



Managing Storage Adapters

This chapter includes the following sections:

- [Managing Storage Adapters, on page 1](#)
- [Managing the Flexible Flash Controller, on page 26](#)
- [Cisco Boot Optimized M.2 Raid Controller, on page 38](#)

Managing Storage Adapters

Self Encrypting Drives (Full Disk Encryption)

Cisco IMC supports self encrypting drives (SED). A special hardware in the drives encrypts incoming data and decrypts outgoing data in real-time. This feature is also called Full Disk Encryption (FDE).

The data on the drive is encrypted on its way into the drive and decrypted on its way out. However, if you lock the drive, no security key is required to retrieve the data.

When a drive is locked, an encryption key is created and stored internally. All data stored on this drive is encrypted using that key, and stored in encrypted form. Once you store the data in this manner, a security key is required in order to un-encrypt and fetch the data from the drive. Unlocking a drive deletes that encryption key and renders the stored data unusable. This is called a Secure Erase. The FDE comprises a key ID and a security key.

The FDE feature supports the following operations:

- Enable and disable security on a controller
- Create a secure virtual drive
- Secure a non-secure drive group
- Unlock foreign configuration drives
- Enable security on a physical drive (JBOD)
- Clear secure SED drives
- Clear secure foreign configuration

Scenarios to consider While Configuring Controller Security in a Dual or Multiple Controllers Environment



Note Dual or Multiple controllers connectivity is available only on some servers.

Controller security can be enabled, disabled, or modified independently. However, local and remote key management applies to all the controllers on the server. Therefore security action involving switching the key management modes must be performed with caution. In a scenario where both controllers are secure, and you decide to move one of the controllers to a different mode, you need to perform the same operation on the other controller as well.

Consider the following two scenarios:

- Scenario 1—Key management is set to remote; both controllers are secure and use remote key management. If you now wish to switch to local key management, switch the key management for each controller and disable remote key management.
- Scenario 2—Key management is set to local; both controllers are secure and use local key management. If you now wish to switch to remote key management, enable remote key management and switch the key management for each controller.

If you do not modify the controller security method on any one of the controllers, it renders the secure key management in an unsupported configuration state.

Enabling Controller Security

This option is available only on some C-series servers.

Before you begin

You must log in with admin privileges to perform this task.

- Step 1** In the **Navigation** pane, click the **Storage** menu.
- Step 2** In the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.
- Step 3** In the **Controller Info** area, click **Enable Drive Security**.
- Step 4** In the **Enable Drive Security** dialog box, update the following fields:

Name	Description
Controller Security field	Indicates that the controller is disabled.
Key Management field	Indicates whether the key is remotely managed or locally managed. This can be one of the following: <ul style="list-style-type: none"> • Remote Key Management radio button— Controller security key is configured or managed using the remote KMIP server. <p>Note If you choose this option, you do not have to specify the existing security key but you have to provide the key ID and the security key for local management.</p> • Local Key Management radio button— Controller security is configured locally.

Name	Description
Security Key Identifier field	The current key ID.
Security Key field	Security key used to enable controller security. If you wish to change the current security key, enter the new key here. Note Once you change the security key, a Secure Key Verification pop-up window appears where you need to enter the current security key to verify it.
Confirm Security Key field	Re-enter the security key.
Suggest button	Suggests the security key or key ID that can be assigned.

Step 5 Click **Save**.

This enables controller security.

Modifying Controller Security

This option is available only on some C-series servers.

Before you begin

- You must log in with admin privileges to perform this task.
- You must have first enabled controller security to modify it.

Step 1 In the **Navigation** pane, click the **Storage** menu.

Step 2 In the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.

Step 3 In the **Controller Info** area, click **Modify Drive Security**.

Step 4 In the **Modify Drive Security** dialog box, update the following fields:

Name	Description
Controller Security field	Indicates whether or not controller security is enabled. This can be one of the following: <ul style="list-style-type: none"> • Enabled— Controller security is enabled. • Disabled— Controller security is disabled.
Security Key Identifier field	The current key ID.
Security Key field	Security key used to enable controller security. If you wish to change the current security key, enter the new key here. Note Once you change the security key, a Secure Key Verification pop-up window appears where you need to enter the current security key to verify it.

