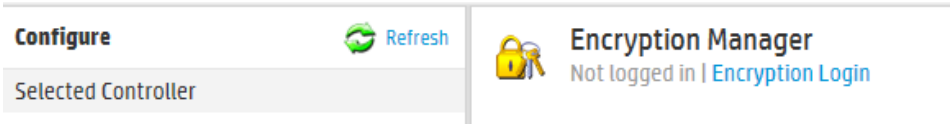
The following sections describe operations under CBE encryption setup.

#### 8.5.1.1. Logging into Encryption Manager (CBE)

To log into Encryption Manager in CBE, perform the following steps:

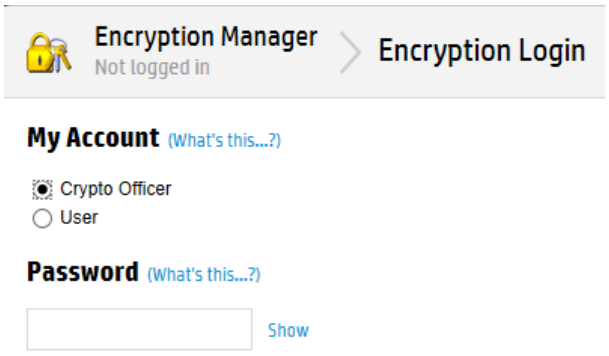
1. Open Encryption Manager.
2. Click **Encryption Login**. A new window appears.

**Figure 8-14.** Encryption Manager




3. Select an account to log in with and type the password in the field provided.

**Figure 8-15.** Encryption Login Window




4. Click **OK** to continue.

#### 8.5.1.2. Managing Passwords

**Notes:** Valid passwords must be 8 to 16 US-ASCII characters long and contain the following:

- At least one lowercase letter
- At least one uppercase letter
- At least one number
- At least one non-alphanumeric character, such as # or $

##### 8.5.1.2.1. Set or Change the Crypto Officer Password

To set or change the Crypto Officer password, perform the following steps:

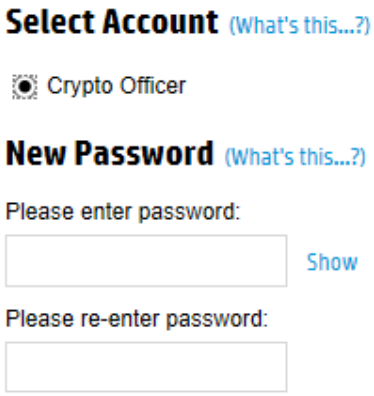
1. Open Encryption Manager.
2. Log in as the Crypto Officer.
3. Under **Accounts**, locate **Crypto Officer Password**, and then click **Set/Change Crypto Officer Password**. A new window appears.

**Figure 8-16.** Set/Change Crypto Officer Password




4. In the boxes under New Password, type a suitable password.

**Figure 8-17.** Crypto Officer—Enter Passoword




5. Click **OK**.

##### 8.5.1.2.2. Set or Change the Password Recovery Question

To set or change the password recovery question, perform the following steps:

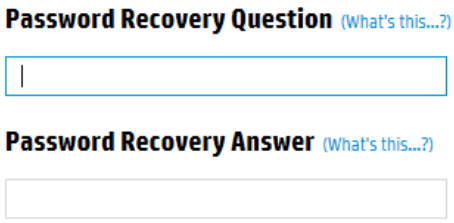
1. Open Encryption Manager.
2. Log in as the Crypto Officer.
3. Under **Accounts**, locate **Crypto Officer Password Recovery Parameters**, and then click **Set/Change Password Recovery Question**.

**Figure 8-18.** Crypto Officer Password Recovery Parameters




A new window appears.

**Figure 8-19.** Set Password Recovery




4. Complete the following fields:
   a. **Password Recovery Question**: Type a question to which only you know the answer.
   b. **Password Recovery Answer**: Type the answer to the question typed above.
5. Click **OK**.

### 8.5.1.2.3. Set or Change User Account Password

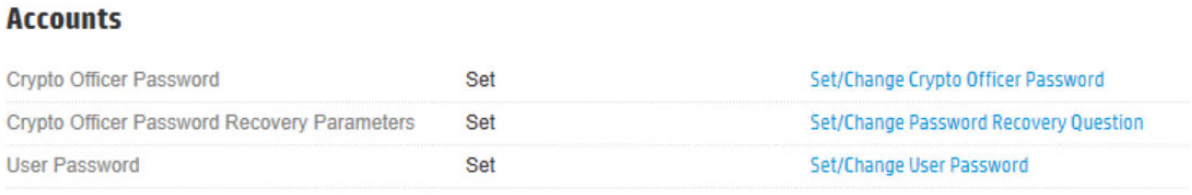
**Important:** If the User password is being set up for the 1st time, you must log in as the Crypto Officer.

The User account is disabled by default until the Crypto Officer sets the User account password for the first time.

To set or change the User account password, perform the following steps:

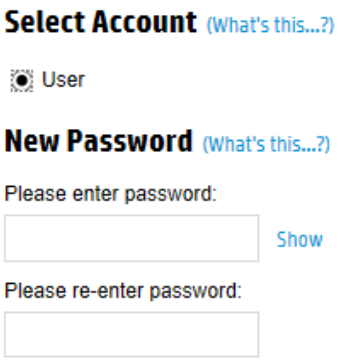
1. Open Encryption Manager.
2. Log in to the Encryption Manager.
3. Under **Accounts**, locate **User Password**, and then click **Set/Change User Password**. A new window appears.

**Figure 8-20.** Set/Change User Password




4. In the boxes under **New Password**, type a suitable password.

**Figure 8-21.** User—New Password Fields




5. Click **OK**.

### 8.5.1.2.4. Set or Change the Controller Password

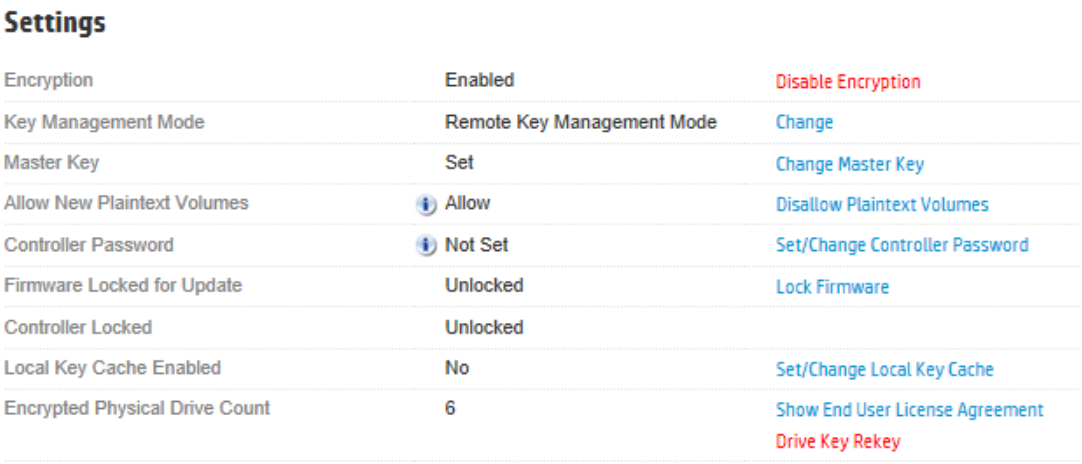
A controller password makes all the encrypted volumes on the controller to be kept offline at start-up until the controller password is entered.

The "Set/Change Password" action enables the controller password feature and sets the initial password. After a password is set, re-execute this action to replace the existing controller password with a new one. This procedure is performed only by the Crypto Officer. You cannot change the controller password while the controller password feature is suspended or while the controller is locked. However, the controller password can be removed by the Crypto Officer and later it can be reset.

To set or change the controller password, perform the following steps:

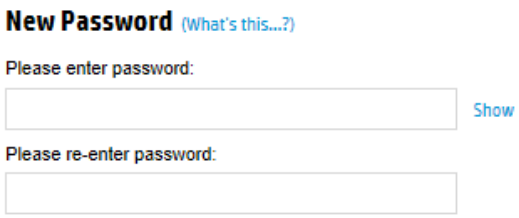
1. Open Encryption Manager.
2. Log in to Encryption Manager.
3. Under **Settings**, locate **Controller Password**. Click **Set/Change Controller Password**. A new window appears.

**Figure 8-22.** Set/Change Controller Password




4. In the boxes under New Password, type a suitable password

**Figure 8-23.** Controller—New Password Fields




5. Click **OK**.

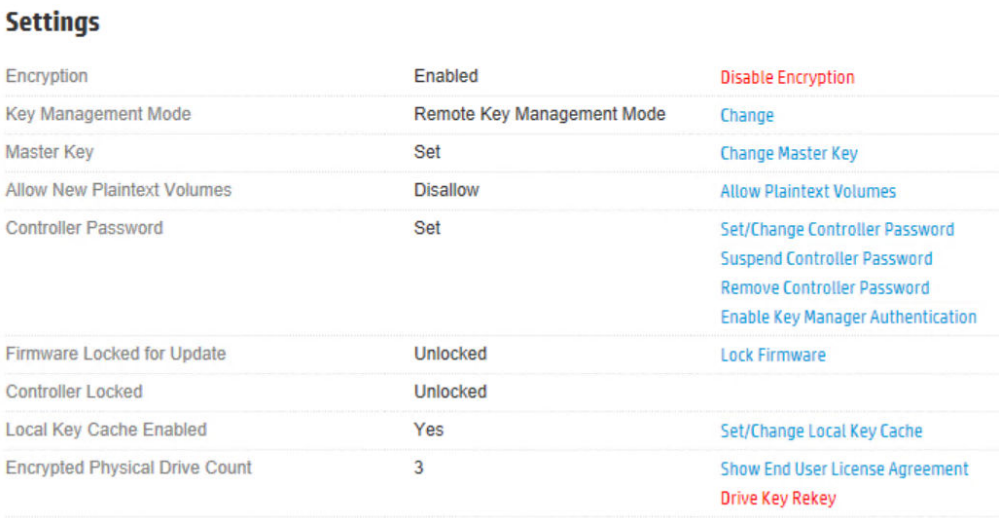
### 8.5.1.2.5. Suspending the Controller Password

If you suspend the controller password, then the controller does not prompt for a password at system startup, and the volumes are allowed online if all the keys are accessible. Once suspended, the controller password feature is resumed without requiring a password reset.

To suspend the controller password, perform the following steps:

1. Open Encryption Manager.
2. Log in to Encryption Manager.
3. Under Settings, locate **Controller Password**, and then click **Suspend Controller Password**. A new window appears, asking if you want to suspend the controller password.

**Figure 8-24.** Suspend Controller Password




4. Click **Yes** to continue.

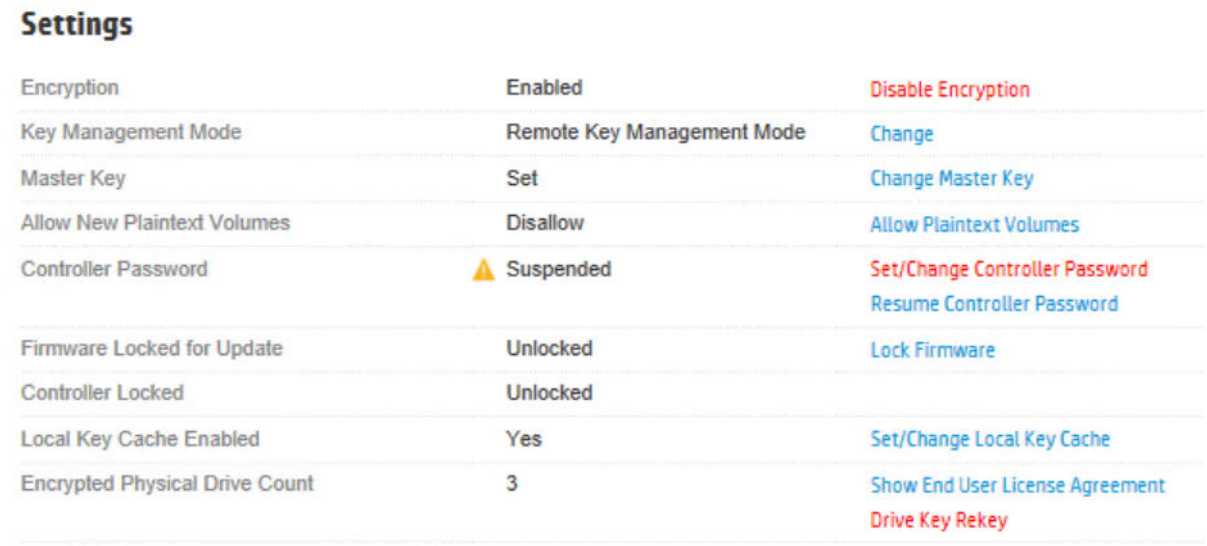
### 8.5.1.2.6. Resuming the Controller Password

When you resume a suspended controller password, the system re-enables password prompt at system startup.

To resume the controller password, perform the following steps:

1. Open Encryption Manager.
2. Log in to Encryption Manager.
3. Under **Settings**, locate **Controller Password**, and then click **Resume Controller Password**. A new window appears, asking if you want to resume the controller password.

**Figure 8-25.** Resume Controller Password




4. Click **Yes** to continue.

### 8.5.1.3. Working with Keys

This section describes how to change the master key encryption, re-key the drive encryption keys, rescan the keys, enable encryption key manager, and volatile keys.

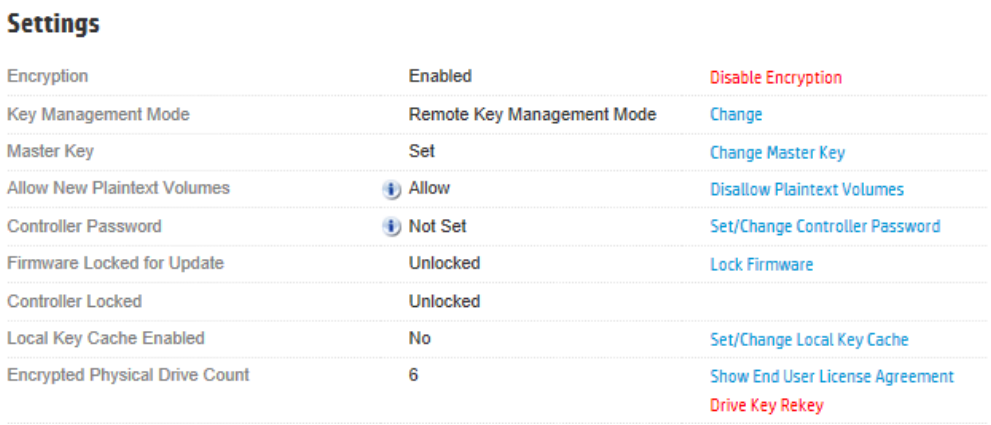
#### 8.5.1.3.1. Changing the Master Encryption Key

**Important:** It is recommended that you keep a record of the Master Encryption Keys when encryption is configured in Local mode. The local Master Encryption Key is not displayed by any available tool or firmware because it is considered a cryptographic secret by FIPS 140-2. Secure Encryption design follows the NIST architecture requirements and does not allow us to assist in the recovery of a lost Master Encryption Key.

To change the Master Encryption Key, perform the following steps:

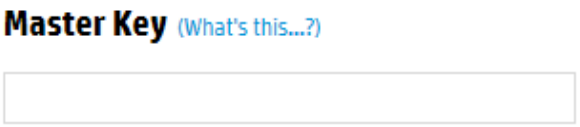
1. Open Encryption Manager.
2. Log in to Encryption Manager.
3. Under **Settings**, locate **Master Key**, and then click **Change Master Key**. A new window appears.

**Figure 8-26.** Change Master Key




4. In the Master Key box, type a suitable value.
   **Note:** While using the Local Key Management mode, the input is any set of printable characters. While using the Remote Key Management mode, the input must be the same name as the key name in the remote key store.

**Figure 8-27.** Enter New Master Key




5. Click **OK**.

#### 8.5.1.3.2. Rekeying the Drive Encryption Keys

This procedure creates a new set of Drive Keys used for encrypting the volume keys on the controller. This task is available to all the roles in the system. For controllers, this feature is disabled by default.

To rekey the Drive Keys, perform the following steps:

1. Open Encryption Manager.
2. Log in to Encryption Manager.
3. Under **Settings**, locate **Encrypted Physical Drive Count**, and then click **Drive Key Rekey**. A prompt appears, indicating new Drive Encryption Keys are created for all physical drives.
4. Click **OK** to continue.

#### 8.5.1.3.3. Rescanning Keys

In Remote mode, this procedure signals the controller to retrieve all the encryption keys from the ESKM. This procedure resolves the potentially locked volumes that could have been locked as a result of failure to initially retrieve the associated keys.

To rescan keys, perform the following steps:

1. Open Encryption Manager.
2. Log in to Encryption Manager.
3. Under **Utilities**, click **Rescan Encryption Keys**. A new window appears, indicating iLO retrieves the keys from the ESKM.

**Figure 8-28.** Utilities—Rescan Encryption Keys




4. Click **OK** to continue.

#### 8.5.1.3.4. Enabling Encryption Key Manager Authentication

When you enable the Encryption Key Manager Authentication, it allows the firmware to bypass prompting the user for the controller password when it is able to contact a verified key manager.

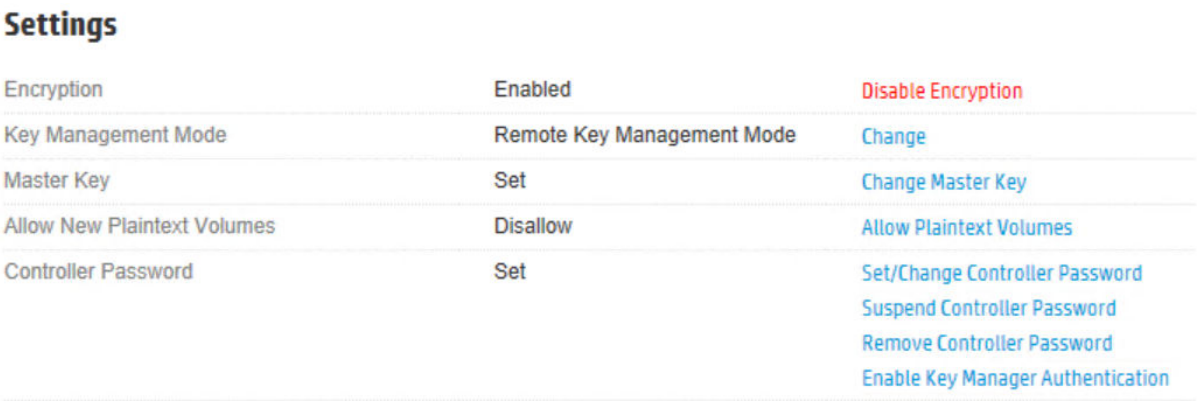
The following encryption settings must be configured before enabling Key Manager Authentication:

- Secure Encryption must be configured to run in Remote Key Management mode
- The controller password must be set
- The 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

To enable Key Manager Authentication, perform the following steps:

1. Open Encryption Manager.
2. Under **Controller Password**, click **Enable Key Manager Authentication**. A new window appears.

**Figure 8-29.** Enable Key Manager Authentication




3. To confirm enabling Encryption Key Manager Authentication, click **Yes**.

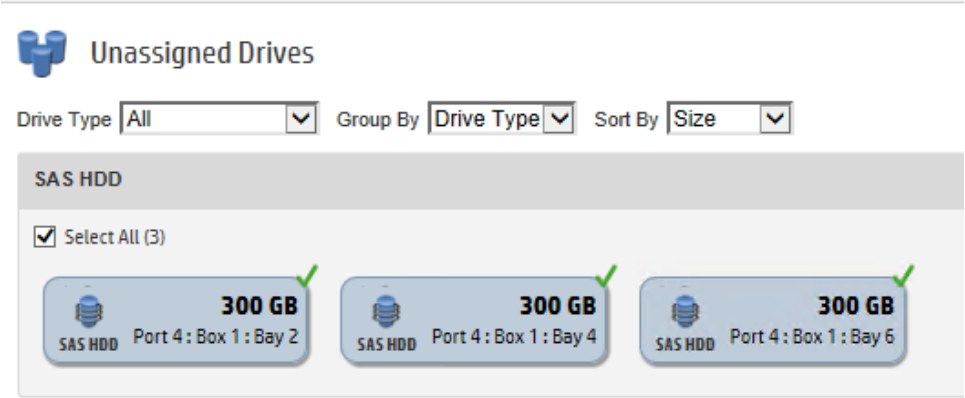
### 8.5.1.4. Creating a Plaintext Volume

**Important:** The controller only allows the creation of new plaintext volumes if it has been configured to do so by the Crypto Officer. To determine if plaintext volume creation is enabled on the controller, see the **Encryption Manager** screen.

To create a plaintext volume, perform the following steps:

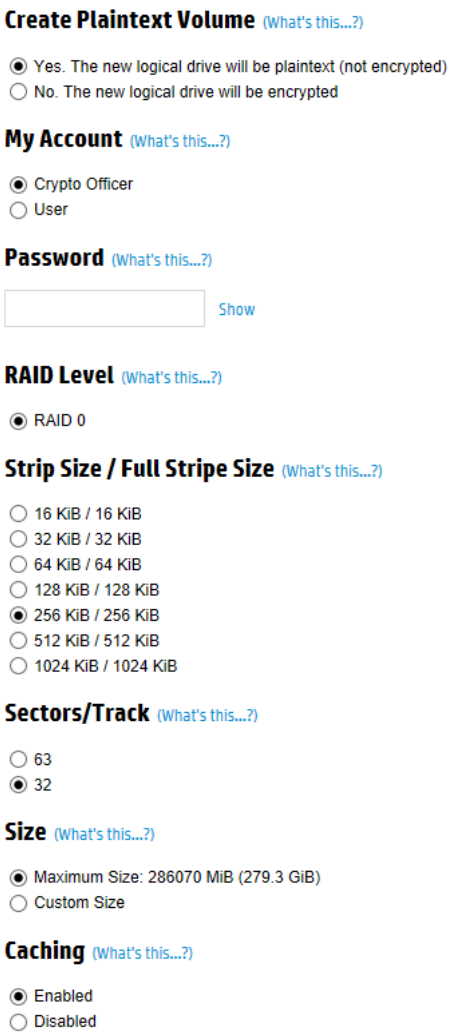
1. Start SSA.
2. Under **Controller Devices**, click **Unassigned Drives**.
3. Select drives.