Name	Description
Confirm Security Key field	Re-enter the security key.
Modify Security Key check box	Note This option appears only for remote key management. If you select this option the security key on the KMIP server is modified.
Suggest button	Suggests the security key or key ID that can be assigned.
Save button	Saves the data.
Cancel button	Cancels the action.

Step 5 Click **Save**.

This modifies the controller security settings.

Disabling Controller Security

This option is available only on some C-series servers.

Before you begin

- You must log in with admin privileges to perform this task.
 - You must have first enabled controller security to disable it.
-

Step 1 In the **Navigation** pane, click the **Storage** menu.

Step 2 In the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.

Step 3 In the **Controller Info** area, click **Disable Drive Security**.

Step 4 Click **OK** in the confirmation pop-up window.

This disables controller security.

Switching Controller Security Between Local and Remote Key Management

This task allows you to switch controller security from local management to remote management, and from remote to local management.

Before you begin

You must log in with admin privileges to perform this task.

Step 1 In the **Navigation** pane, click the **Storage** menu.

Step 2 In the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.

- Step 3** In the **Controller Info** area, to switch the controller security from remote to local management, click **Switch to Local Key Management**.
- Note** When you switch from remote to local key management, ensure that you disable KMIP secure key management first.
- Step 4** (Optional) Similarly, if you want to switch the controller security from local to remote management, click **Switch to Remote Key Management**.
- Step 5** Click **OK** to confirm.
-

Creating Virtual Drive from Unused Physical Drives

Before you begin

You must log in with admin privileges to perform this task.

- Step 1** In the **Navigation** pane, click the **Storage** menu.
- Step 2** In the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.
- Step 3** In the **Actions** area, click **Create Virtual Drive from Unused Physical Drives**.
The **Create Virtual Drive from Unused Physical Drives** dialog box displays.
- Step 4** In the **Create Virtual Drive from Unused Physical Drives** dialog box, select the RAID level for the new virtual drives:
This can be one of the following:
- **Raid 0**—Simple striping.
 - **Raid 1**—Simple mirroring.
 - **Raid 5**—Striping with parity.
 - **Raid 6**—Striping with two parity drives.
 - **Raid 10**—Spanned mirroring.
 - **Raid 50**—Spanned striping with parity.
 - **Raid 60**—Spanned striping with two parity drives.
- Step 5** In the **Create Drive Groups** area, choose one or more physical drives to include in the group.
Use the >> button to add the drives to the **Drive Groups** table. Use the << button to remove physical drives from the drive group.

Note

- The size of the smallest physical drive in the drive group defines the maximum size used for all the physical drives. To ensure maximum use of space for all physical drives, it is recommended that the size of all the drives in the drive group are similar.
- Cisco IMC manages only RAID controllers and not HBAs attached to the server.
- You must have multiple drive groups available to create virtual drives for certain RAID levels. While creating drives for these RAID levels, the create drive option is available only if the required number of drives are selected.

Step 6

In the **Virtual Drive Properties** area, update the following properties:

Name	Description
Virtual Drive Name field	The name of the new virtual drive you want to create.
Read Policy drop-down list	The read-ahead cache mode.
Cache Policy drop-down list	The cache policy used for buffering reads.
Strip Size drop-down list	The size of each strip, in KB.
Write Policy drop-down list	This can be one of the following <ul style="list-style-type: none"> • Write Through— Data is written through the cache and to the physical drives. Performance is improved, because subsequent reads of that data can be satisfied from the cache. • Write Back— Data is stored in the cache, and is only written to the physical drives when space in the cache is needed. Virtual drives requesting this policy fall back to Write Through caching when the BBU cannot guarantee the safety of the cache in the event of a power failure. • Write Back Bad BBU—With this policy, write caching remains Write Back even if the battery backup unit is defective or discharged.
Disk Cache Policy drop-down list	This can be one of the following <ul style="list-style-type: none"> • Unchanged— The disk cache policy is unchanged. • Enabled— Allows IO caching on the disk. • Disabled— Disallows disk caching.
Access Policy drop-down list	This can be one of the following <ul style="list-style-type: none"> • Read Write— Enables host to perform read-write on the VD. • Read Only— Host can only read from the VD. • Blocked— Host can neither read nor write to the VD.

Name	Description
Size field	The size of the virtual drive you want to create. Enter a value and select one of the following units: <ul style="list-style-type: none"> • MB • GB • TB

Step 7 Click the **Generate XML API Request** button to generate an API request.

Step 8 Click **Close**.

Step 9 Click **Create Virtual Drive**.

Creating Virtual Drive from an Existing Drive Group

Before you begin

You must log in with admin privileges to perform this task.

Step 1 In the **Navigation** pane, click the **Storage** menu.

Step 2 In the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.

Step 3 In the **Actions** area, click **Create Virtual Drive from an Existing Virtual Drive Group**.

The **Create Virtual Drive from an Existing Virtual Drive Group** dialog box displays.

Step 4 In the **Create Virtual Drive from an Existing Virtual Drive Group** dialog box, select the virtual drive whose drive group you want to use to create a new virtual drive.

Step 5 In the **Virtual Drive Properties** area, update the following properties:

Name	Description
Virtual Drive Name field	The name of the new virtual drive you want to create.
Read Policy drop-down list	The read-ahead cache mode.
Cache Policy drop-down list	The cache policy used for buffering reads.
Strip Size drop-down list	The size of each strip, in KB.

Name	Description
Write Policy drop-down list	This can be one of the following <ul style="list-style-type: none"> • Write Through— Data is written through the cache and to the physical drives. Performance is improved, because subsequent reads of that data can be satisfied from the cache. • Write Back— Data is stored in the cache, and is only written to the physical drives when space in the cache is needed. Virtual drives requesting this policy fall back to Write Through caching when the BBU cannot guarantee the safety of the cache in the event of a power failure. • Write Back Bad BBU—With this policy, write caching remains Write Back even if the battery backup unit is defective or discharged.
Disk Cache Policy drop-down list	This can be one of the following <ul style="list-style-type: none"> • Unchanged— The disk cache policy is unchanged. • Enabled— Allows IO caching on the disk. • Disabled— Disallows disk caching.
Access Policy drop-down list	This can be one of the following <ul style="list-style-type: none"> • Read Write— Enables host to perform read-write on the VD. • Read Only— Host can only read from the VD. • Blocked— Host can neither read nor write to the VD.
Size field	The size of the virtual drive you want to create. Enter a value and select one of the following units: <ul style="list-style-type: none"> • MB • GB • TB

Step 6 Click the **Generate XML API Request** button to generate an API request.

Step 7 Click **Close**.

Step 8 Click **Create Virtual Drive**.

Setting a Virtual Drive to Transport Ready State

You can move a virtual drive from one MegaRAID controller to another using the **Set Transport Ready** feature. This allows all the pending IOs of the virtual drive to complete their activities, hide the virtual drive from the operating system, flush cache, pause all the background operations, and save the current progress in disk data format, allowing you to move the drive. When you move a virtual drive, all other drives belonging to the same drive group inherit the same change as the moved drive.

When the last configured physical drive on the group is removed from the current controller, the drive group becomes foreign and all foreign configuration rules apply to the group. However, the Transport Ready feature does not change any foreign configuration behavior.

You can also clear a virtual drive from the Transport Ready state. This makes the virtual drive available to the operating systems.

Following restrictions apply to a transport ready virtual drive:

- Only a maximum of 16 transport ready drive groups are currently supported.
- This feature is not supported on high availability.
- A virtual drive cannot be set as transport ready under these conditions:
 - When a virtual drive of a drive group is being reconstructed
 - When a virtual drive of a drive group contains a pinned cache
 - When a virtual drive of a drive group is marked as cacheable or associated with a cachecade virtual drive
 - If a virtual drive is a cachecade virtual drive
 - If a virtual drive is offline
 - If a virtual drive is a bootable virtual drive

Setting a Virtual Drive as Transport Ready

Before you begin

- You must log in with admin privileges to perform this task.
- The virtual drive must be in optimal state to enable transport ready.