**Figure 8-30.** Select Drives




4. Click **Create Array**. A new window appears.

**Figure 8-31.** Create Array




5. Complete the following:
   a. **Create Plaintext Volume**: Select **Yes.**
   b. **My Account**: Select the account to log in with.
   c. **Password**: Enter the account password.
6. Configure the remaining fields based on your requirements.
7. Click **Create Logical Drive**.
8. **Array Details**, **Logical Drives**, **Physical Drives**, and **Device Path** specifications appear. Click **Finish** to complete.

### 8.5.1.5. Converting Plaintext Volumes into Encrypted Volumes

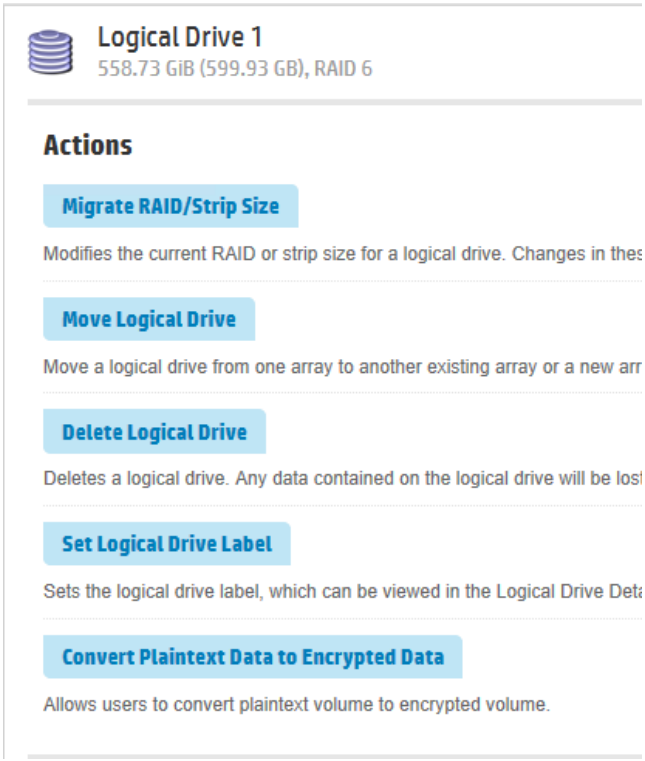
**Note:** In order to preserve existing data, the controller must read and rewrite the entire volume in order to complete the conversion process. Conversion may take some time to complete, especially if there is competing drive activity from the host system.

To convert plaintext volume into encrypted volume, perform the following steps:

1. Open Encryption Manager.
2. Log in to Encryption Manager.
3. Under **Controller Devices**, click **Arrays**.

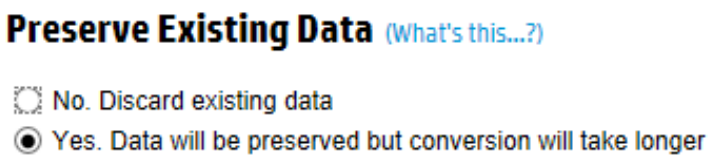
4. Select the plaintext volume.
5. Under **Actions**, click **Convert Plaintext Data to Encrypted Data**.

**Figure 8-32.** SSA Controller Actions




A new window appears.

**Figure 8-33.** Preserve Existing Data




6. Select one of the following:
   a. To preserve existing data, click **Yes**.
   b. To discard existing data, click **No**. If selected, a warning prompt appears after clicking **OK**, confirming your selection. Click **OK** to continue past the warning.
7. Click **OK**. A new window appears, listing the **Logical Drive Details**, **Logical Drive Acceleration Method**, and **Device Path** details.
8. Click **Finish**.

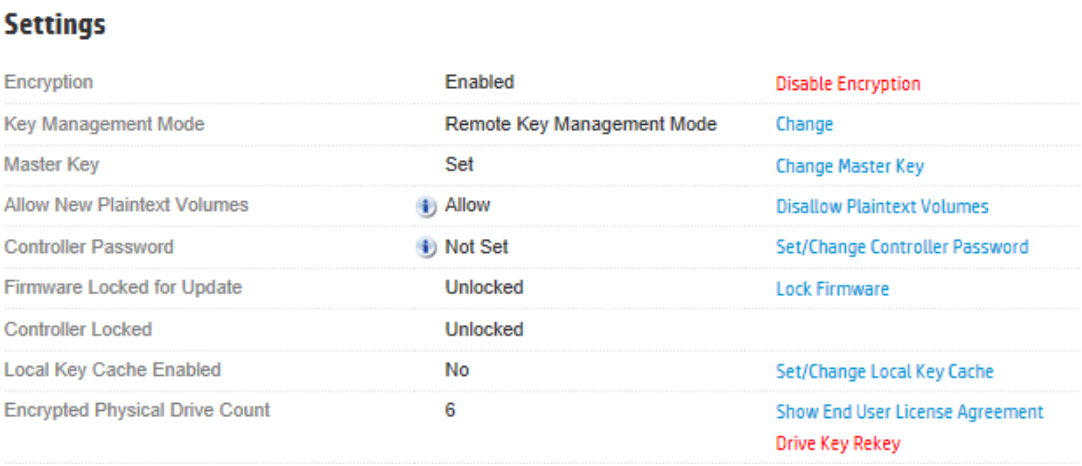
#### 8.5.1.6. Changing Key Management Modes

To change the Key Management modes, perform the following steps:

1. Open Encryption Manager.
2. Log in to Encryption Manager.

3. Under **Settings**, locate **Key Management Mode**. Click **Change**. A new window appears with the key management mode selected.

**Figure 8-34.** Key Management Mode




4. In the Master Key box, type a suitable value.

**Figure 8-35.** Enter Master Key




5. Click **OK**.
6. A warning appears, prompting the user to record the Master Encryption Key. Click **Yes** to continue.

### 8.5.1.7. Enabling/Disabling Plaintext Volumes

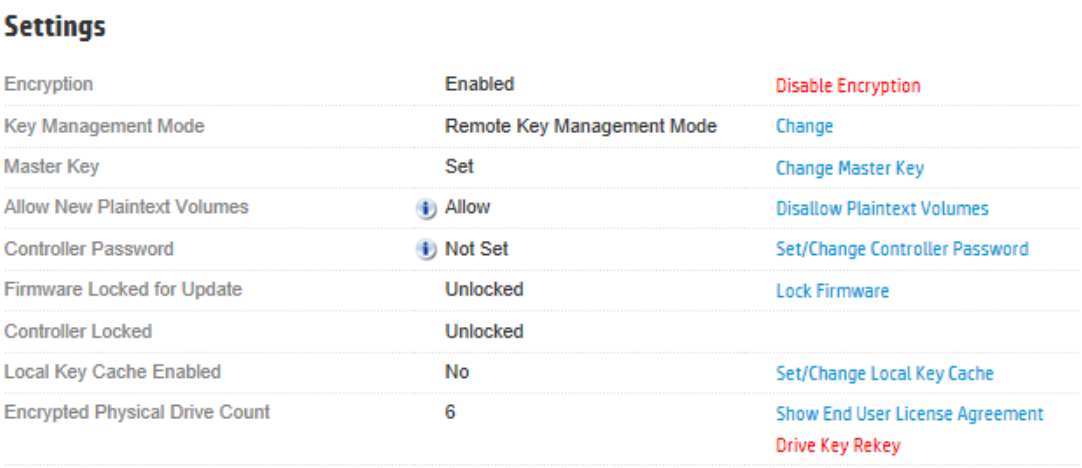
**Important:** Plaintext volumes are unencrypted. The option of allowing or disabling the creation of plaintext volumes depends on the following:

- The type of data to be stored on the plaintext volume
- The level of security you want or need in the system
  It is recommended that you do not enable this option for systems requiring high security or containing highly sensitive data.

To change plaintext volumes permissions after initial configuration, perform the following steps:

1. Open Encryption Manager.
2. Log in as the Crypto Officer.
3. Under **Settings**, locate **Allow New Plaintext Volumes**.

**Figure 8-36.** Enabling/Disabling Plaintext Volumes




4. Do one of the following:
   a. If encryption is disabled, click **Allow Plaintext Volumes**.
   b. If encryption is enabled, click **Disallow Plaintext Volumes**.
5. A prompt appears, asking you to confirm the change. Click **Yes** to continue.

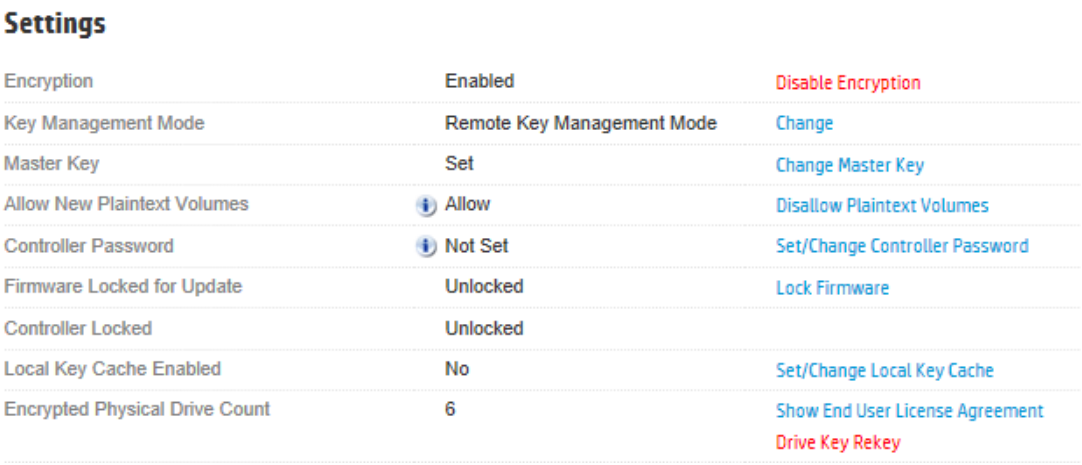
### 8.5.1.8. Enabling/Disabling the Firmware Lock

The firmware lock prevents to update the firmware on the controller, and it is disabled by default. For security purposes, it is recommended to enable the firmware lock function.

To change the firmware lock setting, perform the following steps:

1. Open Encryption Manager.
2. Log in to Encryption Manager.
3. Under **Settings**, locate **Firmware Locked for Update**.

**Figure 8-37.** Firmware Locked for Update



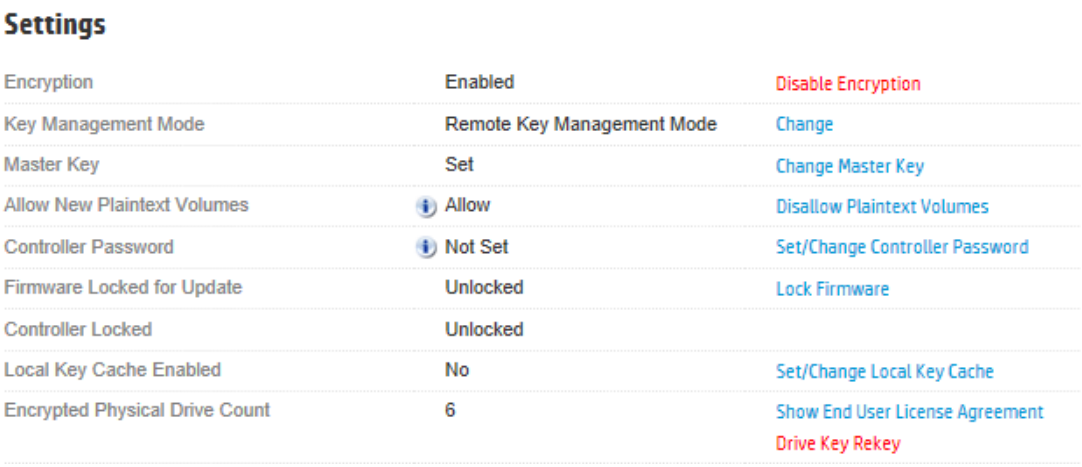

4. Do one of the following:
   a. If unlocked, click **Lock Firmware**.
   b. If locked, click **Unlock Firmware**.
5. A prompt appears, asking you to confirm the change. Click **Yes** to proceed.

### 8.5.1.9. Enabling/Disabling Local Key Cache

Local Key Cache enables the user to store the keys required to decrypt the volume keys in the persistent memory on the controller. When configured for the Remote Key Management mode, the controller normally retrieves the keys from the ESKM at boot time. By storing the key values in the controller, logical drive data can be encrypted and decrypted without the network presence of the ESKM.

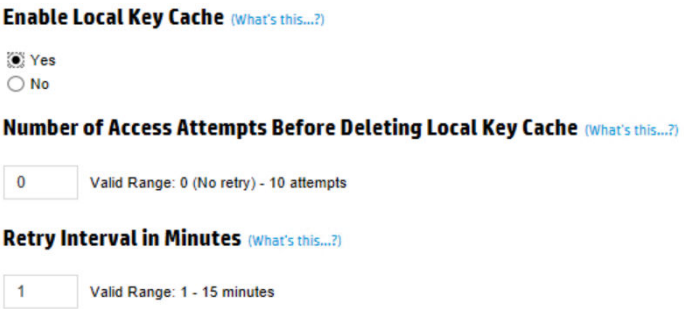
1. Open Encryption Manager.
2. Log in to Encryption Manager.
3. Under **Settings**, locate **Local Key Cache Enabled**,and then click **Set/Change Local Key Cache**.

**Figure 8-38.** Set/Change Local Key Cache




4. Do one of the following:
   a. To disable, click **No**.
   b. To enable, click **Yes**. If you select **Yes**, two new fields appear.

**Figure 8-39.** Enabling/Disabling Local Key Cache




5. Configure the following fields:

**Important:** It is recommended to use the default settings for the number of access attempts. Only change this value if there is a concern that an unintended individual might remove the server from the environment. When the value is set higher than '0', Secure Encryption attempts to locate ESKM the configured number of times during boot. If all attempts fail, the local key cache is deleted prior to boot. All volumes encrypted remains locked until the ESKM is reached and the required keys are retrieved and placed back into the local key cache.

a. **Number of Access Attempts Before Deleting Local Key Cache** – A value of '0' indicates that the Secure Encryption does not check the presence of a key manager, and the key cache remains present on the controller. If the value is greater than '0', Secure Encryption attempts to contact the key manager the number of attempts specified. If any attempt is successful, the encrypted logical drive(s) is unlocked using the keys in the local key cache. If all of the attempts are unsuccessful, then all of the encrypted logical drive(s) remains locked, and the keys in the local key cache are deleted.

b. **Retry Interval in Minutes** – The number of minutes between access attempts.

6. Click **OK**.

#### 8.5.1.10. Importing Drive Sets in Local Key Management Mode

When the Master Encryption Key on an imported drive set is different from the Master Encryption Key on the receiving S Controller, the importing volumes remain offline until user intervention is taken. SSA is used to supply the Master Key name for the importing drives.

In Remote Key Management mode, drives automatically import when the associated key is present on the ESKM. If keys are unable to be retrieved but are confirmed to be on the ESKM, it is possible they are assigned to a different group.

##### 8.5.1.10.1. Importing Drives with Different Master Keys

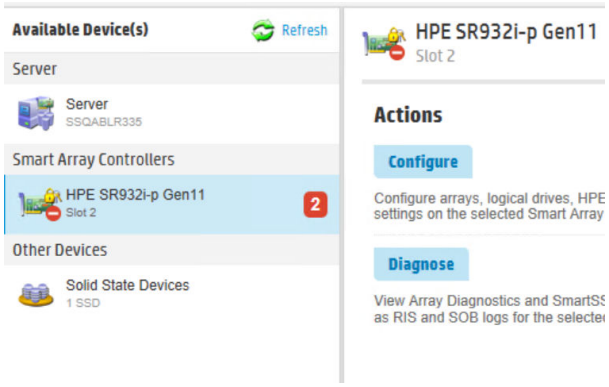
Migrating drives to a non-encrypted controller results in the logical volumes associated with those drives remaining offline until encryption is enabled with the proper Master Encryption Key settings and mode for that volume.

If non-encrypted drives are migrated to an encrypting controller, the controller automatically brings the logical volumes associated with those physical drives online and makes them available for use.

To import drives with a different Master key into a controller when using Local Key Management mode, perform the following steps:

1. Power down the server. For more information, see the documentation that ships with the server.
2. Attach drives. For more information, see the documentation that ships with the drives.
3. Power up the server. For more information, see the documentation that ships with the server.
4. Start SSA.
5. Under **Array Controller(s)**, click the controller assigned to the new drives. Red alert message indicators appears next to it.

**Figure 8-40.** SSA Select Controller Red Alerts



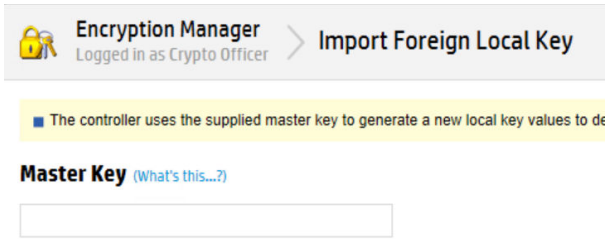

6. Under **Actions**, click **Configure**.
7. From the side menu, click **Encryption Manager**.
8. Log in to Encryption Manager. For more information, see Logging into Encryption Manager (CBE).
9. Under **Utilities**, click **Import Foreign Local Key**.

**Figure 8-41.** SSA Import Foreign Local Key




10. A new screen appears. In the Master Key box, type the new Master Encryption Key name assigned to the drives being imported.

**Figure 8-42.** SSA Import Foreign Local Key—Enter Master key




11. Click **OK**.

The drives are incorporated, unlocked, and assigned the Master Encryption Key of the receiving controller.

### 8.5.2. SED Operations

The following sections describe operations in SED encryption setup.

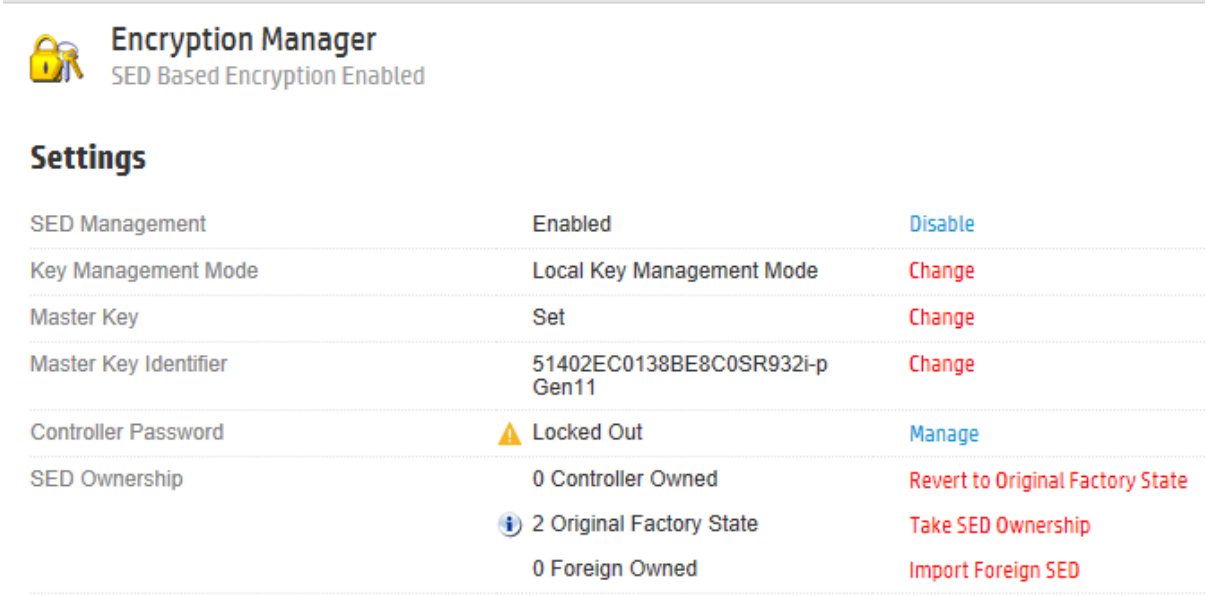
#### 8.5.2.1. Unlocking the Controller Password (SED)

To log into Encryption Manager in SED, perform the following steps:

1. Open Encryption Manager.
2. Under **Setting**, navigate to the **Controller Password** option, and then click **Manage**.

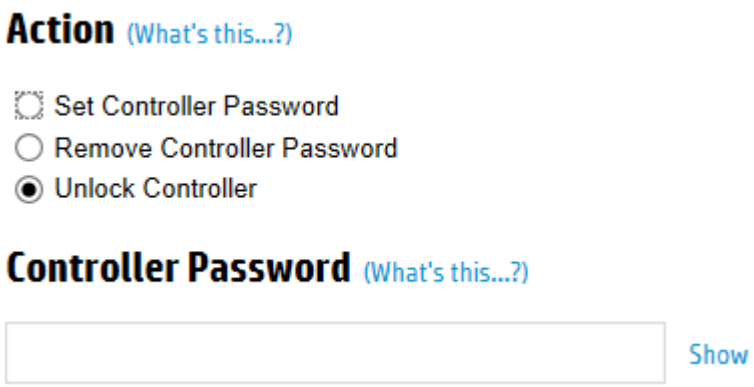
The **Action** window appears.

**Figure 8-43.** Settings—Encryption Manager (SED)




3. Select **Unlock Controller** and type the controller password.

**Figure 8-44.** Action Window




4. Click **OK** to continue.

**Notes:**

- For the Remote Key Management SED setup, use UEFI to set or change the controller password.
- The controller will be locked after three unsuccessful login attempts. To unlock the controller, a reboot is necessary, or the controller password must be removed.