-
- Step 1** In the **Navigation** pane, click the **Storage** menu.
- Step 2** On the **Storage** menu, click the appropriate LSI MegaRAID or HBA Controller.
- Step 3** On the **Work** pane, click the **Virtual Drive Info** tab.
- Step 4** In the **Virtual Drives** area, choose the drive that you want set as transport ready.
- Step 5** In the **Actions** area, click **Set Transport Ready**.
The **Set Transport Ready** dialog box displays.
- Step 6** Update the following properties in the dialog box:

Name	Description
Initialize Type drop-down list	Allows you to select the initialization type using which you can set the selected virtual drive as transport ready. This can be one of the following: <ul style="list-style-type: none"> • Exclude All— Excludes all the dedicated hot spare drives. • Include All— Includes any exclusively available or shared dedicated hot spare drives. • Include Dedicated Hot Spare Drive— Includes exclusive dedicated hot spare drives.
Set Transport Ready button	Sets the selected virtual drive as transport ready.
Cancel button	Cancels the action.

Note When you set a virtual drive to transport ready all the physical drives associated with it are displayed as **Ready to Remove**.

Clearing a Virtual Drive from Transport Ready State

Before you begin

- You must log in with admin privileges to perform this task.
- The virtual drive must be transport ready.

- Step 1** In the **Navigation** pane, click the **Storage** menu.
- Step 2** On the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.
- Step 3** On the **Work** pane, click the **Virtual Drive Info** tab.
- Step 4** In the **Virtual Drives** area, choose the drive to set as transport ready.
- Step 5** In the **Actions** area, click **Clear Transport Ready**.

This reverts the selected transport ready virtual drive to its original optimal state.

Importing Foreign Configuration

When one or more physical drives that have previously been configured with a different controller are inserted into a server, they are identified as foreign configurations. You can import these foreign configurations to a controller.



Important You cannot import a foreign configuration in the following two scenarios:

1. When the secure virtual drive was created on server 1 (from which you want to import the configuration) using the remote key, and on server 2 (to which you want to import) using the local key.
2. When server 2 is configured with another KMIP server, which is not a part of the server 1 KMIP server cluster.

In order to import the foreign configuration in these scenarios, change the controller security on server 2 from local key management to remote key management, and use the same KMIP server from the same cluster where the server 1 KMIP is configured.

Before you begin

You must log in with admin privileges to perform this task.

Step 1 In the **Navigation** pane, click the **Storage** menu.

Step 2 In the **RAID controller** area, the **Controller Info** tab displays by default.

Step 3 In the **Actions** area, click **Import Foreign Config**.

Note If KMIP is not enabled, a **Secure Key Verification** dialog box is displayed, prompting you to enter a security key to initiate the foreign configuration import process.

If KMIP is enabled, the **Secure Key Verification** dialog box is displayed with the following note: *"If drive security has been enabled via remote key management, specifying Security key is optional. Click on verify to start foreign configuration import."*

This allows you to click **Verify** without entering the Security Key, and initiate import.

Step 4 Click **OK** to confirm.

Clearing Foreign Configuration



Important This task clears all foreign configuration on the controller. Also, all configuration information from all physical drives hosting foreign configuration is deleted. This action cannot be reverted.

Before you begin

You must log in with admin privileges to perform this task.

Step 1 In the **Navigation** pane, click the **Storage** menu.

Step 2 On the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.
In the **RAID Controller** area, the **Controller Info** tab displays by default.

Step 3 In the **Actions** area, click **Clear Foreign Config**.

Step 4 Click **OK** to confirm.

Clearing a Boot Drive



Important This task clears the boot drive configuration on the controller. This action cannot be reverted.

Before you begin

You must log in with admin privileges to perform this task.

Step 1 In the **Navigation** pane, click the **Storage** menu.

Step 2 On the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.
In the **RAID Controller** area, the **Controller Info** tab displays by default.

Step 3 In the **Actions** area, click **Clear Boot Drive**.

Step 4 Click **OK** to confirm.

Enabling JBOD Mode

Step 1 In the **Navigation** pane, click the **Storage** menu.

Step 2 On the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.

Step 3 In the **RAID Controller** area, click the **Physical Drive Info** tab.

Step 4 In the **Physical Drives** area, select an unconfigured good drive.

Step 5 In the **Actions** area, click **Enable JBOD**.

Step 6 Click **Ok** to confirm.

Disabling a JBOD



Note This option is available only on some UCS C-Series servers.

Before you begin

JBOD option must be enabled for the selected controller.

-
- Step 1** In the **Navigation** pane, click the **Storage** menu.
 - Step 2** On the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.
 - Step 3** In the **RAID Controller** area, click the **Physical Drive Info** tab.
 - Step 4** In the **Physical Drives** area, select a JBOD drive.
 - Step 5** In the **Actions** area, click **Disable JBOD**.
 - Step 6** Click **Ok** to confirm.
-

Retrieving Storage Firmware Logs for a Controller

This task retrieves the storage firmware logs for the controller and places it in the `/var/log` location. This ensures that this log data is available when Technical Support Data is requested.

Before you begin

You must log in with admin privileges to perform this task.

- Step 1** In the **Navigation** pane, click the **Storage** menu.
- Step 2** In the working area, the **Controller Info** tab displays by default.
- Step 3** In the **Actions** area, click **Get Storage Firmware Log**.
- Step 4** Click **OK** to confirm.

Important Retrieving storage firmware logs for a controller could take up to 2-4 minutes. Until this process is complete, do not initiate exporting technical support data.

Clearing Controller Configuration

Before you begin

You must log in with admin privileges to perform this task.

- Step 1** In the **Navigation** pane, click the **Storage** menu.
- Step 2** In the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.
- Step 3** In the **Controller Info** area, click **Clear All Configuration**.
- Step 4** Click **OK** to confirm.

This clears the existing controller configuration.

Restoring Storage Controller to Factory Defaults

Before you begin

You must log in with admin privileges to perform this task.

- Step 1** In the **Navigation** pane, click the **Storage** menu.
 - Step 2** In the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.
 - Step 3** In the **Controller Info** area, click **Set Factory Defaults**.
 - Step 4** Click **OK** to confirm.
This restores the controller configuration to factory defaults.
-

Preparing a Drive for Removal



Note You can perform this task only on physical drives that display the **Unconfigured Good** status.

Before you begin

You must log in with admin privileges to perform this task.

- Step 1** In the **Navigation** pane, click the **Storage** menu.
 - Step 2** On the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.
 - Step 3** In the **RAID Controller** area, click the **Physical Drive Info** tab.
 - Step 4** In the **Physical Drives** area, select the drive you want to remove.
 - Step 5** In the **Actions** area, click **Prepare for Removal**.
 - Step 6** Click **OK** to confirm.
-

Undo Preparing a Drive for Removal

Before you begin

You must log in with admin privileges to perform this task.

- Step 1** In the **Navigation** pane, click the **Storage** menu.
- Step 2** On the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.
- Step 3** On the **RAID Controller** area, click the **Physical Drive Info** tab.
- Step 4** In the **Physical Drives** area, select a drive with a status of **Ready to Remove**.

Step 5 In the **Actions** area, click **Undo Prepare for Removal**.

Step 6 Click **OK** to confirm.

Making a Dedicated Hot Spare

Before you begin

You must log in with admin privileges to perform this task.

Step 1 In the **Navigation** pane, click the **Storage** tab.

Step 2 On the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.

Step 3 On the **RAID Controller** area, click the **Physical Drive Info** tab.

Step 4 In the **Physical Drives** area, select an unconfigured good drive you want to make a dedicated hot spare.

Step 5 In the **Actions** area, click **Make Dedicated Hot Spare**.

The **Make Dedicated Hot Spare** dialog box displays.

Step 6 In the **Virtual Drive Details** area, update the following properties:

Name	Description
Virtual Drive Number drop-down list	Select the virtual drive to which you want to dedicate the physical drive as hot spare.
Virtual Drive Name field	The name of the selected virtual drive.
Make Dedicated Hot Spare button	Creates the dedicated hot spare.
Cancel button	Closes the dialog box without saving any changes made while the dialog box was open.