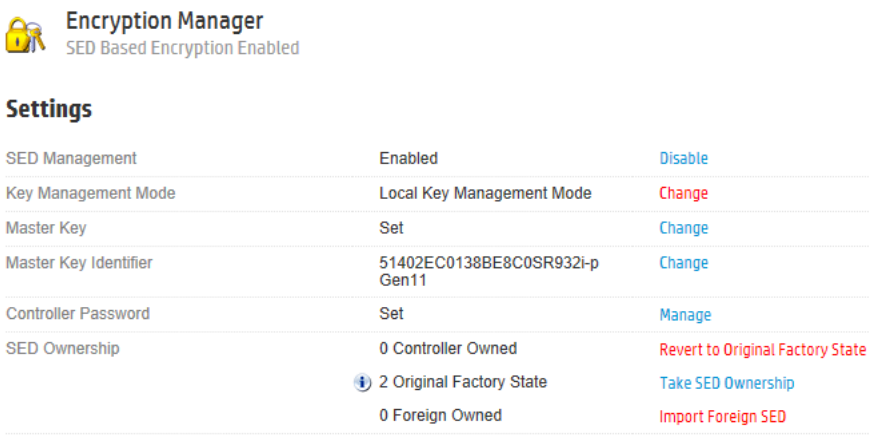
### 8.5.2.2. Disabling SED Based Encryption

**Note:**
To disable the SED encryption feature, delete all secured logical drives from the controller.

To disable SED based encryption, perform the following steps:

1. Open Encryption Manager.
2. Under **Settings**, locate **SED Management**, and then click **Disable**. A new window appears.

**Figure 8-45.** SSA Settings—SED Encryption Setup



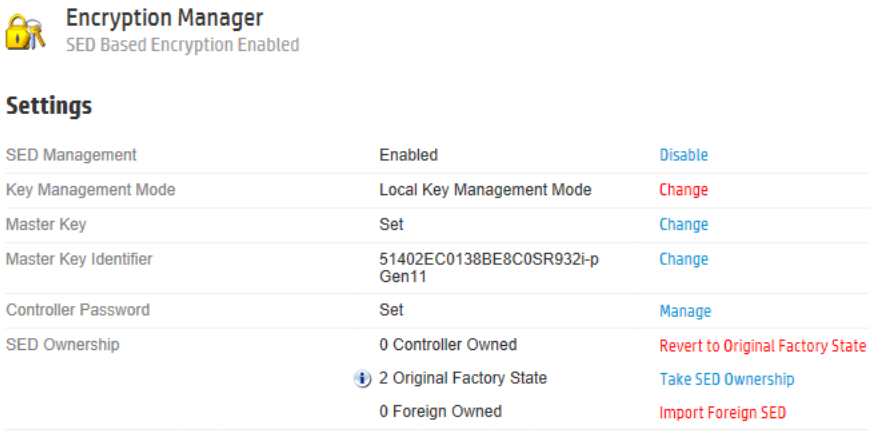

3. Click **Yes**.

### 8.5.2.3. Changing Master Key

To change the master key, perform the following steps:

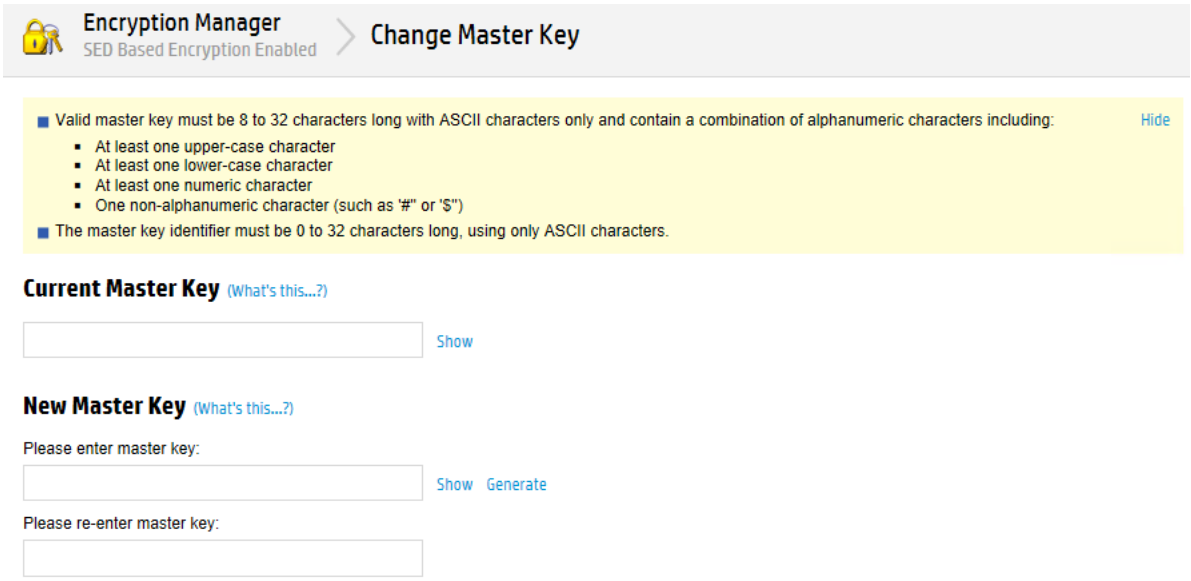
1. Open Encryption Manager.
2. Under Settings, locate **Master Key**, and then click **Change**.

**Figure 8-46.** SSA Settings—SED Encryption Setup




3. Type the **Current Master Key** and the **New Master Key** in the respective boxes.

**Figure 8-47.** Change Master Key




4. Click **OK**.

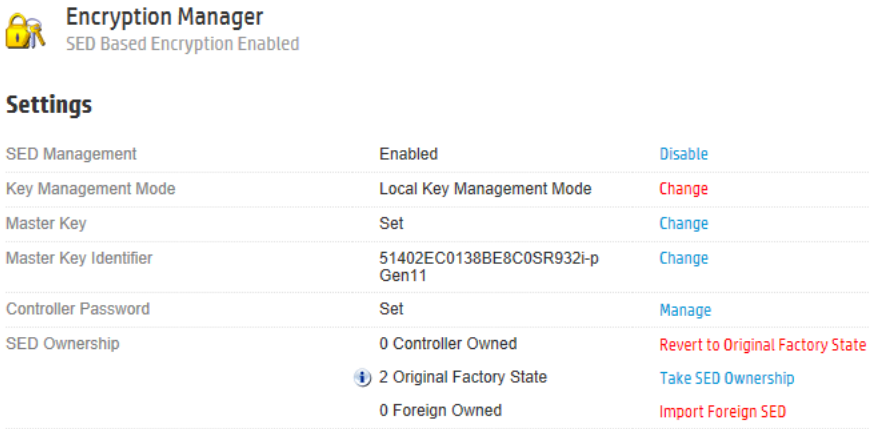
**Note:** For the Remote Key Management SED setup, users are unable to set or change the Master Key and the option is disabled as denoted by the red "Change" text. Instead, the UEFI HII rekey operation is used to have the KMS server generate a new Master Key and Master Key ID.

### 8.5.2.4. Changing Master Key Identifier

To change the master key identifier, perform the following steps:

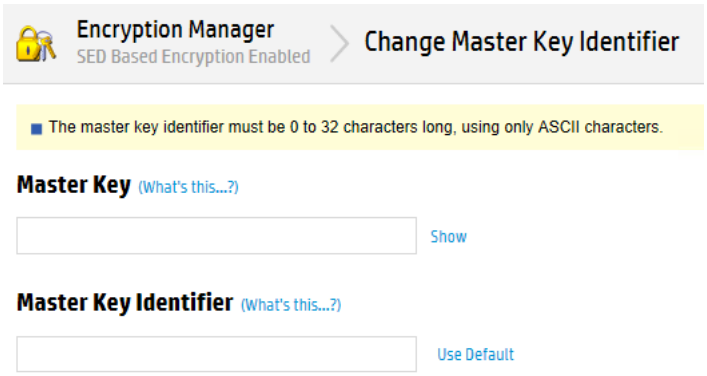
1. Open Encryption Manager.
2. Under Settings, locate **Master Key Identifier**, and then click **Change**.

**Figure 8-48.** SSA Settings—SED Encryption Setup




3. Type the **Master Key** and the new **Master Key Identifier** in the respective boxes.

**Figure 8-49.** Change Master Key Identifier




4. Click **OK**.

**Note:** For the Remote Key Management SED setup, users are unable to set or change the Master Key Identifier. Instead, the UEFI HII rekey operations is used to have the KMS server generate a new Master Key and Master Key ID.

### 8.5.2.5. Managing Controller Password

The following sections describe how to manage the controller password.

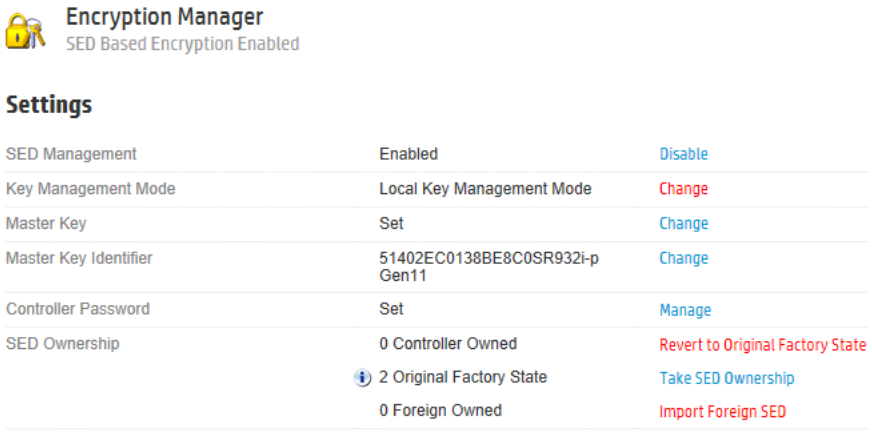
**Note:** For the Remote Key Management SED setup, use UEFI to set or change the controller password.

#### 8.5.2.5.1. Set the Controller Password

To set or change the controller password, perform the following steps:

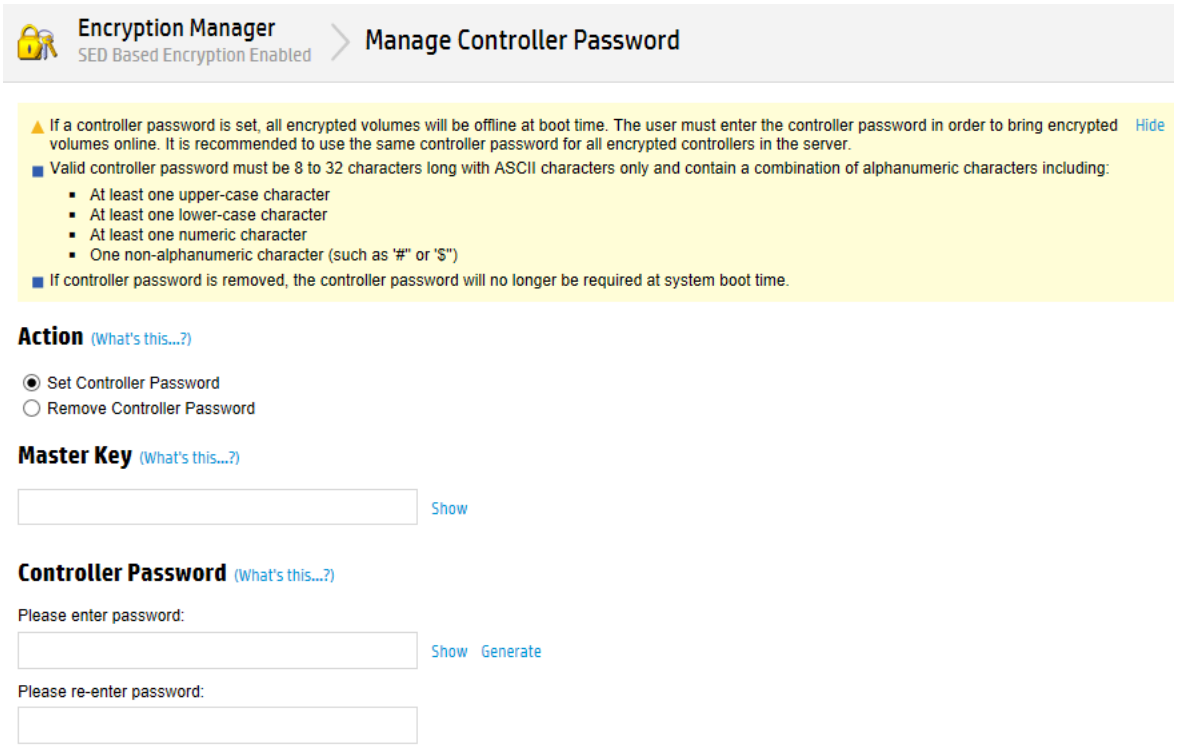
1. Open Encryption Manager.
2. Under Settings, locate **Controller Password,** and then click **Manage**.

**Figure 8-50.** SSA Settings—SED Encryption Setup




3. Under **Action**, click **Set Controller Password.**
4. In the **Master Key** box type the new Master Encryption Key and in the **Controller Password** boxes, type suitable password.

**Figure 8-51.** Manage Controller Password




5. Click **OK**.

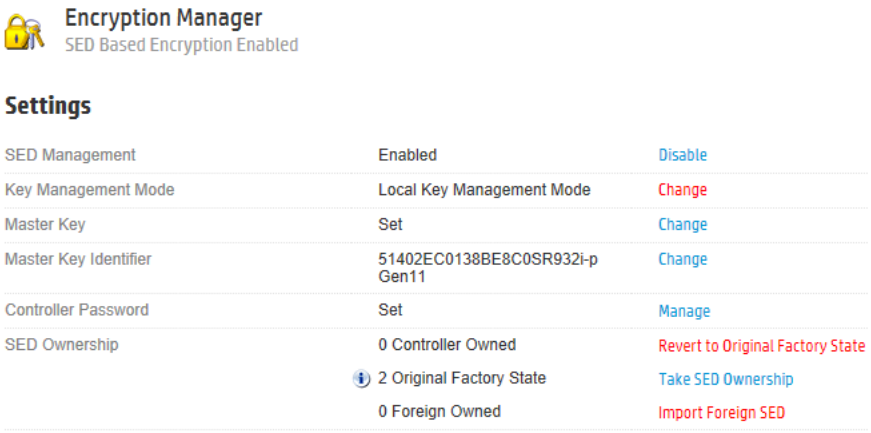
**Note:** For the Remote Key Management SED setup, use UEFI to set or change the controller password.

### 8.5.2.5.2. Removing Controller Password

To remove the controller password, perform the following steps:

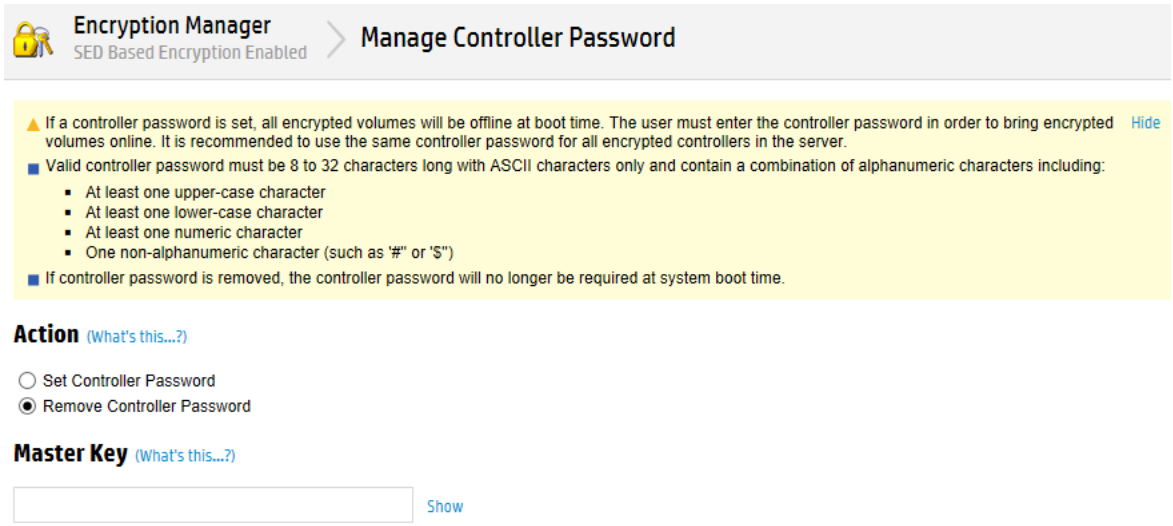
1. Open Encryption Manager.
2. Under Settings, locate **Controller Password**,and then click **Manage**.

**Figure 8-52.** SSA Settings—SED Encryption Setup




3. Under Action, click **Remove Controller Password**. In the **Master Key** box, type a suitable value.

**Figure 8-53.** Remove Controller Password




4. Click **OK**.

### 8.5.2.6. Managing SED Ownership

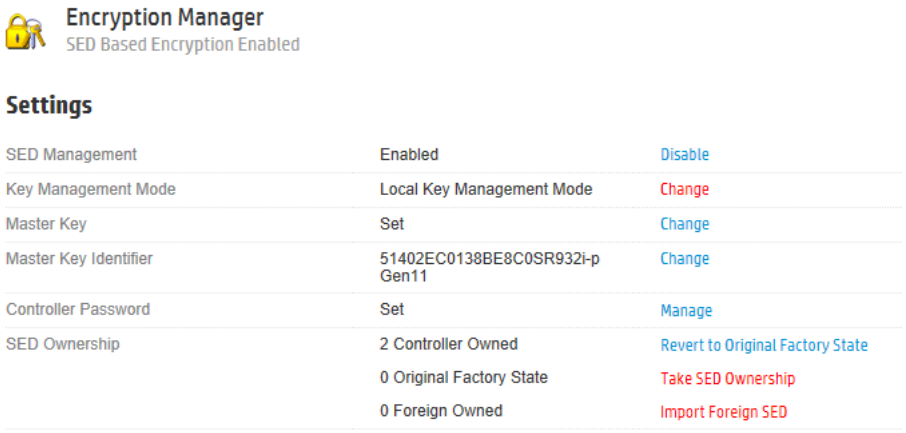
The following sections describe how to manage the SED ownership.

#### 8.5.2.6.1. Revert to the Original Factory State

To revert SEDs owned to the original factory state, perform the following steps:

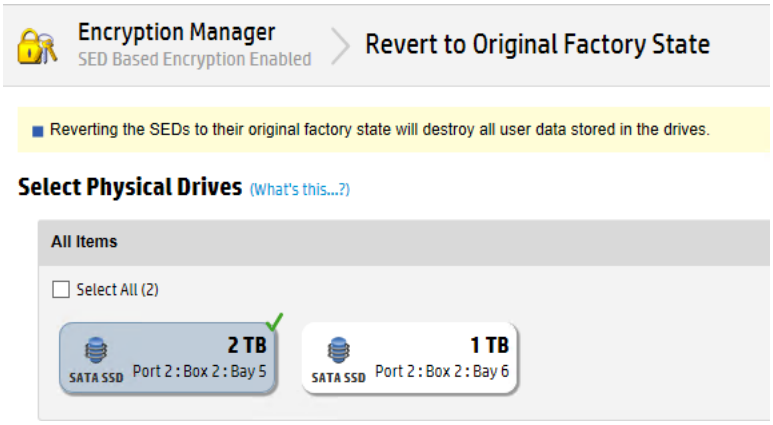
1. Open Encryption Manager.
2. Under **Settings**, locate **SED Ownership**, and then click **Revert to Original Factory State**.

**Figure 8-54.** SSA Settings—SED Encryption Setup




3. Under **Select Physical Drives**, select the SED drive to revert to the original factory state.

**Figure 8-55.** Revery to Original Factory State—Select Physical Drives



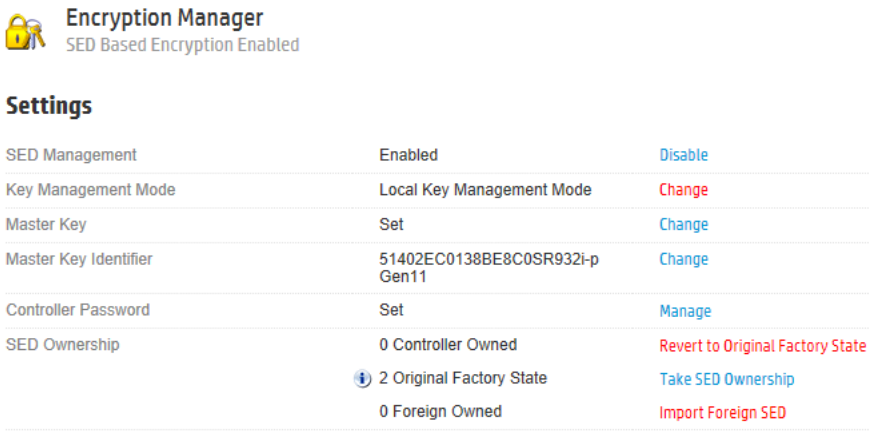

4. Click **OK**.

### 8.5.2.6.2. Taking SED Ownership

In this operation, the controller takes the ownership of the selected SED drives and manages their encryption settings. To take SED ownership, perform the following steps:

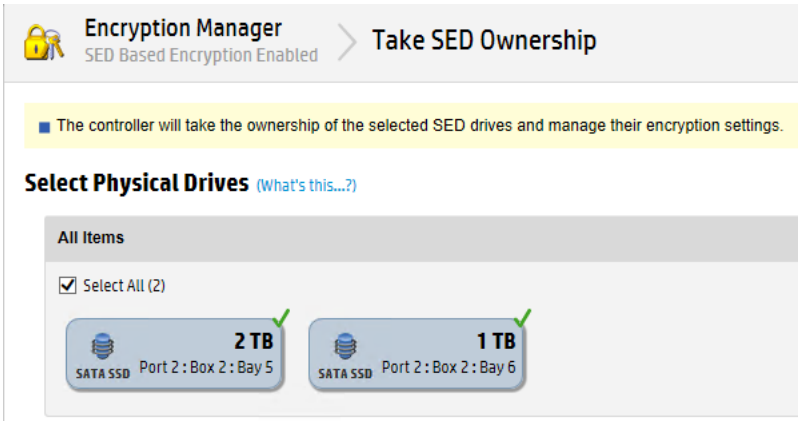
1. Open Encryption Manager.
2. Under **Settings**, locate **SED Ownership**. Click **Take SED Ownership**.

**Figure 8-56.** SSA Settings—SED Encryption Setup




3. Under **Select Physical Drives**, select the SED drive to take ownership.

**Figure 8-57.** Take SED Ownership—Select Physical Drives




4. Click **OK**.

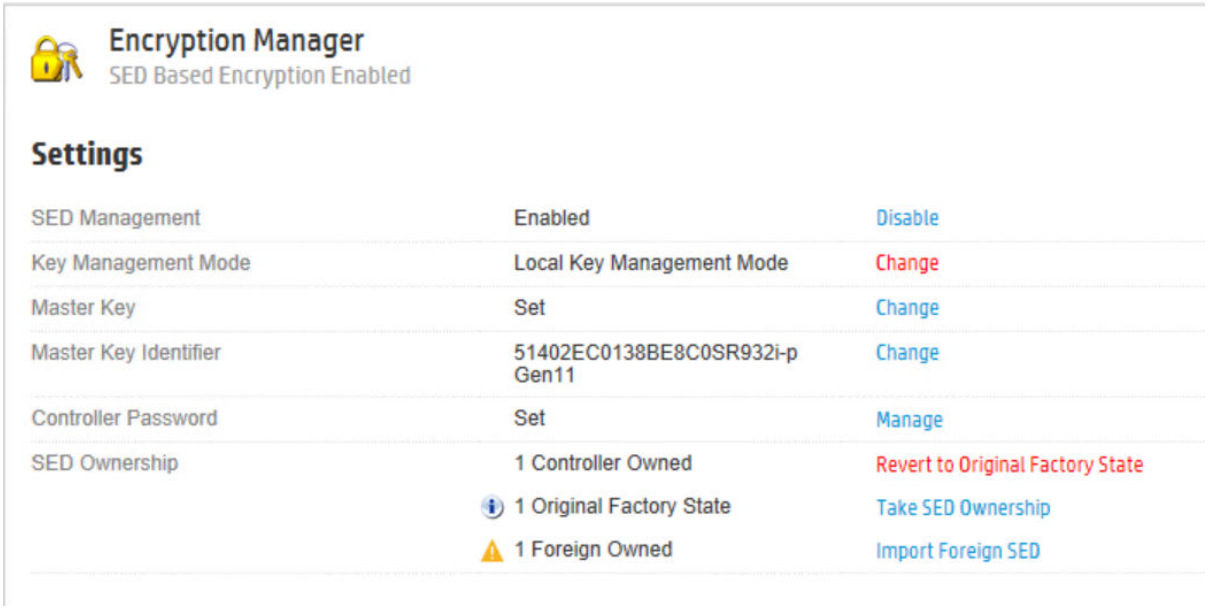
### 8.5.2.6.3. Importing Foreign SED

If you migrate encrypted drives to a non-encrypted controller, the logical volumes associated with those drives become offline until encryption is enabled with the proper Master Encryption Key settings and mode for that volume. If you migrate non-encrypted drives to an encrypting controller, the controller automatically brings the logical volumes associated with those physical drives online and makes them available for use.

To import foreign SED, perform the following steps:

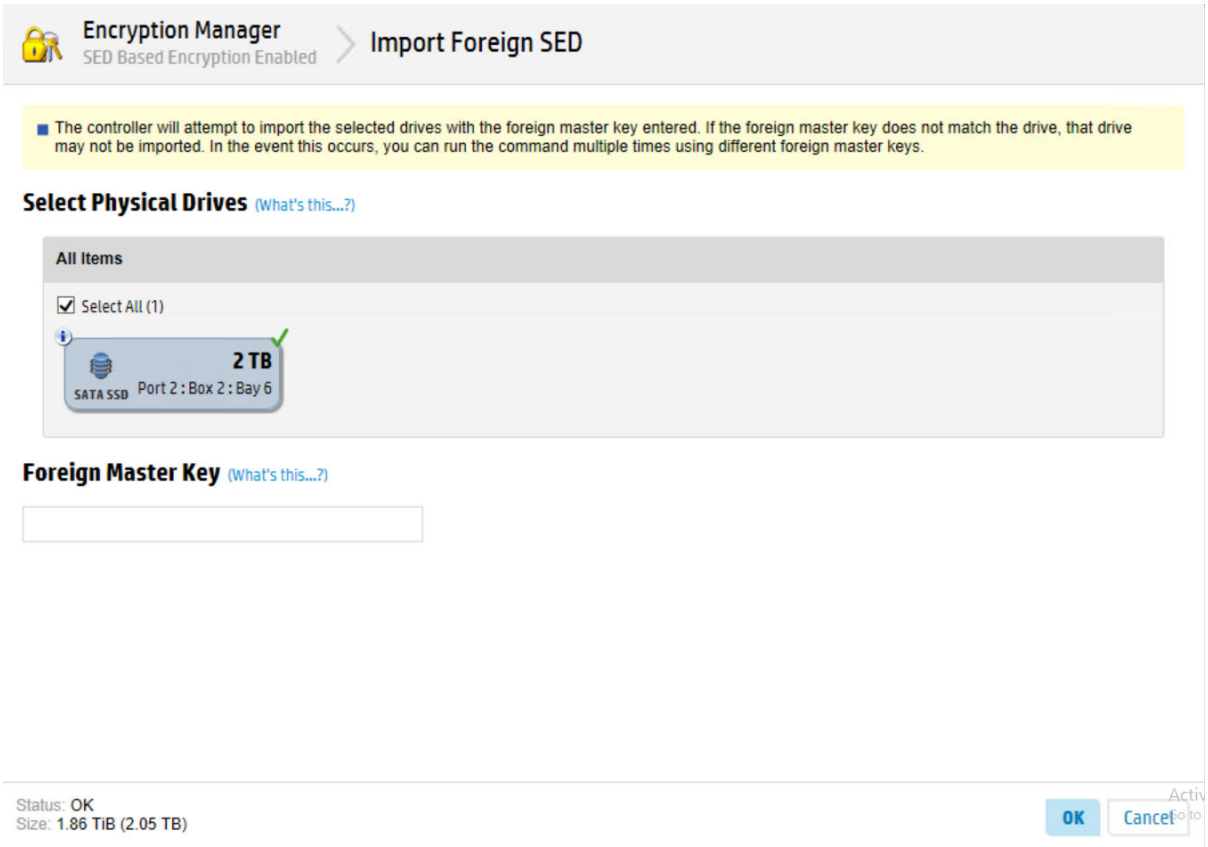
1. Open Encryption Manager.
2. Under **Settings**, locate **SED Ownership**, and then click **Import Foreign SED**.

**Figure 8-58.** SSA Settings—SED Encryption Setup




3. Under **Select Physical Drives**, click the foreign SED drive to take ownership.
4. Type the Master Key of the foreign SED drive.

**Figure 8-59.** Import Foriegn Key—Select Physical Drives




5. Click **OK**.

**Notes:**

- The user may need to run import multiple times
- The user can run import SED from an array or a physical drive level
- For a remoted foreign SED, the master key is imported with HEX format

### 8.5.2.7. Creating Logical Drive (SED)

This section explains how to create a logical device using SED drives. It is not allowed to mix drives to create a logical device. For example, the following characteristics are not mixed:

- Mixing of block sizes (512 Bytes and 4K)
- Interface types, drive types (HDD and SSD)
- SED types (Opal and Enterprise).

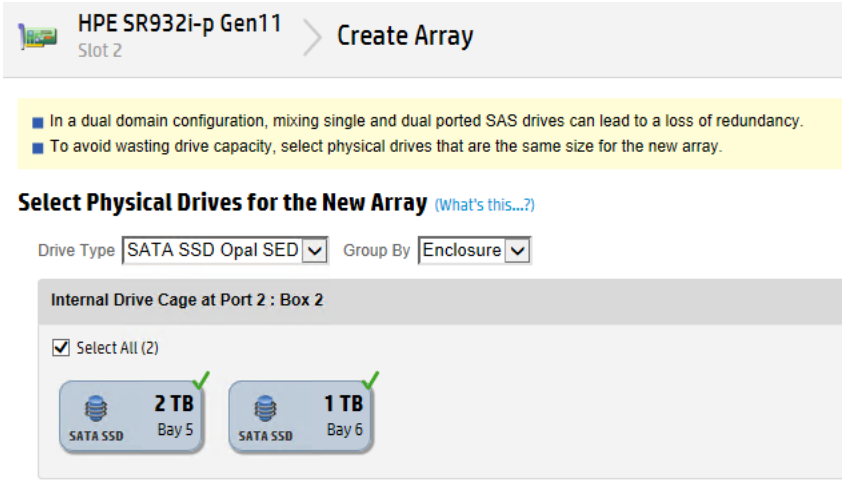
To create a logical device, the SED drive must be either in OFS or owned.

To create a logical drive (SED), perform the following steps:

1. Open SSA.
2. Open the **Configure** panel by any of the following ways:
   a. Choose a device and click **Configure** in the quick navigation menu.
   b. From the Home screen, select an available device, and then click **Configure** under the available options.
3. Under **Controller Devices**, select **Logical Devices**.
4. Under **Action** panel, click **Create Array**.

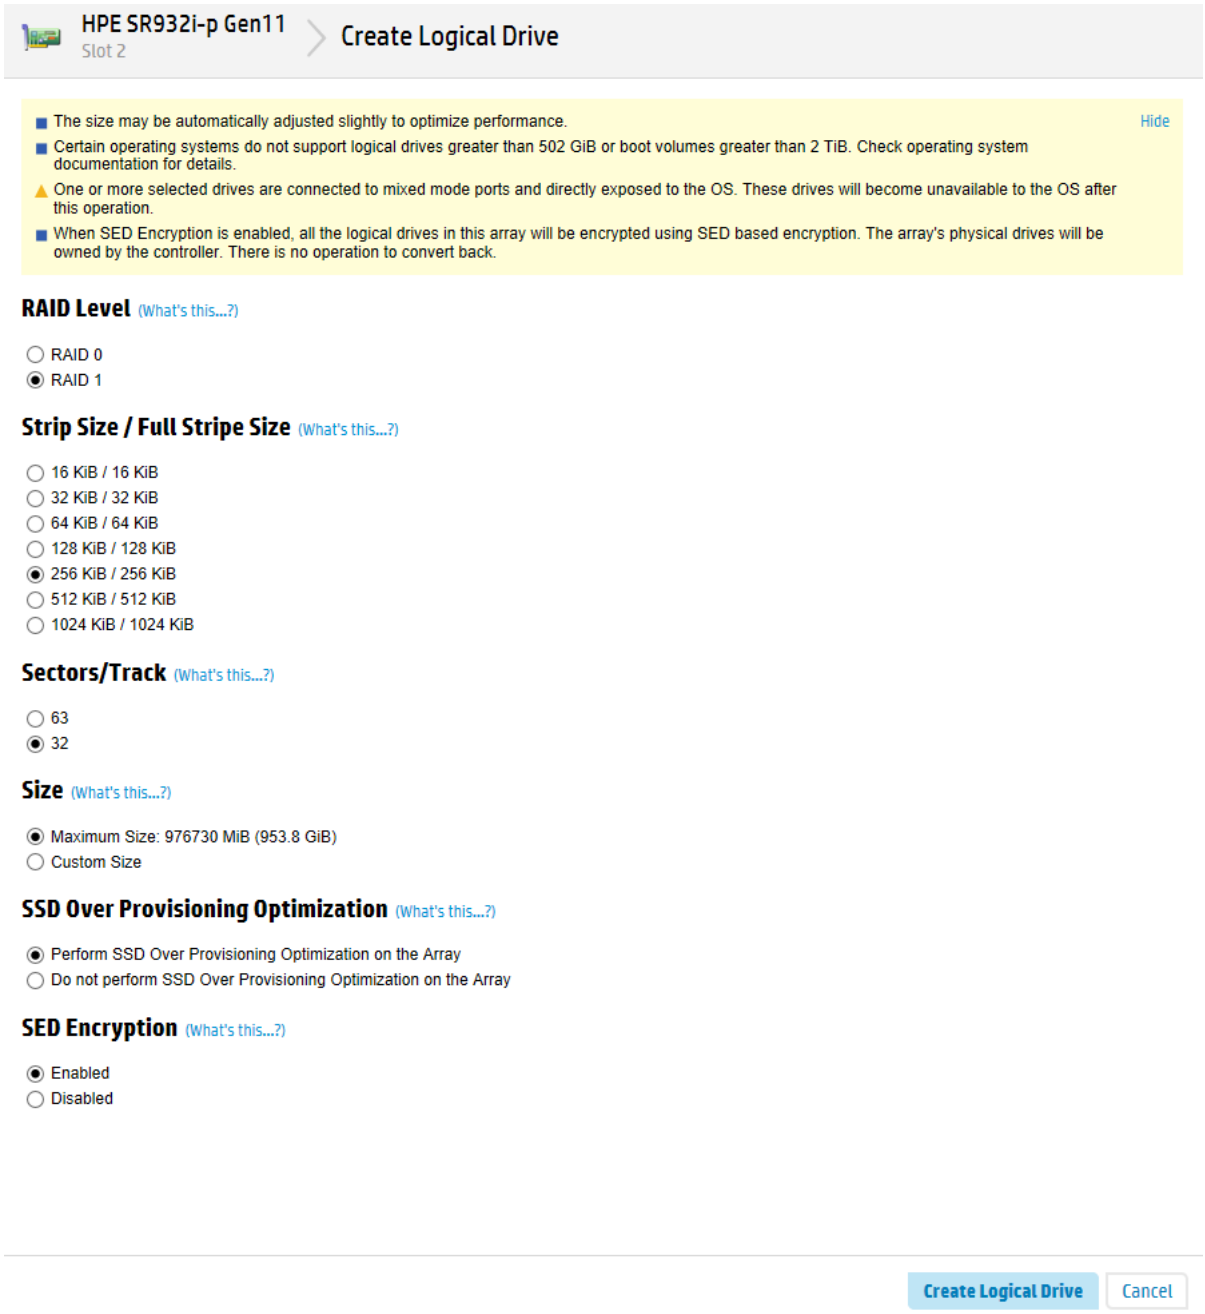
5. Under **Select Physical Drives for the New Array**, select **Drive Type**. Select the SED drive to create a logical drive. A new window appears.

**Figure 8-60.** Creating Logical Drive (SED)




6. Click **Yes**.
7. Configure **RAID Level**, S**trip Size/Full Stripe Size**, **Sectors/Track**, **Size**, **SSD Over Provisioning Optimization**, and **SED Encryption**. See the following figure.

**Figure 8-61.** Creating Logical Drive (SED)—Selection Window




8. Click **Create Logical Drive**. A new window appears.
9. Click **Yes**.
10. Click **Finish**.

## 8.6. Maintenance

This section describes the maintenance of the controllers, the drives, and the groups, as well as how to view the log information and run the queries.

### 8.6.1. Maintaining Controllers

This section describes the maintenance of the controllers.

#### 8.6.1.1. Clearing the Controller

To clear all logical drives and arrays on controllers, perform the following steps:

1. Start SSA.
2. Select the controller to be cleared.
3. Under **Actions**, click **Clear Configuration**. A new window appears, confirming your request to clear the controller's configuration.
4. To continue, click **Clear**.
5. A new window appears, displaying controller settings and configuration. To continue, click **Finish**.

#### 8.6.1.2. Replacing an Encrypted Controller with CBE

If some or all of the drives managed by the controller being replaced are encrypted, you must re-configure the replacement controller with the same settings and Key Management mode you used for the controller you are replacing. For more information, see the documentation that ships with the controller.

In Local Key Management mode, you must provide the correct Master Encryption Key name that matches the one used for the attached drives.

In Remote Key Management mode, any valid Master Encryption Key name is acceptable, since the Master Encryption Key names are part of the drive configuration information stored on each drive.

#### 8.6.1.3. Replacing a Server while Retaining the Controller with CBE

If you retain the same controller and the physical disks, then there are no encryption-related tasks to complete.

If Remote Key Management mode is in use, the previous iLO configuration for key management must be applied to the new server.

For more information on configuring iLO, see Configuring iLO.

#### 8.6.1.4. Preconfiguring Replacement Components

It is possible to configure replacement controllers in advance for encryption. After installing the SmartRAID Controller, enable encryption on the controller. For more information, see Configuration.

After the server is powered down, the controller is physically removed and set aside for later use.

#### 8.6.1.5. Flashing Firmware with CBE

If the firmware lock function is enabled, the firmware lock on the controller must be unlocked before attempting to flash the controller. To disable the firmware lock function, see Enabling/Disabling the Firmware Lock.

### 8.6.2. Drives

This section describes the maintenance of the drives.

#### 8.6.2.1. Replacing a Physical Drive

To replace a drive, see the *Server Maintenance and Service Guide*.

### 8.6.3. Displaying Log Information

The event log displays events for all the controllers in the system and does not differentiate between the events produced by different controllers.

When operating Secure Encryption in Remote mode, you can access the ESKM events log for information on key retrieval and exchange, including the following:

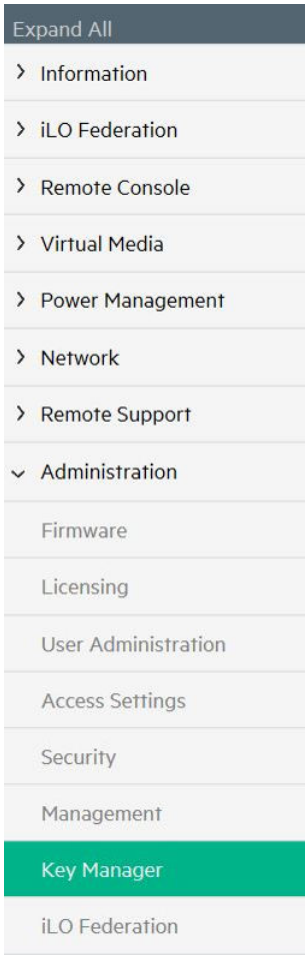
- Connection status

- Master Encryption Key retrieval
- Drive Key retrieval
- Drive Key save requests
- Drive Key deletion

To view the event log, perform the following steps:

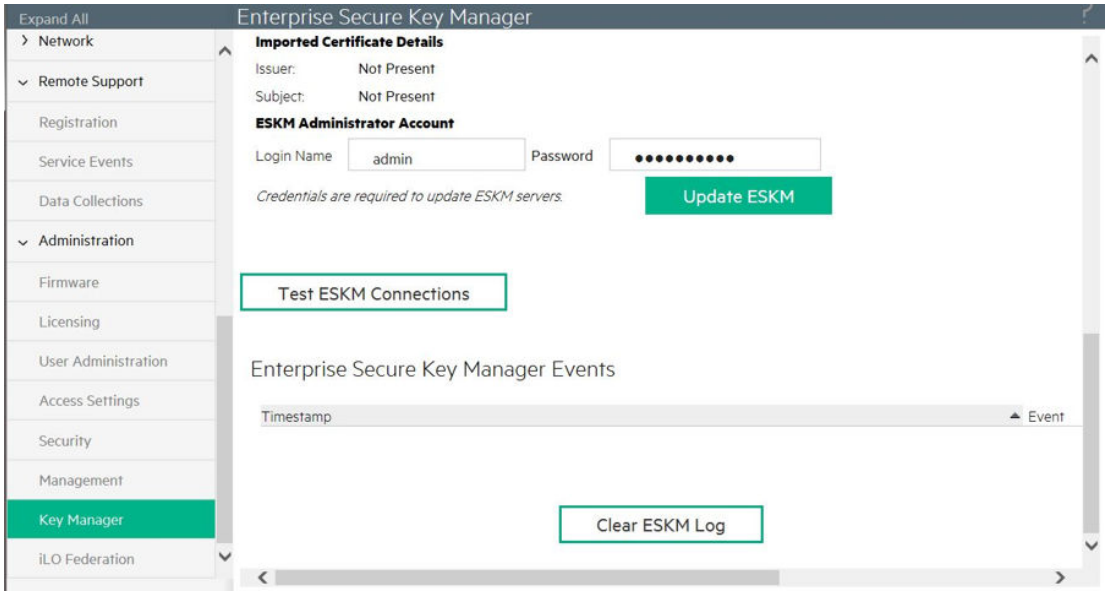
1. Log in to iLO using your server's credentials.
2. From the left side panel, expand the **Administration** menu.

**Figure 8-62.** iLO Key Manager



3. Click **Key Manager**. The **Enterprise Secure Key Manager Events** appears at the bottom of the screen.

**Figure 8-63.** iLO Key Manager Event Log




To refresh the list of events, navigate away from the page and return or click Test ESKM Connections.

# 9. Troubleshooting

The following sections describe troubleshooting for issues in SSA diagnostics utility CLI, 512e Physical Drive support, drive arrays and fault tolerance methods, and to diagnose array problems.

## 9.1. SSA Diagnostics Utility CLI

The SSA Diagnostics Utility CLI is a standalone tool that collects all possible information about storage devices in the system, detects all problems, and provides a detailed configuration report. The following sections describe the SSA Diagnostics Utility CLI.

### 9.1.1. About the Utility

The SSA 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 SSA, which can be run in an offline environment.

The utility generates two types of reports:

- Array diagnostic report:
  This report contains information about devices those are attached to SmartRAID controllers, such as array controllers, storage enclosures, drive cages, as well as logical, physical, and tape drives. For supported SSDs, 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 SSDs attached to the system.

For more information, see Reported Information.

### 9.1.2. 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 SSDs attached to the system.

For discovered SSDs, the report summary page provides the following calculated totals:

- Total Solid State Drives with Wearout Status
- Total SmartRAID Solid State Drives
- Total Non-SmartRAID Solid State 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
- SmartRAID Controllers
- Non SmartRAID Controllers

When generated report is viewed in a browser, the report page displays the following fields in the SmartSSD Status table.