Step 7 Click **Make Dedicated Hot Spare** to confirm.

Making a Global Hot Spare

Before you begin

You must log in with admin privileges to perform this task.

Step 1 In the **Navigation** pane, click the **Storage** tab.

Step 2 On the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.

Step 3 In the **RAID Controller** area, click the **Physical Drive Info** tab.

Step 4 In the **Physical Drives** area, select an unconfigured good drive you want to make a global hot spare.

Step 5 In the **Actions** area, click **Make Global Hot Spare**.

Removing a Drive from Hot Spare Pools

Before you begin

You must log in with admin privileges to perform this task.

Step 1 In the **Navigation** pane, click the **Storage** menu.

Step 2 On the **Storage** tab, click the appropriate LSI MegaRAID controller.

Step 3 On the **Work** pane, click the **Physical Drive Info** tab.

Step 4 In the **Physical Drives** area, select the global or dedicated hot spare you want to remove from the hot spare pools.

Step 5 In the **Actions** area, click **Remove From Hot Spare Pools**.

Toggling Physical Drive Status

Before you begin

- You must log in with admin privileges to perform this task.
 - The controller must support the JBOD mode and the JBOD mode must be enabled.
-

Step 1 In the **Navigation** pane, click the **Storage** tab.

Step 2 On the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.

Step 3 In the **RAID Controller** area, click the **Physical Drive Info** tab.

Step 4 In the **Physical Drives** area, select the drive you want to set as unconfigured good.

Step 5 In the **Actions** area, click **Set State as Unconfigured Good**.

Step 6 Click **OK** to confirm that the JBOD mode be disabled.

The **Set State as JBOD** option is enabled.

Step 7 To enable the JBOD mode for the physical drive, click **Set State as JBOD**.

Step 8 Click **OK** to confirm.

The **Set State as Unconfigured Good** option is enabled.

Setting a Physical Drive as a Controller Boot Drive

Before you begin

- You must log in with admin privileges to perform this task.

- The controller must support the JBOD mode and the JBOD mode must be enabled.

-
- Step 1** In the **Navigation** pane, click the **Storage** menu.
- Step 2** On the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.
- Step 3** In the **RAID Controller** area, click the **Physical Drive Info** tab.
- Step 4** In the **Physical Drives** area, select the drive you want to set as boot drive for the controller.
- Step 5** In the **Actions** area, click **Set as Boot Drive**.
- Step 6** Click **OK** to confirm.
-

Initializing a Virtual Drive

All data on a virtual drive is lost when you initialize the drive. Before you run an initialization, back up any data on the virtual drive that you want to save.

Before you begin

You must log in with admin privileges to perform this task.

-
- Step 1** In the **Navigation** pane, click the **Storage** menu.
- Step 2** On the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.
- Step 3** In the **RAID Controller** area, click the **Virtual Drive Info** tab.
- Step 4** In the **Virtual Drives** area, choose the drive that you want to initialize.
- Step 5** In the **Actions** area, click **Initialize**.
- The **Initialize Virtual Drive** dialog box displays.
- Step 6** Choose the type of initialization you want to use for the virtual drive.
- This can be one of the following:
- **Fast Initialize**—This option allows you to start writing data to the virtual drive immediately.
 - **Full Initialize**—A complete initialization is done on the new configuration. You cannot write data to the new virtual drive until the initialization is complete.
- Step 7** Click **Initialize VD** to initialize the drive, or **Cancel** to close the dialog box without making any changes.
- Step 8** To view the status of the task running on the drive, in the **Operations** area, click **Refresh**.

The following details are displayed:

Name	Description
Operation	Name of the operation that is in progress on the drive.
Progress in %	Progress of the operation, in percentage complete.

Name	Description
Elapsed Time in secs	The number of seconds that have elapsed since the operation began.

Set as Boot Drive

Before you begin

You must log in with admin privileges to perform this task.

- Step 1** In the **Navigation** pane, click the **Storage** menu.
- Step 2** On the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.
- Step 3** In the **RAID Controller** area, click the **Virtual Drive Info** tab.
- Step 4** In the **Virtual Drives** area, choose the drive from which the controller must boot.
- Step 5** In the **Actions** area, click **Set as Boot Drive**.
- Step 6** Click **OK** to confirm.