**Table 9-1.** SmartSSD Status Table

| Field | Description |
|---|---|
| SSD Wear Status | Indicates the SSD's wear status with one of the following messages:<br>• OK<br>• Not Supported<br>• The SmartSSD Wear Gauge log is full. Wear Gauge parameters are not available.<br>• SSD has less than 5% usage remaining before wearout.<br>• SSD has less than 2% usage remaining before wearout.<br>• SSD has less than an estimated 56 days before it reaches the maximum usage limit for writes (wearout) and must be replaced as soon as possible.<br>• 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 must be replaced as soon as possible.<br>• 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 must be replaced as soon as possible.<br>• SSD has reached the maximum rated usage limit for writes (wearout) and must be replaced immediately. |
| Power Cycles | Indicates the number of times the SSD has powered on from the powered off state |
| Power On Hours | Indicates the number of hours the SSD has been powered on |
| Estimated Life Remaining Based On Workload To Date | Indicates an estimate of the number of days the SSD has before SSD Utilization reaches 100%. This field is displayed only if the value shown for Usage Remaining is less than 100%. |
| Usage Remaining | Indicates the percentage of the SSD that has not worn out. Usage remaining is equal to the difference of 100 and the SSD Utilization percentage. |
| SSD Utilization | Indicates the percentage of the SSD that has worn out. This field is displayed only if the value shown for Usage Remaining is less than 100%. |

### 9.1.3. Installing the Utility

To install the utility, perform the following steps:

1. Browse to the Smart Storage Administrator website.
2. Click **Download software**.
3. Select an OS.
4. Identify the preferred software and version, and then click **Download**.
5. Save, and then run, the executable file.
   By default, the software installs at `C:\Program Files\Smart Storage Administrator\`.

### 9.1.4. Launching the Utility in CLI Mode

To launch the utility in CLI mode, perform the following steps:

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

   For more information, see the *Launching the Utility in CLI Mode* section in *Smart Storage Administrator Command Line Interface Guide*.

For more options, use the following command:

```
ssaducli -help
```

### 9.1.5. Diagnostic Report Procedures

The following sections describe the procedure to identify and view a diagnostic report.

#### 9.1.5.1. Viewing the Diagnostic Report

To view the diagnostic report, perform the following steps:

1. Launch the utility. For more information, see Launching the Utility in CLI Mode.
2. Browse to the .zip file you created using the utility.
3. Open the HTML file to view the report.

#### 9.1.5.2. 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

To identify and view diagnostic report files, perform the following steps:

1. Extract `ADUReportViewer.htm` to a directory.
2. Open `ADUReportViewer.htm` in the browser.

### 9.1.6. SmartSSD Wear Gauge Report Procedures

The following sections describe the procedure to view SmartSSD Guage report.

#### 9.1.6.1. Viewing the SmartSSD Wear Gauge Report

To view the SmartSSD wear guage report, perform the following steps:

1. Launch the utility. For more information, see Launching the Utility in CLI Mode
2. Browse to the .zip file you created using the utility.
3. Open the HTML file to view the report.

#### 9.1.6.2. 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 identify and view SmartSSD Wear Guage Report files, perform the following steps:

1. Extract the following files to a single directory:
   - `SmartSSDWearGaugeReport.json`
   - `SmartSSDWearGaugeReport.htm`
2. Open `SmartSSDWearGaugeReport.htm` in the browser.

## 9.2. 512e Physical Drive Support

SSA is able to detect and correct performance issues due to nonoptimal logical drive alignment for the 512e physical drives.

The following scenarios indicate that drive support is needed:

- Multiple logical drives exist in a single array
- An array consists of one or more 512e physical drives.
- At least one of the logical drives in the array is not aligned on a native block boundary. For current 512e drives, the native block boundary is 4K.

SSA displays a warning indicating that the logical drive is not optimally aligned and that performance of the logical drive is not optimal. Additionally, if one or more of the following scenarios are met, the array presents a **Realign Logical Drive** button:

- There is enough free space in the array to move the logical drive to be aligned to the native 4K boundary
- The controller can perform the transformation (requires a cache module with a fully charged energy pack).
- The controller does NOT have SmartCache enabled.

## 9.3. Drive Arrays and Fault-Tolerance Methods

The following sections describe drive arrays and fault-tolerance methods.

### 9.3.1. Drive Arrays

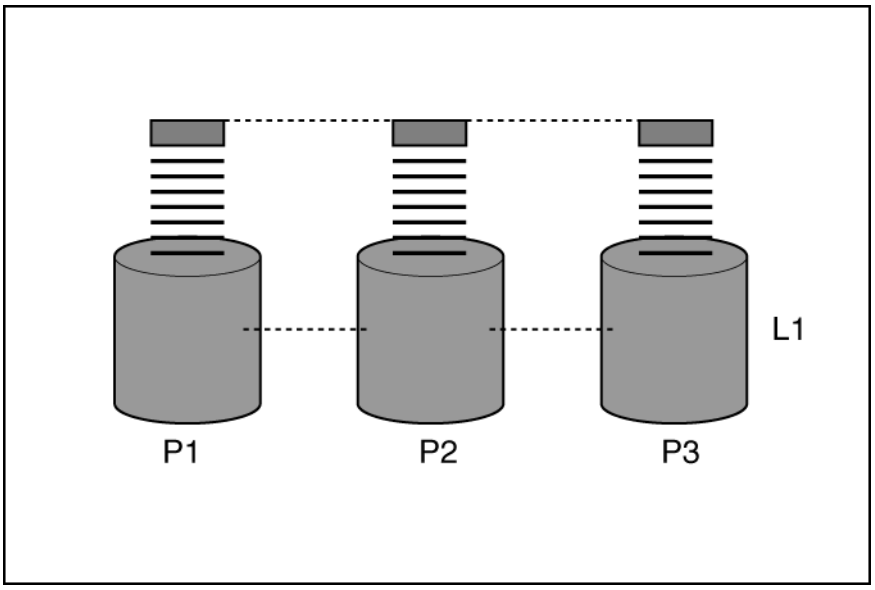
The capacity and performance of a single physical (hard) drive is adequate for home users. However, business users demand higher storage capacities, higher data transfer rates, and greater protection against data loss when drives fail.

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 is still transferable 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

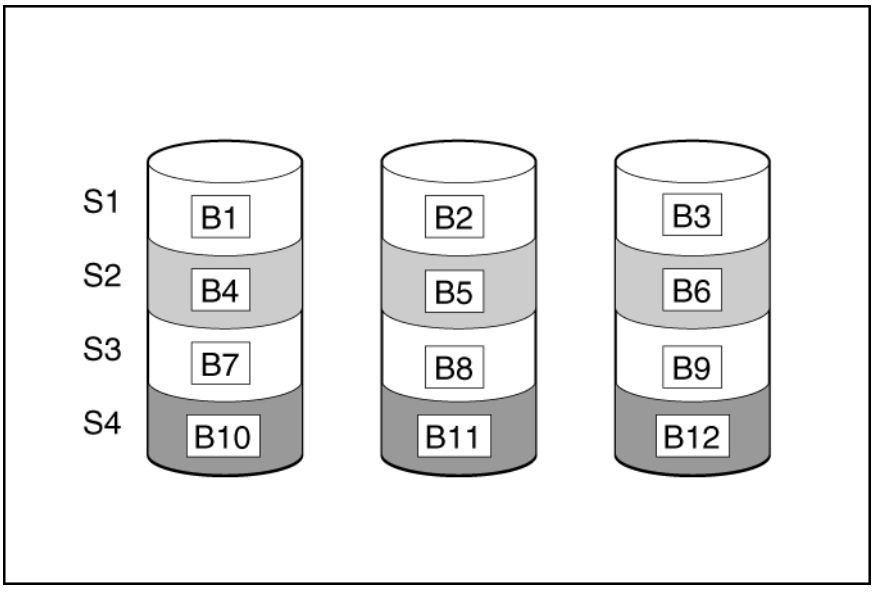
denoted by L*n* in the figures in this section). Then, the read/write heads of all the constituent physical drives are active simultaneously, reducing the total time required for data transfer. The following figure shows the logical drive.

**Figure 9-1.** Logical Drive




The same amount of data is written to each drive during any given time interval because the read/write heads are active simultaneously. Each unit of data is called a block (denoted by B*n* in the figure), and the adjacent blocks form a set of data stripes (S*n*) across all the physical drives that comprise the logical drive.

**Figure 9-2.** Data Striping

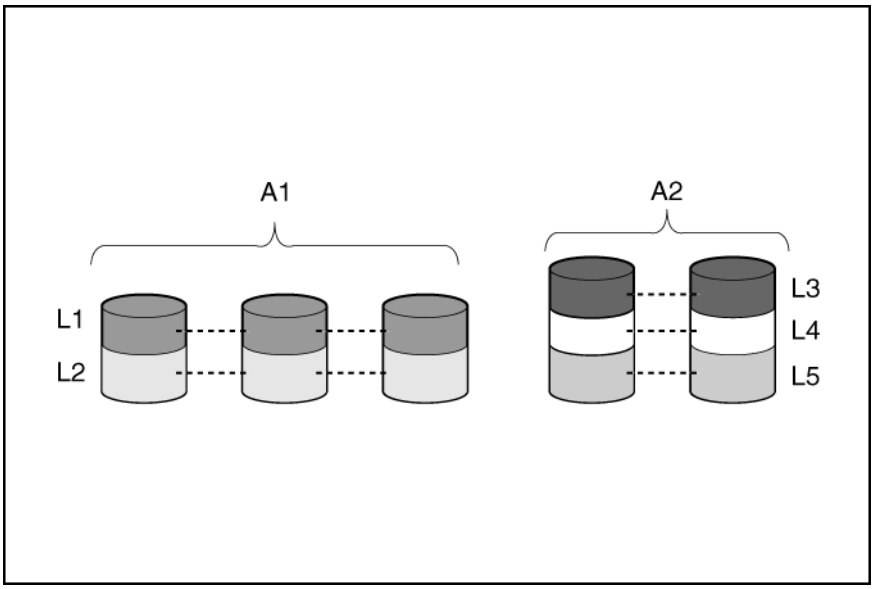



For data in the logical drive to be readable, the data block sequence must be the same in every stripe. This sequencing process is performed by the array controller, which sends the data blocks to the drive write heads in the correct order.

A natural consequence of the striping process is that each physical drive in a given logical drive contains 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.

**Figure 9-3.** Two Arrays




Each logical drive in an array is distributed across all the physical drives within the array. A logical drive can also extend across more than one port on the same controller, but it cannot extend across more than one controller.

Drive failure, although rare, is potentially catastrophic. For arrays that are configured as shown in the previous figure, failure of any physical drive in the array causes every logical drive in the array to suffer irretrievable data loss. To protect against data loss due to physical drive failure, logical drives are configured with fault tolerance.

For any configuration except RAID 0, further protection against data loss can be achieved by assigning a drive as an **online spare** (or **hot spare**). This drive has 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 still fails.)

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.

### Related Links

Fault-Tolerance Methods

### 9.3.2. 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 the RAID 0 logical drives in the same array also fail.
- RAID 1 tolerates the failure of one drive
- RAID 10 configurations tolerate multiple drive failures as long as no failed drives are mirrored to another failed drive
- RAID 5 configurations tolerate one drive failure
- RAID 50 configurations tolerate one failed drive in each parity group
- RAID 6 configurations tolerate two failed drives at a given time
- RAID 60 configurations tolerate two failed drives in each parity group
- RAID 1 (Triple) and RAID 10 (Triple) configurations tolerate multiple drive failures if no more than two drives, mirrored to one another, fail.

### 9.3.3. Fault-Tolerance Methods

Several fault-tolerance methods exist. The hardware-based RAID methods are mostly used with the SmartRAID controllers.

Alternative fault-tolerance methods are also available. However, hardware-based RAID methods provide a much more robust and controlled fault-tolerance environment, so these alternative methods are seldom used.

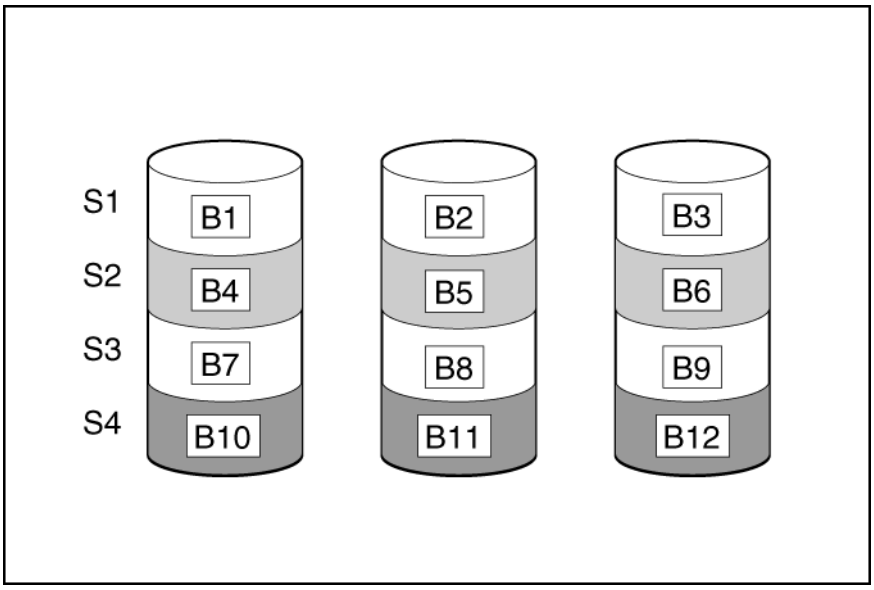
**Related Links**

Alternative Fault-Tolerance Methods

#### 9.3.3.1. RAID 0

A RAID 0 configuration provides data striping, but there is no protection against data loss when a drive fails. However, it is useful for rapid storage of large amounts of noncritical data (for example, for printing or image editing) or when cost is the most important consideration. The minimum number of drives required is one. The following figure shows data stripping.

**Figure 9-4.** Data Striping



This method has the following benefits:

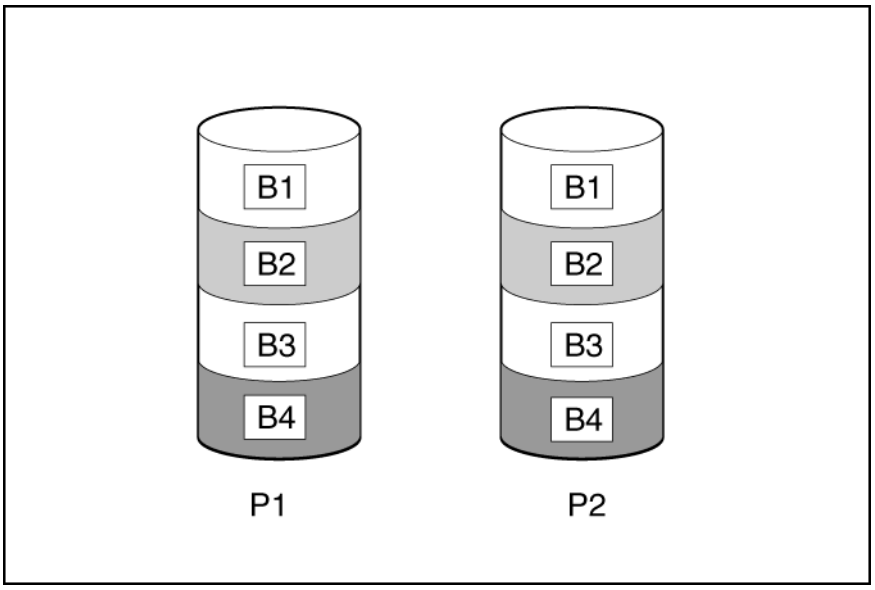
- Useful when performance and low cost are more important than data protection
- Has the highest write performance of all RAID methods
- Has the lowest cost per unit of stored data of all RAID methods
- All drive capacity is used to store data (none allocated for fault tolerance)

### 9.3.3.2. RAID 1 and RAID 1+0 (RAID 10)

In RAID 1 and RAID 1+0 (RAID 10) configurations, data is duplicated to a second drive. The usable capacity is C x (n/2) where C is the drive capacity with n drives in the array. A minimum of two drives is required. The following figure shows drive mirroring.

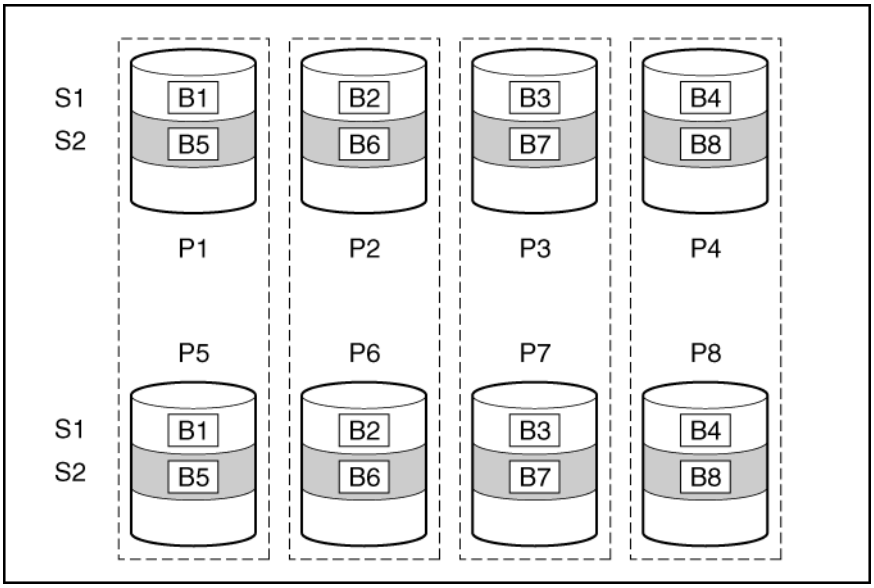
When the array contains only two physical drives, the fault-tolerance method is known as RAID 1.

**Figure 9-5.** Drive Mirroring




When the array has more than two physical drives, drives are mirrored in pairs, and the fault-tolerance method is known as RAID 1+0 or RAID 10. If a physical drive fails, the remaining drive in the mirrored pair still provides all the necessary data. Several drives in the array can fail without incurring data loss, as long as no two failed drives belong to the same mirrored pair. The total drive count must increment by 2 drives. A minimum of four drives is required. The following figure shows mirroring multiple devices.

**Figure 9-6.** Mirroring Multiple Devices




This method has the following benefits:

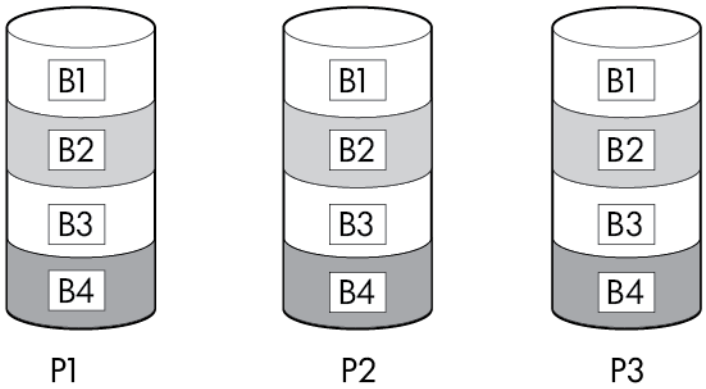
- It is useful when high performance and data protection are more important than usable capacity.
- This method has the highest write performance of any fault-tolerant configuration.
- No data is lost when a drive fails, as long as no failed drive is mirrored to another failed drive.
- Up to half of the physical drives in the array can fail.

#### 9.3.3.3. RAID 1 (Triple) and RAID 10 (Triple)

In RAID 1 (Triple) and RAID 10 (Triple) configurations, data is duplicated to two additional drives. The usable capacity is C x (n / 3) where C is the drive capacity with n drives in the array. A minimum of 3 drives is required.

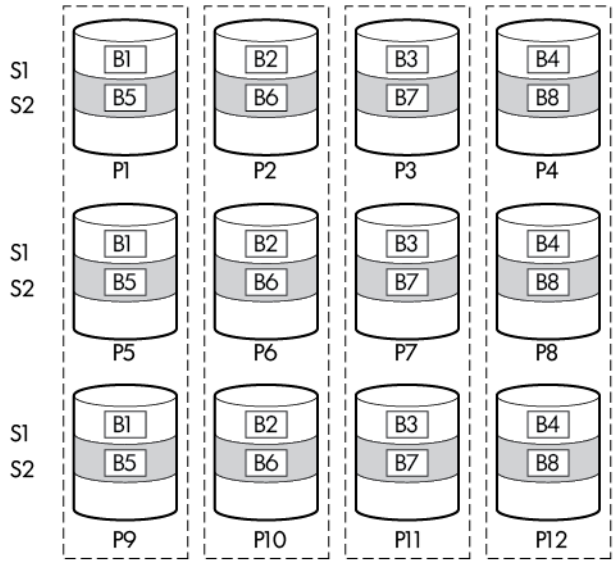
When the array contains only three physical drives, the fault-tolerance method is known as RAID 1 (Triple). The following figure shows RAID 1(Triple).

**Figure 9-7.** RAID 1(Triple)



When the array has more than six physical drives, drives are mirrored in trios, the fault-tolerance method is known as RAID 10 (Triple). If a physical drive fails, the remaining two drives in the mirrored trio can still provide all the necessary data. Several drives in the array can fail without incurring data loss, as long as no three failed drives belong to the same mirrored trio. The total drive count must increment by 3 drives. The following figure shows the RAID 10 (Triple).

**Figure 9-8.** RAID 10 (Triple)



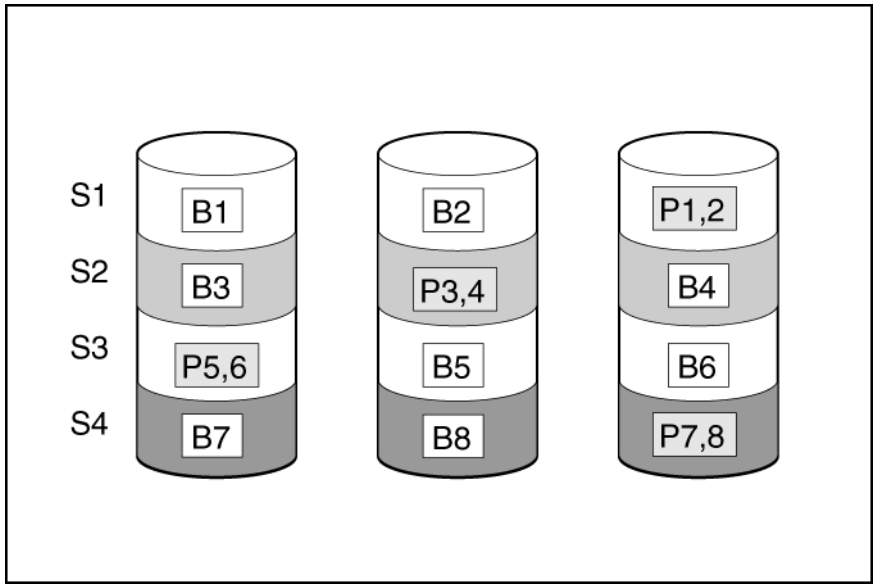

This method has the following benefits:

- It is useful when high performance and data protection are more important than usable capacity
- This method has the highest read performance of any configuration due to load balancing
- This method has the highest data protection of any configuration
- No data is lost when two drives fail, as long as no two failed drives are mirrored to another failed drive
- Up to two-thirds of the physical drives in the array can fail

#### 9.3.3.4. RAID 5

RAID 5 protects data using parity (denoted by Px,y in the figure). Parity data is calculated by summing (XOR) the data from each drive within the stripe. The strips of parity data are distributed evenly over every physical drive within the logical drive. When a physical drive fails, data that was on the failed drive can be recovered from the remaining parity data and user data on the other drives in the array. The usable capacity is C x (n −1) where C is the drive capacity with n drives in the array. A minimum of three drives is required. The following figure shows the RAID 5.

**Figure 9-9.** RAID 5




This method has the following benefits:

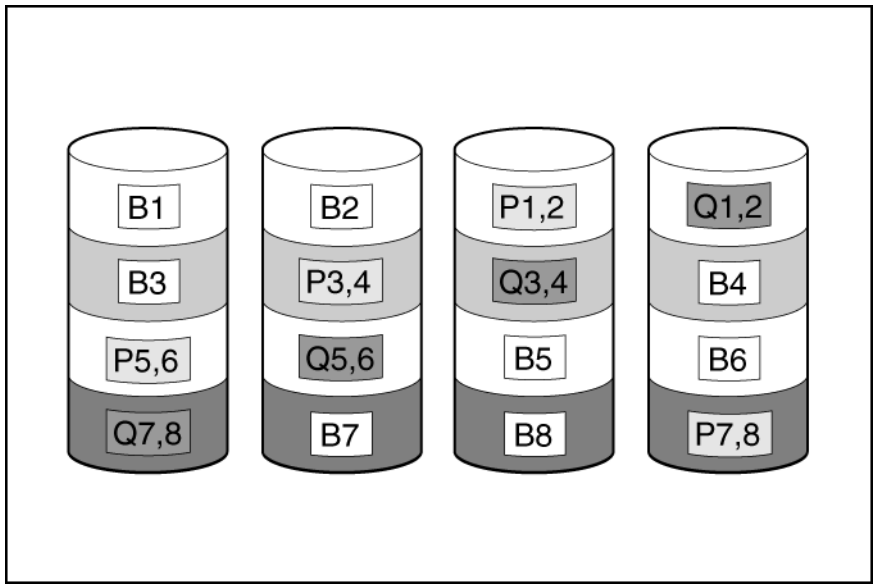
- It is useful when usable capacity, write performance, and data protection are equally important
- It has the highest usable capacity of any fault-tolerant configuration
- Data is not lost if one physical drive fails

### 9.3.3.5. RAID 6

RAID 6 protects the data using double parity. With RAID 6, two different sets of parity data are used (denoted by Px,y and Qx,y in the figure), allowing data to still be preserved if two drives fail. Each set of parity data uses a capacity equivalent to that of one of the constituent drives. The usable capacity is C x (n − 2) where C is the drive capacity with n drives in the array.

A minimum of 4 drives is required.

**Figure 9-10.** RAID 6(ADG)



This method is most useful when data loss is unacceptable but cost is also an important factor. The probability that data loss occurs when an array is configured with RAID 6 (ADG) is less than it would be if it were configured with RAID 5.

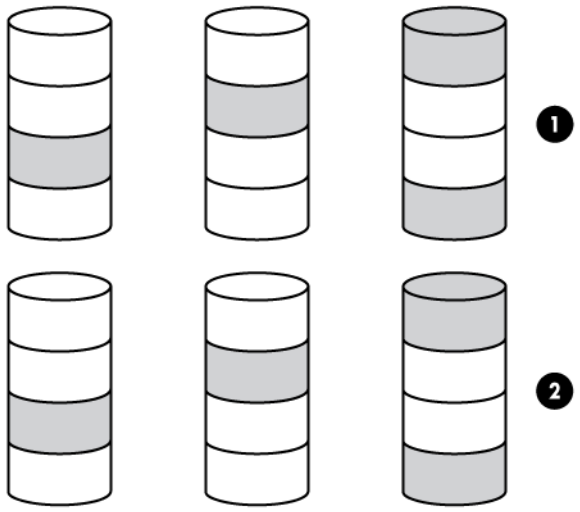
This method has the following benefits:

- It is useful when data protection and usable capacity are more important than write performance
- It allows any two drives to fail without loss of data

### 9.3.3.6. RAID 50

RAID 50 is a nested RAID method in which the constituent drives are organized into several identical RAID 5 logical drive sets (parity groups). The smallest possible RAID 50 configuration has six drives organized into two parity groups of three drives each. The following figure shows RAID 50 arrays.

**Figure 9-11.** RAID 50




For any given number of drives, data loss is least likely to occur when the drives are arranged into the configuration that has the largest possible number of parity groups. For example, four parity groups of three drives are more secure than three parity groups of four drives. However, less data can be stored on the array with the larger number of parity groups.

All data is lost if a second drive fails in the same parity group before data from the first failed drive has finished rebuilding. A greater percentage of array capacity is used to store redundant or parity data than with non-nested RAID methods (for example, RAID 5). A minimum of six drives is required.

This method has the following benefits:

- Higher performance than for RAID 5, especially during writes
- Better fault tolerance than either RAID 0 or RAID 5
- Up to n physical drives can fail (where n is the number of parity groups) without loss of data, as long as the failed drives are in different parity groups.

### 9.3.3.7. RAID 60

RAID 60 is a nested RAID method in which the constituent drives are organized into several identical RAID 6 logical drive sets (parity groups). The smallest possible RAID 60 configuration has eight drives organized into two parity groups of four drives each.

For any given number of hard drives, data loss is least likely to occur when the drives are arranged into the configuration that has the largest possible number of parity groups. For example, five parity

groups of four drives are more secure than four parity groups of five drives. However, less data can be stored on the array with the larger number of parity groups.

The number of physical drives must be exactly divisible by the number of parity groups. Therefore, the number of parity groups that you can specify is restricted by the number of physical drives. The maximum number of parity groups possible for a particular number of physical drives is the total number of drives divided by the minimum number of drives necessary for that RAID level (three for RAID 50, 4 for RAID 60).

A minimum of 8 drives is required.

All data is lost if a third drive in a parity group fails before one of the other failed drives in the parity group has finished rebuilding. A greater percentage of array capacity is used to store redundant or parity data than with non-nested RAID methods.

This method has the following benefits:

- Higher performance than for RAID 6, especially during writes
- Better fault tolerance than RAID 0, 5, 50, or 6
- Up to 2n physical drives can fail (where n is the number of parity groups) without loss of data, as long as no more than two failed drives are in the same parity group.

### 9.3.3.8. Comparing the Hardware-Based RAID Methods

Not all controllers support all RAID levels.

The following table provides the comparison between the hardware-based RAID methods.

**Table 9-2.** Comparing the Hardware-Based RAID Methods

| Item | RAID 0 | RAID 1+0 | RAID 5 | RAID 6 (ADG) | RAID 1(0) (Triple) |
|---|---|---|---|---|---|
| Alternative name | Striping (no fault tolerance) | Mirroring | Distributed Data Guarding | Advanced Data Guarding | Advanced Data Mirroring |
| Formula for number of drives usable for data (n = total number of drives in array) | n | n / 2 | n − 1 | n − 2 | n / 3 |
| Percentage of drive space usable[1] | 100% | 50% | 67% to 93% | 50% to 96% | 33% |
| Minimum number of physical drives | 1 | 2 | 3 | 4 | 3 |
| Tolerates failure of one physical drive | No | Yes | Yes | Yes | Yes |
| Tolerates simultaneous failure of more than one physical drive | No | Only if no two failed drives are in the same mirrored pair | No | Yes | Only if no three drives are in the same mirror group[2] |
| Read performance | High | High | High | High | High |
| Write performance | High | Medium | Low | Low | Medium |
| Relative cost | Low | High | Medium | Medium | Very high |

**Notes:**

1. Values for the percentage of drive space usable are calculated with these assumptions:
   a. All physical drives in the array have the same capacity.
   b. Online spares are not used.
   c. No more than 14 physical drives are used per array for RAID 5.
   d. No more than 56 drives are used with RAID 6 (ADG).
2. Mirror groups include the physical drives in each mirror.

#### 9.3.3.9. Selecting a RAID Method

Not all controllers support all the RAID levels. To determine the RAID capabilities of your controller, see the model-specific information for your controller on the Hewlett Packard Enterprise website.

The following table lists the important criterion to select the RAID level.

**Table 9-3.** Criterion to Select RAID Method

| Most Important Criterion | Also Important | Suggested RAID Level |
|---|---|---|
| Fault tolerance | Cost effectiveness<br>I/O performance | RAID 6<br>RAID 10 (Triple), RAID 1+0, RAID 50, RAID 60 |
| Cost effectiveness | Fault tolerance<br>I/O performance | RAID 6<br>RAID 5 (RAID 0 if fault tolerance is not required) |
| I/O performance | Cost effectiveness<br>Fault tolerance | RAID 5 (RAID 0 if fault tolerance is not required)<br>RAID 10 (Triple), RAID 1+0, RAID 50, RAID 60 |

#### 9.3.3.10. Alternative Fault-Tolerance Methods

Your OS may also support the following software-based RAID or controller duplexing:

- **Software-based RAID** resembles hardware-based RAID, except that the OS works with logical drives as if they are 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 and identical sets of drives containing identical data. In the unlikely event of a controller failure, the remaining controller and drives 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 the maximum storage capacity. For further implementation details, see the OS documentation.

## 9.4. Diagnosing Array Problems

The following sections describe the diagnostic tools to troubleshoot array problems and troubleshooting resources for servers.

### 9.4.1. Diagnostic Tools

To troubleshoot array problems and generate feedback about arrays, use the following diagnostic tools:

- **SSA**
  For information about error messages, see the *ProLiant Servers Troubleshooting Guide*.
- **SSA 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 SSA Diagnostics Utility CLI.

## 9.5. Secured Encryption Issues

This section provides troubleshooting to resolve common issues, including recovering the Master Key Encryption name, restoring a logical drive that is offline, exporting the Master Key, testing the connection between iLO and the ESKM, and clearing the encryption configuration.

### 9.5.1. Common Issues with CBE

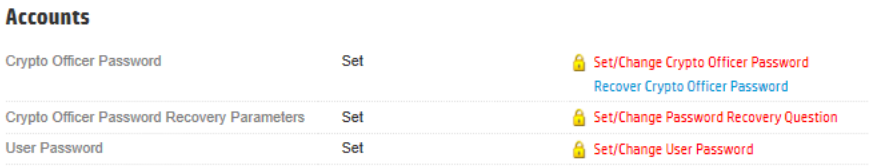
This section describes troubleshooting for the common issues such as lost or forgotten Crypto Officer passwords, controller passwords, and Master Keys.

#### 9.5.1.1. Lost or Forgotten Crypto Officer Password

If you have lost or forgotten Crypto Officer password, perform the following steps:

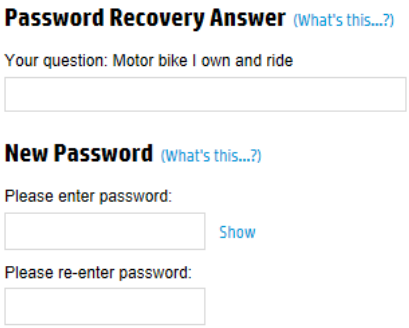
1. Open Encryption Manager.
2. Under **Accounts**, locate **Crypto Officer Password**, and then click **Recover Crypto Officer Password**.

**Figure 9-12.** SSA Recover Crypto Officer Password




A new window appears.

**Figure 9-13.** Recover Crypto Officer Password




3. Do the following:
   a. In the **Password Recovery Answer** box, answer the security question.
   b. In the boxes under **New Password**, type a suitable password.
4. Click **OK**.

##### 9.5.1.1.1. Lost or Forgotten Controller Password

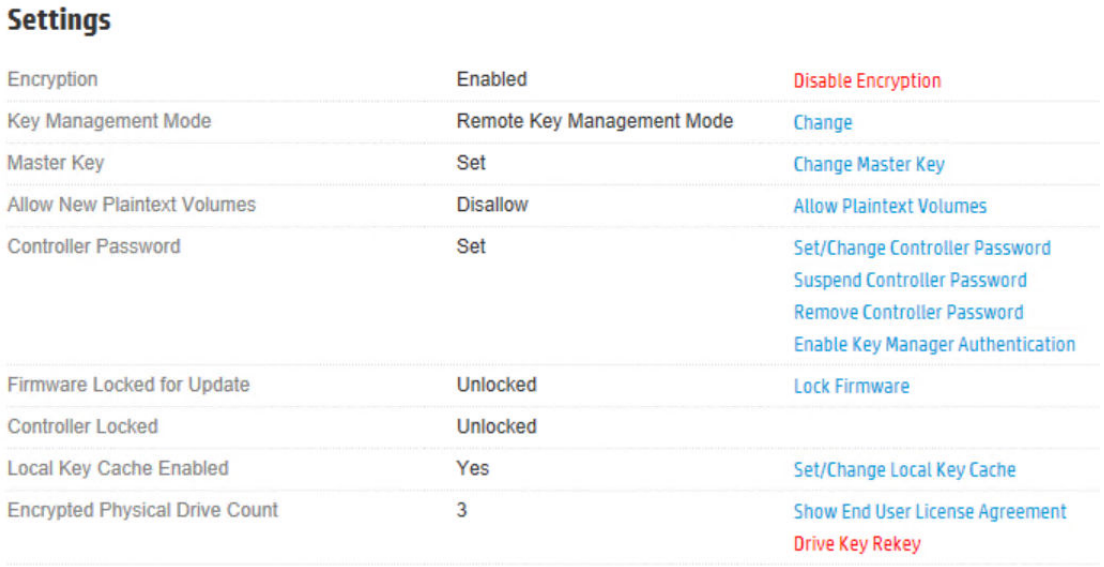
Use the controller password to protect the data in the event of a storage system theft. Once enabled, the controller does not unlock the encrypted volumes until the correct controller password has been provided. If the controller password is lost or forgotten, the controller remains locked and all the encrypted volumes are offline and inaccessible.

If the OS logical drive is encrypted, offline SSA is required to perform the following steps.

To clear the controller password, perform the following steps:

1. Open Encryption Manager .
2. Log in as the Crypto Officer.
3. Under **Settings**, locate **Controller Password**, and then click **Remove Controller Password**.

**Figure 9-14.** SSA Settings with Controller Password Options




4. A window appears, asking you to confirm that you want to remove the controller password. Click **Yes**.
5. Click on **Change Master Key** and type the Master Encryption Key used for encryption. For more information, see Changing the Master Encryption Key.
6. Enable Secure Encryption, then reboot the server.

Volumes appear online and are available.

#### 9.5.1.1.2. Lost or forgotten Master Key

**CAUTION** It is recommended to maintain a backup of the Master Encryption Key in a secure location. In some instances, it is possible that a missing key renders your data inaccessible. If operating Secure Encryption in Remote Key Management Mode, it recommended that you back up the ESKM regularly.

**Local Mode**

While operating Secure Encryption in Local mode, securing the Master Encryption Key value is critical in accessing the encrypted logical drive data. If the controller requires replacement or if the physical drives are moved to another controller, a matching Master Key is required to gain access to the data. Master Keys are not recoverable if lost. If the Master Key is lost or forgotten, you must perform a data restore operation from the backup media to regain access to the data.

**Remote Mode**

This section describes locating a lost or forgotten Master Encryption Key in the Remote mode using ESKM and iLO.

### Locating the Key using the ESKM

To locate a lost or forgotten Master Encryption Key using the ESKM and for more information about key query, see the HPE Enterprise Secure Key Manager - Document List in the Resources.

**Note:**

If you cannot locate the Master Key name, it may have been accidentally deleted from the ESKM. You may be able to locate the key by using an ESKM backup.

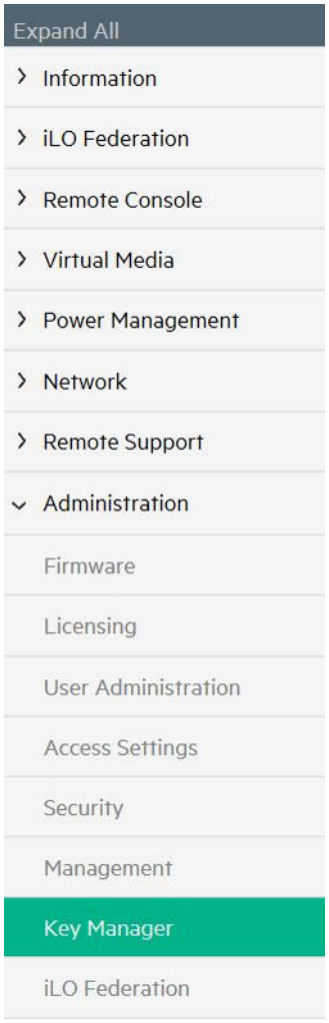
## Locating the Key using iLO

iLO utilizes an event log to list the recent key activity. If the lost or forgotten key was recently modified, it might appear in the event log.

To locate a lost or forgotten Master Encryption Key using iLO, perform the following steps:

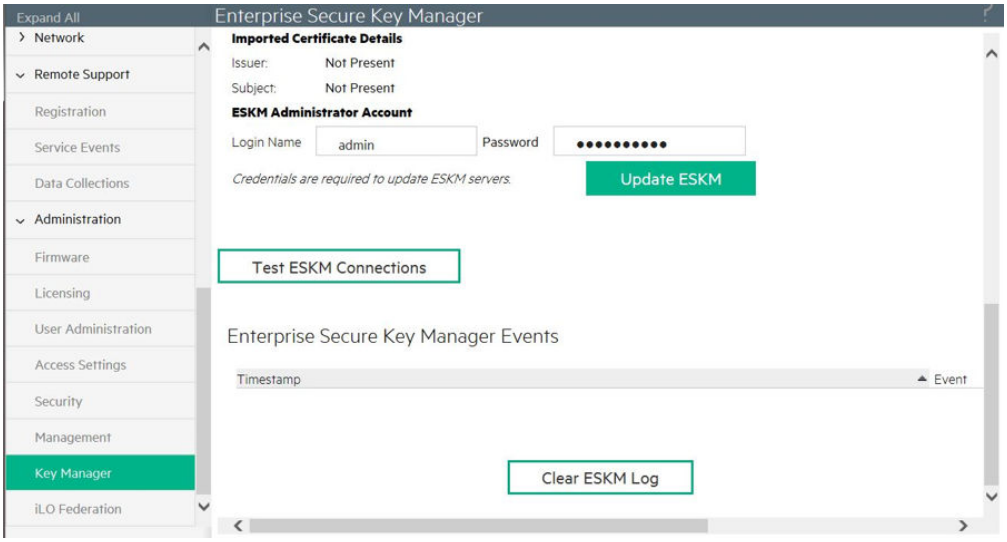
1. Log in to iLO using your server's credentials.
2. From the left side panel, expand the **Administration** menu.

**Figure 9-15.** iLO Key Manager




3. Click **Key Manager**. The **Enterprise Secure Key Manager Events** appears at the bottom of the screen. Review the event log for the missing key.

**Figure 9-16.** iLO Key Manager Event Log




#### 9.5.1.2. Forgotten Which Master Key Goes with Which Drive

Recovery of the Master Encryption Key name corresponding to a specific set of drives is possible when operating Secure Encryption in Remote Key Management mode.

To recover the Master Encryption Key name, perform the following steps:

1. Log in to the ESKM.
2. Run a key query with the following search parameters:
   a. **Choose Keys Where** list: click **Custom: Server_Name**. Two new fields appear.
   b. In the second list menu, select **Equals**.
   c. In the third field, type the name of the server to be associated with the Master Encryption Key.
   d. Under **Custom Attributes**, click **Master_Key**.

#### 9.5.1.3. Logical Drives Remain Offline

If the cryptographic information is missing, the logical drives remain offline after a system start. General causes include a missing, incorrect, or inaccessible key. Restoring the cryptographic information to match the attached drives results in the appropriate access to the logical drive.

The following is a list of possible causes:

- The encryption is not enabled
- The matching Master Encryption Key is missing or incorrect
- The controller password was enabled but is not entered or is incorrect

#### 9.5.1.4. Master Key not Exporting

This issue occurs only in Remote Key Management mode. The problem appears as either a locked controller or as locked volumes.

The following is a list of possible causes:

- A network problem prevents key retrieval from the ESKM
- Lost or incorrect iLO configuration
- Missing or incorrectly configured Master Encryption Key

The following is a list of possible resolutions:

- Troubleshoot the network connection between iLO and the ESKM. For more information, see Testing the Connection Between iLO and the ESKM.
- Ensure that the Master Encryption Key exists. For more information, see Locating the Key using the ESKM.
- Ensure that the Master Encryption Key is in the correct group.

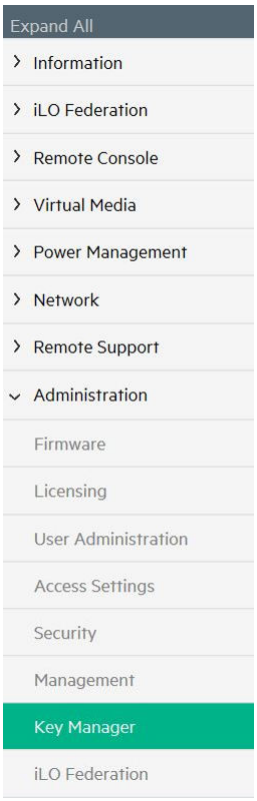
#### 9.5.1.5. Testing the Connection Between iLO and the ESKM

iLO connects and manages key exchanges between the controller and ESKM. If you suspect iLO has lost its connection to the ESKM, you can test the connection in iLO.

To test the connection between iLO and the ESKM, perform the following steps:

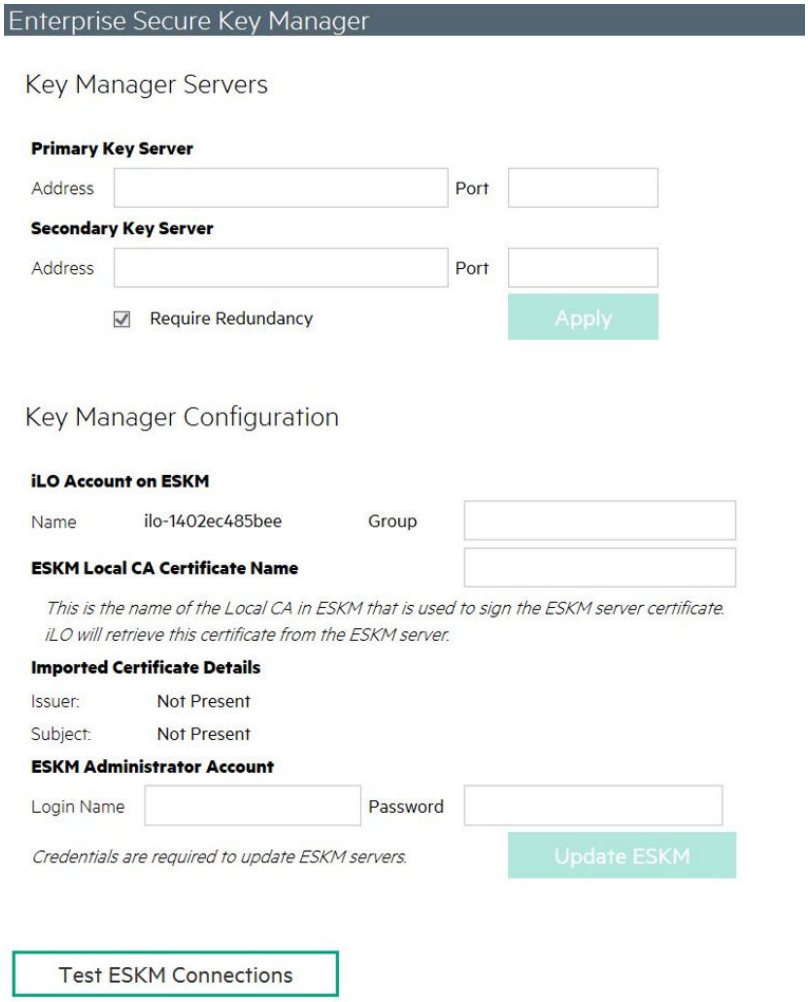
1. Log into iLO using your server's credentials.
2. From the left side panel, expand the **Administration** menu, and then click **Key Manager**.

**Figure 9-17.** iLO Key Manager



The following screen appears.

**Figure 9-18.** iLO Key Manager Configuration Window




3. Under **Key Manager Configuration**, click **Test ESKM Connections**:
   - If iLO is connected to the ESKM, a green checkmark appears indicating the key managers are accessible.
   - If the connection has been lost, you need to re-configure iLO to communicate with the ESKM. For more information, see Connecting iLO to the ESKM.

### 9.5.1.6. Potential Errors Encountered

The following table lists the errors that are encountered when configuring or operating Secure Encryption.

**Table 9-4.** Potential Errors Encountered

| Error | Description | Action |
|---|---|---|
| Remote key manager communication failure | Slot X Encryption Failure – Communication issue prevents the drive keys from being retrieved. Encrypted logical drives are offline. System may not boot. | To troubleshoot, see the **Key Manager** page in iLO interface. |
| Incorrect or missing Master Key on Remote key manager | Slot X Encryption Failure – Master Encryption Key is incorrect or not retrieved from ESKM. Encrypted logical drives may be offline. System may not boot. | Correct the problem on the ESKM. |

**Table 9-4.** Potential Errors Encountered (continued)

| Error | Description | Action |
|---|---|---|
| Volume Key decryption failure | Invalid Drive Encryption Keys on ESKM. Encrypted logical drives may be offline. System may not boot. | Restore the correct version of the Drive Encryption Key on the ESKM. |
| Unable to establish communication with controller | Communication issue prevents keys from being retrieved. Dependent encrypted logical drives are offline. System may not boot. | Reset the controller by rebooting the server. |
| Missing local Master Key | Imported encrypted logical drives are offline; the matching local Master Encryption Key is required. System may not boot. | Use SSA to enter the local Master Encryption Key. |
| Controller password failure | All the encrypted local drives are offline due to failure to enter proper controller password. | Reboot the server and enter the proper controller password, or unlock the controller using SSA. |
| Controller encryption not enabled | Encrypted logical drives are present but encryption is not yet enabled. Encrypted logical drives are offline. | Use Smart Storage Administrator to enable encryption. |
| Encryption parameters not set | Encryption is enabled for the controller but the Master Encryption Key name is not set. | Use Encryption Manager to set the Master Key name for the controller and reboot. |
| Controller/logical drive encryption type mismatch | Key Management mode mismatch between controller and drives. Dependent encrypted drives offline. | Use Encryption Manager to match Key Management modes. For more information, see Importing Drives with Different Master Keys. |
| Encryption failure - unsupported system ROM detected | Unsupported system ROM is detected. Encrypted logical drives may be offline. System may not boot. | Update the system ROM to a version supporting encryption. |
| Encrypted logical drives on non-encrypting controller | Encrypted logical drives are offline. Encryption feature is not available on this controller. | Move drives to a controller with encryption support or delete the logical drives. |
| Encryption failure - unsupported iLO firmware detected | Unsupported iLO firmware is detected. Encrypted drive may be offline. System may not boot. | Update iLO firmware to a version supporting encryption. |
| NVRAM failure | Non-volatile storage is corrupted. Critical Security Parameters erased per policy. Encrypted drives are offline. | Use SSA to reestablish CSPs. |
| Encryption engine self-test failure | Encryption engine hardware failure. Encrypted logical drives are offline until the problem is corrected. | Replace the controller to bring encrypted drives online. |
| Unable to create a plaintext volume | While logged into the system, you are unable to create a plaintext volume. | Verify that Encryption Manager has been set to allow the creation of future plaintext volumes. |

### 9.5.1.7. Clearing the Encryption Configuration

**Important:** Clearing all encryption settings clears all secrets, keys, and passwords from the controller. Secure Encryption is returned to a factory-new state.

To clear all encryption settings, perform the following steps:

1. Clear the controller. For more information, see Clearing the Controller.

   **Important:** Clearing the controller is not necessary if there are no encrypted drives present or if SSA is operating in an offline mode.

2. Log in to Encryption Manager. For more information, see Logging into Encryption Manager (CBE).
3. Under **Utilities**, click **Clear Encryption Configuration**.

**Figure 9-19.** SSA Utilities




4. A prompt appears, indicating all encryption settings is cleared from the controller. To continue, click **Clear**.

### 9.5.2. Common Issues with SED

This section describes troubleshooting for the common issues such as lost or forgotten Master Keys and removing controller password.

#### 9.5.2.1. Lost or Forgotten Master Key

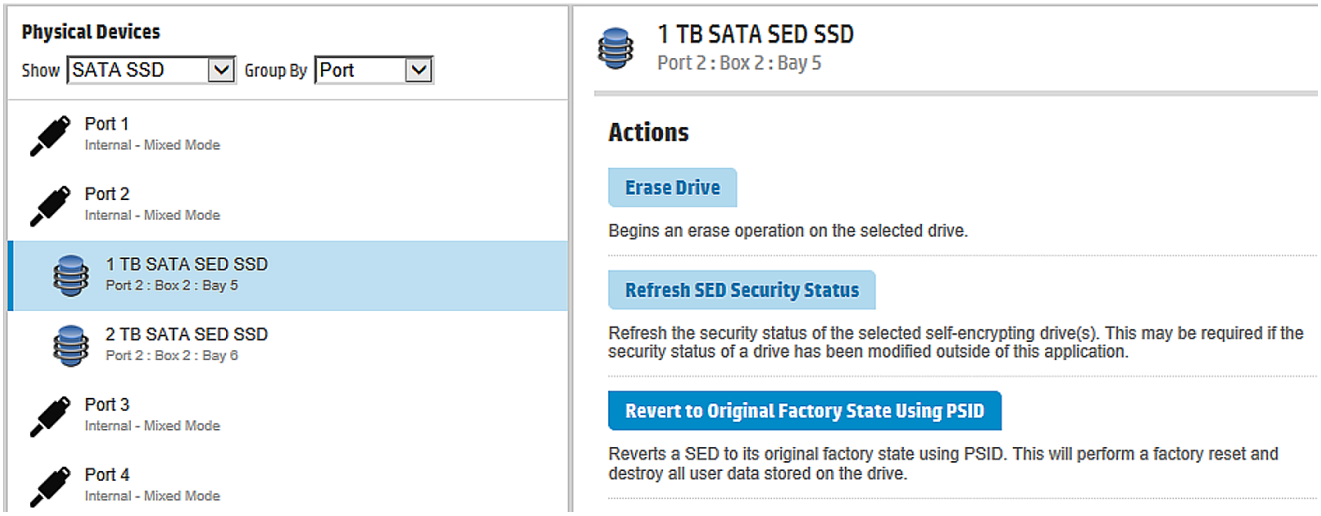
If you have lost or forgotten the master key, you can use the PSID to revert the drive to its original factory state, effectively erasing all the data on the drive.

If the OS logical drive is encrypted, perform the following steps to erase the data:

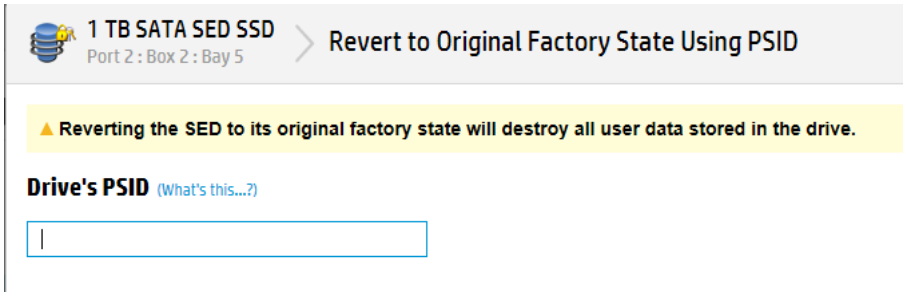
To reset a drive to original factory state, perform the following steps:

1. Open Encryption Manager.
2. Under **Controller Devices**, click **Physical Devices**.
3. Select the SED drive that you want to reset to Original Factory State. Click **Revert to Original Factory State using PSID**. A new window appears.

**Figure 9-20.** Revert to Original Factory State




4. Enter the PSID mentioned in the drive.




**Note:** PSID is a unique, static, 32-character key embedded in the drive.

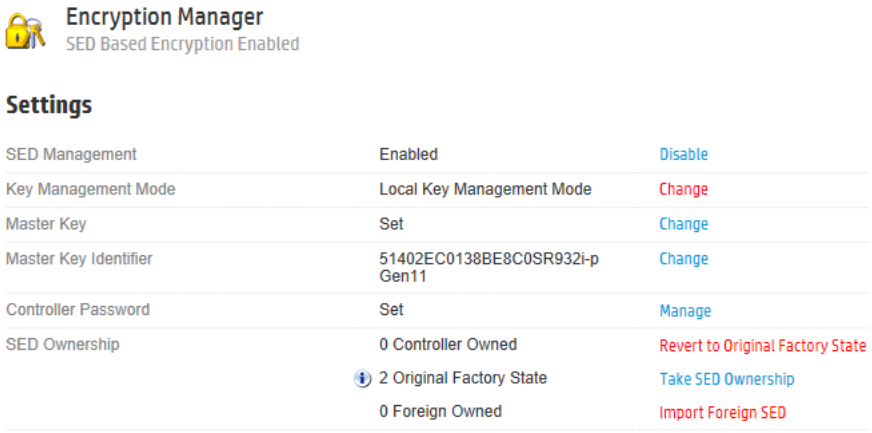
5. When prompted to erase data, click **Yes**.

### 9.5.2.2. Removing Controller Password

To remove the controller password, perform the following steps:

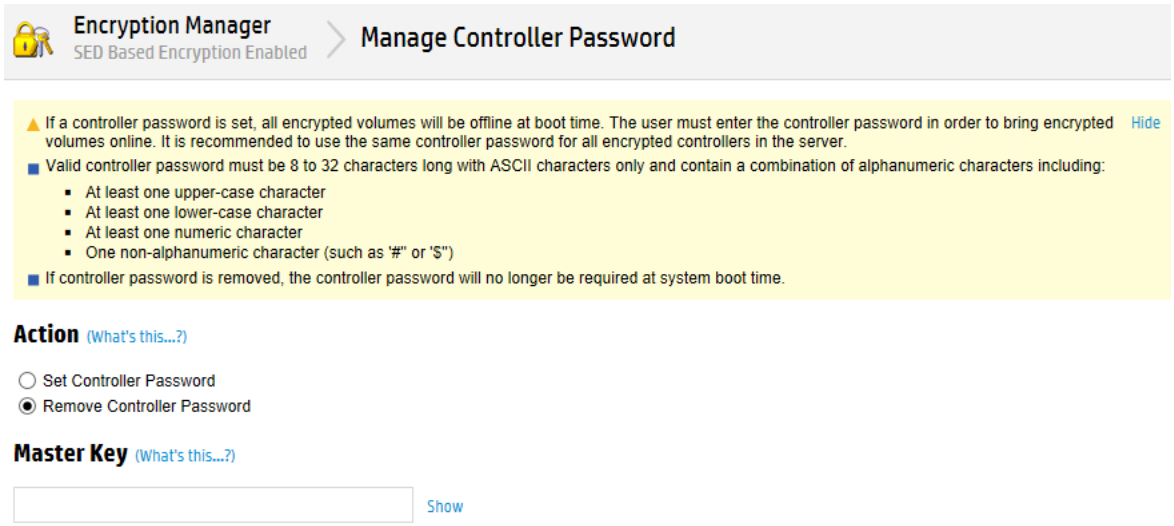
1. Open Encryption Manager.
2. Under Settings, locate **Controller Password**,and then click **Manage**.

**Figure 9-21.** SSA Settings—SED Encryption Setup




3. Under Action, click **Remove Controller Password**. In the **Master Key** box, type a suitable value.

**Figure 9-22.** Remove Controller Password




4. Click **OK**.

# 10. Glossary

**Active Spare**
An Active Spare is a spare that is currently in use by an Array that contains a failed Physical Drive. A spare may also become an Active Spare if the current Active Spare fails and there are multiple spares available.

**Adapter ID**
A unique identification code for a host controller that is printed on the controller. This is also called a World-Wide Name (WWN).

**Anti-Freeze**
Anti-Freeze is a Sanitize Lock setting that prevents physical disks from being frozen. This may require the physical disks to be power cycled, hot-plugged, or complete a sanitize operation.

**Array**
A group of physical drives configured into one or more logical drives. Arrayed drives have significant performance and data protection advantages over non-arrayed drives.

**Auto Replace Drive**
An auto replace drive will permanently take over for a failed data drive. As soon as a data drive fails, the spare drive becomes a data drive and the data drive becomes a spare drive. The advantage of auto replace drives is that the information in the failed data drive only has to be recovered once. This type of spare cannot be shared between multiple arrays.

**Boot Controller**
The boot controller is the first controller that the system looks at to find a bootable logical drive/volume after a system power-on.

**Bootable Logical Drive/Volume**
A bootable logical drive/volume is a logical drive/volume that the system can attempt to boot from after a system power-on. A controller or storage system can have up to two bootable logical drives/volumes, where one is a primary boot logical drive/volume and the other a secondary boot logical drive/volume. When the system looks at a controller or storage system for a boot logical drive/volume, it will first attempt to boot from a primary boot logical drive/volume, and if that fails, then it will attempt to boot from a secondary boot logical drive/volume

**Cache Hit Rate**
The calculated ratio of the count of the number of times the cache was successfully read from or written to over the total number of times the cache was accessed

**Cache Hits**
The count of how many times the cache was successfully read from or written to.

**Cache Line Size**
The cache line size is the data block size used by SSD caching. It can impact the cache performance and maximum size supported. A larger cache line size can support a larger maximum cache size. Some controllers may support only the default option which is 64 KiB.

**Cache Misses**
The count of how many times the cache was unable to be read from or written to.

**Cache Write Policy**
How SmartCache handles data writes when caching. There are 2 different policies: Write-back and write-through

**Caching**

An internal part of the controller that dramatically improves performance of disk read and write operations by providing a buffer. Backup power sources, ECC memory and, in some systems, flash memory, protect the data.

**Capacity Expansion**
A feature that allows an increase in storage capacity of a drive array with the addition of one or more physical drives to the array. With the added space on the array, one or more new logical drives can be created. This feature is available only on controllers that support expansion.

**Connection Name**
A user-defined name for a connection from a server to a controller. The connection is made by means of a host controller that is installed in the server. The Connection Name is used as a convenient way to identify a connection for a controller instead of relying entirely on the Adapter ID of the host controller.

**Controller key**
A key created by the controller and permanently saved to the Remote Key Manager after being wrapped by the Master Encryption Key. This key is used on a temporary basis to alleviate potential bottlenecks to the Remote Key Manager during volume creation/change events. Use of a Controller Key is on a temporary basis only and is ultimately transitioned through a rekey operation to the appropriate Drive Encryption Key.

**Controller-secured region**
The section of a device where data and Critical Security Parameters can exist in an unencrypted format. This boundary must be secured against tampering as acquiring this sensitive data may result in unauthorized access to data.

**Critical Security Parameters (CSPs)**
An industry standard term referring to security related information such as keys, passwords, and so forth, whose disclosure would compromise an encrypted system.

**Crypto Officer**
Personnel who have permission to access the full range of encryption functions available on the controller. This includes turning encryption on and off, resetting keys, importing Master Encryption Keys, and so forth.

**Dedicated Spare Drive**
A dedicated spare drive temporarily takes over for a failed data drive. When the failed data drive is replaced, the spare drive becomes idle again. This type of spare drive can be shared between multiple arrays.

**Drive array**
The group of physical drives containing a logical volume.

**Drive encryption key**
Key generated by the SmartRAID controller for each physical drive that contains at least one encrypted logical drive. Use the Drive Encryption Key for each physical drive is to encrypt (wrap) the Volume Encryption Keys for all of the logical drives resident on that physical drive.

**Drive key caching**
In Remote mode, the Drive Encryption Keys are typically stored on the Remote Key Manager. However, it is possible to enable the controller to cache all of these Drive Encryption Keys necessary to decrypt attached logical drives within the controller-secured region. This option is available to the user through SSA.

**Encrypted data**
Data that has been encrypted through the use of an encryption key.

**Encryption Setup Types**

During encryption initial setup, you can choose one of the two setup types. The Full Setup lets you customize configurations and make changes later on. The Express Local Encryption only allows local key management mode and no changes can be made afterwards.

**Encryption User**
The Encryption User is an encryption user account with restricted access to encryption settings. The account is enabled when a password is set by the Crypto Officer.

**ESKM**
Enterprise Secure Key Manager

**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%.

**Express Local Encryption**
One of the setup types for encryption. Express Local Encryption is a simple way to set up the encryption in Local Key Management Mode. It takes all the default settings. You don't have crypto officer account or master key. You cannot make any configuration changes after the setup is done unless you clear the encryption configuration and start over.

If you have existing logical drive with plaintext data, it provides your with the only chance to covert them to encrypted data during the setup. The other setup type is Full Setup.

**Failure Spare Activation**
This spare activation mode is the default for controllers. Spare drives can be assigned to arrays containing at least one fault tolerant logical drive. When a physical drive fails on a fault tolerant logical drive, the assigned spare will become active.

**FIPS**
Federal Information Processing Standard

**Flexible Latency Scheduler**
Flexible Latency Scheduler (FLS) is a controller option where the controller can re-prioritize I/O requests to prevent some requests to HDDs from timing out. Under normal operation (when FLS is disabled, or in controllers that don't support FLS), the controller will sort incoming requests in order to minimize the amount of travel for the HDD's read heads (A.K.A. elevator sorting). This strategy works well for workloads that access sequential data, or workloads that require multiple requests from localized sectors in the drive. For highly random workloads, such as transaction processing, some requests will end up on the wrong side of the disk platter and, due to their high latency, will be marked as timed out. When FLS is enabled, it will detect these high-latency requests and apply a cutoff value, after which it will suspend elevator sorting and service the request right away.

**Foreign Owned**
For SED Based Encryption, a SED is foreign owned when its master key is different from the masterkey of the current controller. This can happen when: The SED was migrated from a different controller; The SED was previously owned by the connected controller but was removed and the master key changed during that time.

You can use Import Foreign SED operation to take the ownership of a foreign SED.

**Freeze**
Freeze is a Sanitize Lock setting that causes the physical disks to transition to a frozen state and will cause any subsequent sanitize operation commands to be aborted. This may require the physical disks to be power cycled, hot-plugged, or complete a sanitize operation

**Full Setup**
One of the setup types for encryption. Full Setup lets you fully configure encryption settings. You can set up crypto officer and user accounts, Key Management Mode, Master Key, etc. After initial setup,

you can log into Encryption Manager and make changes to the configuration. The other setup type is Express Local Encryption.

**GUID Partition Table**
GUID Partition Table (GPT) is a scheme for the layout of the partition table on a logical drive. GPT was introduced as a part of the Extensible Firmware Interface (EFI) initiative. It provides a more flexible mechanism for partitioning drives than the older Master Boot Record (MBR) partitioning scheme, which restricts a disk partition's size to a maximum of 2 TiB (= $2^{41}$ bytes). GPT drives can grow to a very large size. In theory, the maximum size can be up to $2^{64}$ x 512 = $2^{73}$ bytes.

**HBA Mode**
Host Bus Adapter (HBA) Mode allows drives attached to a controller to be directly accessible from the operating system. Most advanced array operations are not available in HBA Mode.

**Notes:**

- The configuration must be cleared and encryption must be disabled before HBA Mode can be enabled.
- This functionality is only available on supported controllers.
- For newer controllers, this feature has been replaced by port mode (if available).

**HIPAA**
Health Insurance Portability and Accountability Act

**HITECH**
Health Information Technology for Economic and Clinical Health

**Local Key Cache**
Encryption keys are stored in a cache locally to allow access to encrypted volumes when the remote key server is offline.

**Note:**
This only applies when encryption is operating in Remote Key Management Mode.

**Local Key Management Mode**
Encryption keys are locally generated using the Master Key.

**SSA**
Smart Storage Administrator

**iLO 4**
Integrated Lights-Out 4

**Local Master Encryption Key**
The equivalent of a Master Encryption Key in Local mode. The Local Master Encryption Key name is stored in non-volatile memory within the controller-secured region and used to generate a Local Master Encryption Key for wrapping the Drive Encryption Keys.

**Logical Drive**
An equal area from all physical drives in a drive array grouped together logically to act as a single hard drive. Logical drives are configured with software utilities to enhance the performance and usability of drive arrays.

**Logical Drive Extension**
This allows you to increase the size of an existing logical drive without disturbing the data on the logical drive. If an existing logical drive is full of data, you can extend the logical drive when there is free space on the array. If there is no free space on the array, you can add drives to the array and proceed to extend the logical drive. This feature is only available for certain controllers and should only be used with certain operating systems.

**Master Boot Record**
Master Boot Record (MBR) is the 512-byte boot sector that is the first sector of a partitioned logical drive. It contains Master Partition Table (MBR Partition Table) and a program the BIOS uses to boot an OS from the drive.

**Master Encryption Key**
A two-part key established on the Remote Key Manager. This key consists of both a name and a value. The name consists of a maximum of 64 characters and is used to uniquely identify this key to all controllers within a given Security Domain. The Master Encryption Key value is a 256-bit quantity used by controllers to wrap Drive Encryption and Controller Keys for secure storage on the controller and import into the Remote Key Manager.

**Maximum Boot Size**
Max Boot or Maximum Boot Size determines the number of sectors used for the logical drive. When Max Boot is disabled, the logical drive is created with 32 sectors per track. In this configuration, the largest boot drive which can be created is 4 GiB. With Max Boot enabled, the controller creates the logical drive with 63 sectors per track which will allow you to create a boot drive which is up to 8 GiB in size. We suggest only enabling Max Boot on the drive from which you will boot your server, as a slight performance gain is seen using 32 sectors per track.

The maximum boot size option is initially disabled. Disabling maximum boot size means that the logical drive will report the default of 32 sectors per track to BIOS calls (int13h). Enabling the maximum boot size increases the number of sectors reported in BIOS calls to the maximum of 63 in order to increase the number of blocks available. Enabling maximum boot size may be necessary to create large boot partitions for some operating systems. For example, enabling maximum boot size on a logical drive in Windows NT 4.0 allows you to create a bootable partition with a maximum size of 8 GiB, instead of the 4 GiB maximum size allowed when maximum boot size is disabled. When a logical drive larger than 255 GiB is created, a sector size of 63 is reported to BIOS calls regardless of whether or not the maximum boot size is enabled.

**Maximum Drive Request Queue Depth**
This is the maximum number of physical drive requests that the controller will submit to a drive at a given time.

**MBR Partition Table**
The MBR Partition Table is a scheme for creating multiple slices (partitions) of data within a logical drive. When a logical drive has been partitioned with the MBR partition table scheme, the master boot record (MBR) contains a table that defines the start addresses and lengths of up to four partitions. MBR partition tables suffer from the limitation that both the partition length and partition start address are stored as 32-bit quantities. Because the sector size is $2^9$ (= 512) bytes, this implies that neither the maximum size of a partition nor the maximum start address (both in bytes) can exceed $2^{32} \times 2^9 = 2^{41}$ bytes = 2 TiB. If the logical drive is larger than 2 TiB, only the first 2 TiB will be visible. For drives greater than 2 TiB in size, it is recommended to use a different partitioning scheme, such as GPT. Note that the BIOS in most servers is only capable of booting from a logical drives partitioned using the MBR Partition Table scheme.

**Migration**
A feature that allows you to change the fault tolerance level or strip size of a configured logical drive without incurring any data loss.

**Note:**
This functionality is only available on supported controllers.

**Mixed Mode**
Mixed Mode combines features from both RAID Mode and HBA Mode. It supports array operations and also allows the unassigned drives accessible directly from the OS.

**Multi-Actuator Physical Drive**

Multi-actuator drive contains two or more independent actuators that transfer data concurrently. It enables concurrent I/O streams to and from the host for increased performance by enabling parallelism of data flows in and out of a single hard drive.

Each actuator is mapped to a separate logical unit and assigned a Logical Unit Number (LUN).

**Note:**
This functionality is only available on supported controllers.

**NIST**
National Institute of Standards and Technology

**NPG**
Number of Parity Groups

**NVRAM**
Nonvolatile memory

**OFS**
Original Factory State

**Online Firmware Activation**
Feature with the ability to flash controller firmware and, once successfully completed, activate immediately without the need to reset or reboot the operating system or server. The current controller firmware and operating system driver must both support Online Firmware Activation.

**Online Recovery Server**
A controller that has been set by the System Configuration Utility to Online Recovery Server mode is a controller that has the ability to dynamically move storage devices from a failed server to an active server. In effect, the storage devices are hot plugged from one system and hot plugged into the new one.

**Online Spare**
A physical drive used in RAID 1/RAID 1+0 - Drive Mirroring, RAID 4 - Data Guarding, RAID 5 - Distributed Data Guarding, and RAID 6 - Advanced Data Guarding to provide drive replacement for a failed drive without user intervention. The spare immediately replaces a failed drive as soon as the failure occurs. The controller automatically begins rebuilding the data from the failed drive on the spare to return to a fault tolerant state. The failed drive can be replaced while the system is operating at top performance. The drawback is that the drive is not used while inactive and this reduces the amount of usable storage capacity.

**Parallel Surface Scan Count**
The number of volumes on which the controller will perform Surface Scan Analysis in parallel.

**Parity Group**
A Parity Group (PG) is a sub-volume that is a member of a compound RAID volume. RAID 50 combines RAID levels by creating multiple RAID 5 parity groups and then striping (RAID 0) data across all groups. RAID 60 combines RAID levels by creating multiple RAID 6 parity groups and then striping (RAID 0) data across all groups.

**Parity Initialization**
RAID levels that use parity (RAID 5, RAID 6, RAID 50, and RAID 60) require that the parity blocks be initialized to valid values. Valid parity data is required to enable enhanced data protection through background surface scan analysis and higher performance write operations.

There are two initialization methods available:

The default method initializes parity blocks in the background while the logical drive is available for access by the operating system.

The Rapid Offline Initialization method works by overwriting both the data and parity blocks in the foreground. The logical drive remains invisible and unavailable to the operating system until the parity initialization process completes.

**Parity RAID Degraded Mode Performance Optimization**
This setting applies to RAID 5/RAID 50/RAID 6/RAID 60 volumes in degraded mode only. Enabling this setting directs the controller to attempt to improve performance of large read requests by buffering physical drive requests. Disabling this feature forces the controller to read from the same drives multiple times.

**PCI-DSS**
Payment Card Industry Data Security Standard

**Plaintext**
Data in unencrypted form.

**Plaintext Volume**
With encryption enabled, any volume that is not encrypted is labelled as a plaintext volume.

**Physical Drive**
A hard disk drive, which can be connected to a controller and used for storage of data.

**Physical Drive Request Elevator Sort**
The elevator sort is a feature in which the controller can re-order requests to a physical drive in order to minimize the amount of seeking the drive must perform. Enabling the request elevator sort improves seek times and disabling the elevator sort improves throughput.

**Port**
A synonym for a SCSI bus or channel on a controller. Physical drives are connected to a controller through a port.

**Port Discovery Protocol**
The protocol used by a port to discover a connected backplane. Available port discovery protocols are: Auto Detect, UBM, SGPIO, VPP, and Direct-Attached Cable.

Auto Detect—The controller firmware attempts to automatically detect the discovery protocol of the backplane attached to the port.

UBM—The controller firmware uses the UBM protocol to communicate with the backplane attached to the port.SGPIOThe controller firmware uses SGPIO to communicate with the backplane attached to the port.

VPP—The controller firmware uses the VPP protocol to communicate with the backplane attached to the port.

Direct-Attached Cable—The controller firmware uses the direct-attached cable protocol to support cable attached drives. Number of targets (drives) must be set to match the cable's capabilities.

**Notes:**

- If the port discovery protocol is not configured correctly, some features of the backplane do not function as expected.
- A reboot is required for the new port discovery protocol to take effect.

**Power On Hours**
Indicates the number of hours the SSD has been powered on for.

**PSID**
PSID stands for Physical Security Identifier. The PSID is unique to each drive. It is printed on the disk label and visible to anyone with physical access to the SED.

**Primary Boot Logical Drive/Volume**
When booting from a controller, the primary boot logical drive/volume is the first logical drive/volume that the system attempts to boot from.

**RAID 0**
RAID stands for Redundant Array of Inexpensive Disks. RAID 0 indicates that there is no fault tolerance method used. However, the data is striped across all physical drives in the array for rapid access.

If you select this option for any of your logical drives, you will experience data loss for that logical drive if one physical drive fails. However, because none of the capacity of the logical drives is used for redundant data, this method offers the best processing speed and capacity. You may consider assigning RAID 0 to drives that require large capacity and high speed, but pose no safety risk

**RAID 1**
RAID stands for Redundant Array of Inexpensive Disks. RAID 1 (drive mirroring) creates fault tolerance by storing two sets of duplicate data on a pair of disk drives.

If a drive fails, the mirror drive provides a backup copy of the files and normal system operations are not interrupted.

**RAID 1+0**
RAID stands for Redundant Array of Inexpensive Disks. RAID 1+0 (drive mirroring with striping) is a fault tolerance method that uses 50 percent of drive storage capacity to provide greater data reliability by storing a duplicate of all user data. Half the physical drives in the array are duplicated or "mirrored" by the other half.

RAID 1+0 first mirrors each drive in the array to another, then stripes the data across the mirrored pairs.

Drive mirroring creates fault tolerance by storing two sets of duplicate data on a pair of disk drives. There must be an even number of drives for RAID 1+0. This is the most costly fault tolerance method.

If a drive fails, the mirror drive provides a backup copy of the files and normal system operations are not interrupted. The mirroring feature requires a minimum of two drives, and in a multiple drive configuration (four or more drives), mirroring can withstand multiple simultaneous drive failures as long as the failed drives are not mirrored to each other.

**RAID 1 Triple**
RAID stands for Redundant Array of Inexpensive Disks. RAID 1 Triple creates fault tolerance by maintaining redundant copies of data using three disk drives. All three drives contain mirrored duplicated user data.

If a drive fails, the remaining drives provide backup copies of the files and normal system operations are not interrupted.

**RAID 10 Triple**
RAID stands for Redundant Array of Inexpensive Disks. RAID 10 Triple creates fault tolerance by maintaining redundant copies of data using at least six disk drives. Data is striped across two or more sets of RAID 1 Triple drives for rapid access.

If a drive fails, the remaining drives provide backup copies of the files and normal system operations are not interrupted.

**RAID 5**
RAID stands for Redundant Array of Inexpensive Disks. RAID 5 (distributed data guarding) is a fault tolerance method that stores parity data across all the physical drives in the array, which allows more simultaneous read operations and higher performance than RAID 4 - Data Guarding. If a drive fails, the controller uses the parity data and the data on the remaining drives to reconstruct

data from the failed drive. This allows the system to continue operating with a slightly reduced performance until you replace the failed drive.

RAID 5 requires an array with a minimum of 3 physical drives. The capacity of the logical drive used for fault tolerance depends on the number of physical drives in the array. For example, in an array containing 3 physical drives, only 33 percent of the total logical drive storage capacity is used for parity data; while a 14-drive configuration uses only 7 percent.

**RAID 50**
RAID stands for Redundant Array of Inexpensive Disks. RAID 50 is a fault tolerance method that combines the reliability of RAID 5 - Distributed Data Guarding with the increased performance of RAID 0 (striping). A RAID 50 volume is composed of two or more RAID 5 sub-volumes (parity groups) where data is striped across each parity group as if it were a single physical drive. Each RAID 5 parity group can sustain a single drive failure without incurring data loss.

RAID 50 requires an array with a minimum of 6 physical drives. The capacity of the logical drive used for fault tolerance depends on the number of physical drives and the number of parity groups in the array. For example, in an array containing 6 physical drives and 2 parity groups, only 33 percent of the total logical drive storage capacity is used for parity data; while a 28-drive, 2 parity group, configuration uses only 7 percent.

**RAID 6**
RAID stands for Redundant Array of Inexpensive Disks. This fault tolerance method provides the highest level of data protection. It is similar to RAID 5 in that parity data is distributed across all drives in the array, except that multiple separate sets of parity data are used in RAID 6, and the capacity of multiple drives is used to store the parity data. Assuming the capacity of two drives is used for parity data, the system will continue to operate even if two drives fail simultaneously, whereas RAID 4 and RAID 5 can only sustain failure of a single drive. The fault tolerance of RAID 6 configurations is actually higher than that of RAID 1+0 configurations because in RAID 1+0 there is a chance that two drives mirrored to each other will fail simultaneously.

RAID 6 read performance is similar to that of RAID 5, since all drives can service read operations. However, the write performance is lower with RAID 6 than with RAID 5, because parity data must be updated on multiple drives. Performance is further reduced in a degraded state.

RAID 6 requires an array with a minimum of 2+P physical drives, where P is the number of drives used for parity data; normally, P= 2. The percentage of total drive capacity used for fault tolerance is equal to the number of drives used for parity data divided by the total number of physical drives. For example, in an array containing a total of five physical drives (two of which are used for parity), 40 percent of the total logical drive storage capacity is used for fault tolerance. A 14-drive configuration (again using two drives for parity) uses only 14 percent of total capacity for fault tolerance.

**RAID 60**
RAID stands for Redundant Array of Inexpensive Disks. RAID 60 is a fault tolerance method that combines the reliability of RAID 6 - Advanced Data Guarding with the increased performance of RAID 0 (striping). A RAID 60 volume is composed of two or more RAID 6 sub-volumes (parity groups) where data is striped across each parity group as if it were a single physical drive. Each RAID 6 parity group can sustain up to two drive failures without incurring data loss.

RAID 60 requires an array with a minimum of 8 physical drives. The capacity of the logical drive used for fault tolerance depends on the number of physical drives and the number of parity groups in the array. For example, in an array containing 8 physical drives and 2 parity groups, 50 percent of the total logical drive storage capacity is used for parity data; while a 28-drive, 2 parity group, configuration uses only 14 percent.

**RAID 6/60 Alternate Inconsistency Repair Policy**
An inconsistency arises when, during a surface analysis scan, the controller detects that the parity information does not match the data present on the drives. Disabling the alternate repair policy

directs the controller to always update the parity information, leaving the data untouched. Enabling the alternate repair policy allows the controller to update the data on the drives based on the parity information. This behavior applies to RAID 6 and RAID 60 volumes only.

**RAID Mode**
RAID Mode allows controllers to perform array configuration operations on any drives attached to the controller. Any data on drives that were exposed to the OS in HBA Mode will not be available in RAID Mode.

**Notes:**

- This functionality is only available on supported controllers.
- For newer controllers, this feature has been replaced by port mode (if available).

**RAID Overhead**
A pre-defined space set aside for RAID redundant information on a logical drive.

**Rebuild Priority**
After a failed drive has been replaced, the level of priority that rebuilding the data from the failed drive should have over handling current requests from the operating system.

**Redundant Controllers**
A pair of controllers that have been installed into a system and share a single storage system. The controllers are interconnected either through an Inter-Controller Link (ICL) for 64-Bit or Extended PCI Controllers or internally for Fibre Channel Controllers.

The primary controller of the pair handles all communications and control of the storage system and its attached drives. If the primary controller is no longer able to issue read or write commands to these drives, the secondary controller assumes control.

**Remote Key Manager**
A server used to store, backup and retrieve keys for a group of controllers in a data center.

**Sanitize Lock None**
Sanitize Lock None is a Sanitize Lock Setting that causes the physical disks to transition to the "None" state. This means that the controller will not send freeze or anti-freeze commands to any drive. This may require the physical disks to be power cycled, hot-plugged, or complete a sanitize operation.

**SCSI**
Small Computer Systems Interface.

**SCSI ID**
A unique ID assigned to each SCSI device connected to the same SCSI channel. The ID number uniquely defines each peripheral device address and determines the device priority on the bus. ID 7 (SCSI controller) is the highest priority; ID 0 is the lowest.

**Secondary Boot Logical Drive/Volume**
If booting from the primary boot logical drive/volume fails, then the system attempts to boot from the logical drive/volume designated as the secondary boot logical drive/volume on that boot controller.

**SED**
A SED is a type of hard drive that automatically and continuously encrypts the data on the drive without any user interaction. If a SED becomes locked, the volumes on the array may become degraded or inaccessible. If this occurs, unlock the SED(s) and warm-boot the server.

**Smart Storage Administrator**
Smart Storage Administrator (SSA) is the software tool that performs configuration and diagnostic actions on Smart RAID/HBA controllers.

**SmartCache**
SmartCache is a feature that allows SmartCache cache to be assigned to logical drives for performance acceleration

**SmartCache Cache**
SmartCache cache is physical memory on a drive allocated for caching on its associated logical drive. This caching helps to accelerate performance.

**SSD I/O Bypass**
When SSD I/O Bypass is enabled for an array, and the Operating System driver supports the SSD I/O Bypass feature, the driver can bypass the hardware RAID stack and read (and write) information from (and to) the SSDs directly. This improves the latency of accessing information from the drives.

**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 data drive. The optimization process may take some time, during which the system may seem unresponsive. Some controllers may not support this option.

**SSD Utilization**
Indicates the percentage of the SSD that has worn out.

**Strip Size / Full Stripe Size**
A stripe is a collection of contiguous data that is distributed evenly across all physical drives in a logical drive. A strip represents the portion of a stripe that is written to a single physical drive. The strip size is selected to optimize the performance of the operating system. Strip size is synonymous with distribution factor. Note that in 8.50 and earlier versions of ACU, the term stripe size was used instead of strip size. This is a change of labeling and does not signify a change in functionality

**Stripe Size**
A stripe is a collection of contiguous data that is distributed evenly across all physical drives in a logical drive. The stripe size, or full stripe size, is the combined size of all strips across all physical drives, excluding parity/redundancy drives. For example, a RAID 5 logical drive in an array with 5 drives and a strip size of 256 KiB will have a stripe size of 1 MiB. The strip size (and consequently the stripe size) is selected to optimize the performance of the operating system.

**Surface Scan Analysis**
Surface Scan Analysis is a background process that scans hard drives for bad sectors in fault tolerant logical drives. In RAID 5, RAID 50, RAID 6, and RAID 60 configurations, Surface Scan also verifies the consistency of parity data. This process assures that you can recover all data successfully if a drive failure occurs in the future.

**Transformation Priority**
After choosing to modify an array or logical drive configuration, the level of priority that transformation should have over handling current operating system requests. This setting applies to the following operations: Extend Logical Drive, Migrate RAID/Strip Size, Expand Array, Shrink Array and Move Array. Note that not all of these operations are available on all controllers.

**Transient Drive**
A transient drive is a physical drive that is in transition from being a member of an array to being an unassigned drive as a result of a Shrink Array or Move Array operation.

**UEFI KMS**
UEFI KMS is a Key Management Service (KMS) protocol running on UEFI pre-boot system interface environment. This provides services to generate, store, retrieve and manage cryptographic keys.

**Usable Space**
Space on the array available to the user for logical drives.

If an array is created with different sized physical drives, some of the space on the larger capacity drives will not be usable.

**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.

**Write-Back**
A caching method where data is not copied to the data volume until absolutely necessary. Write-back may accelerate performance compared to the write-through policy because it reduces the number of write operations to data volumes. This performance improvement comes a slight risk that data may be lost if the cache volume fails.

**Write-Through**
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.

**Volume Encryption Key**
The key used in conjunction with hardware-based algorithms to perform the encryption of data resident on logical volumes.

# 11. Resources

This section provides the list of resources:

- Enterprise Secure Key Manager (ESKM)
  - HPE Enterprise Secure Key Manager - Document List

# 12. Revision History

The revision history describes the changes that were implemented in the document. The changes are listed by revision, starting with the most current publication.

**Table 12-1.** Revision History

| Revision | Date | Description |
|---|---|---|
| H | 07/26 | The following is a summary of changes in revision H of this document:<br>• Added a note in Configuring the Controller to Enable CBE (Local Key Encryption) |
| G | 05/25 | The following is a summary of changes in revision G of this document:<br>• Updated section Creating an Array |
| F | 02/25 | The following is a summary of changes in revision F of this document:<br>• Updated sections— Array and Controller Configuration and Accessing SSA in the Offline Environment<br>• Updated notes in the Unlocking the Controller Password (SED) |
| E | 10/2024 | The following is a summary of changes in revision E of this document:<br>• Updated the instances of references to 'Enterprise Secure Key Manager User Guide and Installation and Replacement Guide' to 'HPE Enterprise Secure Key Manager - Document List' in the following sections:<br>– Encryption Features<br>– Configuring Remote Key Management Mode<br>– Configuring the ESKM<br>– Locating the Key using the ESKM<br>• Added Resources |
| D | 06/2024 | The following is a summary of changes in revision D of this document:<br>• Added Expand an Array and Extend the Logical Drive |
| C | 12/2023 | The following is a summary of changes in revision C of this document:<br>• Updated the instances of 'maxCache' to 'SmartCache' in the following sections:<br>– Enable SmartRAID SmartCache<br>– Glossary |
| B | 11/2023 | The following is a summary of changes in revision B of this document:<br>• Updated section Configuring the ESKM<br>• Removed the following sections Logging in to the ESKM, Adding a User, Adding a Group, Assigning a User to a Group, Creating Keys, About Keys, Creating a Master Key, Placing a Key in a Group, Running a Key Query, Assigning a Key to a Group, Groups, Locating Groups Associated with a Drive, Query by Drive Serial Number, Query by Previous Server Name, and Running Queries<br>• Added Configuring the Controller to Enable SED (Remote Key Management Mode)<br>• Updated Table 8-4 to include Remote Key Management mode<br>• Updated the following sections to include a note on Remote Key Management SED setup:<br>– Unlocking the Controller Password (SED)<br>– Changing Master Key<br>– Changing Master Key Identifier<br>– Managing Controller Password<br>– Set the Controller Password<br>• Added a note in Importing Foreign SED |
| A | 02/2023 | Initial revision. |