Editing a Virtual Drive

- Step 1** In the **Navigation** pane, click the **Storage** menu.
- Step 2** On the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.
- Step 3** In the **RAID Controller** area, click the **Virtual Drive Info** tab.
- Step 4** In the **Virtual Drives** area, click **Edit Virtual Drive**.
- Step 5** Review the instructions, and then click **OK**.
The **Edit Virtual Drive** dialog box displays before prompting you to take a backup of your data.
- Step 6** From the **Select RAID Level to migrate** drop-down list, choose a RAID level.
See the following table for RAID migration criteria:

Name	Description
<p>Select RAID Level to migrate drop-down list</p>	<p>Select the RAID level to which you want to migrate. Migrations are allowed for the following RAID levels:</p> <ul style="list-style-type: none"> • RAID 0 to RAID 1 • RAID 0 to RAID 5 • RAID 0 to RAID 6 • RAID 1 to RAID 0 • RAID 1 to RAID 5 • RAID 1 to RAID 6 • RAID 5 to RAID 0 • RAID 6 to RAID 0 • RAID 6 to RAID 5 <p>When you are migrating from one raid level to another, the data arms of the new RAID level should be equal to or greater than the existing one.</p> <p>In case of RAID 6, the data arms will be number of drives minus two, as RAID 6 has double distributed parity. For example, when you create RAID 6 with eight drives, the number of data arms will be $8 - 2 = 6$. In this case, if you are migrating from RAID 6 to RAID 0, RAID 0 must have a minimum of six drives. If you select lesser number of drives then Edit or Save button will be disabled.</p> <p>If you are adding, you can migrate to RAID 0 as you will not be deleting any drives.</p> <p>Note RAID level migration is not supported in the following cases:</p> <ul style="list-style-type: none"> • When there are multiple virtual drives in a RAID group. • With a combination of SSD/HDD RAID groups.

Step 7 From the **Write Policy** drop-down list in the **Virtual Drive Properties** area, choose one of the following:

- **Write Through**— Data is written through the cache and to the physical drives. Performance is improved, because subsequent reads of that data can be satisfied from the cache.
- **Write Back**— Data is stored in the cache, and is only written to the physical drives when space in the cache is needed. Virtual drives requesting this policy fall back to **Write Through** caching when the BBU cannot guarantee the safety of the cache in the event of a power failure.
- **Write Back Bad BBU**—With this policy, write caching remains **Write Back** even if the battery backup unit is defective or discharged.

Step 8 Click **Save Changes**.

Deleting a Virtual Drive



Important This task deletes a virtual drive, including the drives that run the booted operating system. So back up any data that you want to retain before you delete a virtual drive.

Before you begin

You must log in with admin privileges to perform this task.

-
- Step 1** In the **Navigation** pane, click the **Storage** menu.
 - Step 2** On the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.
 - Step 3** In the **RAID Controller** area, click the **Virtual Drive Info** tab.
 - Step 4** In the **Virtual Drives** area, select the virtual drive you want to delete.
 - Step 5** In the **Actions** area, click **Delete Virtual Drive**.
 - Step 6** Click **OK** to confirm.
-

Hiding a Virtual Drive

Before you begin

You must log in with admin privileges to perform this task.

-
- Step 1** In the **Navigation** pane, click the **Storage** menu.
 - Step 2** On the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.
 - Step 3** On the **RAID Controller** area, click the **Virtual Drive Info** tab.
 - Step 4** In the **Virtual Drives** area, select the virtual drive you want to hide.
 - Step 5** In the **Actions** area, click **Hide Drive**.
 - Step 6** Click **OK** to confirm.
-

Starting Learn Cycles for a Battery Backup Unit

Before you begin

You must log in with admin privileges to perform this task.

-
- Step 1** In the **Navigation** pane, click the **Storage** menu.
 - Step 2** On the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.
 - Step 3** In the **RAID Controller** area, click the **Battery Backup Unit** tab.

Step 4 From the **Actions** pane, click **Start Learn Cycle**.

A dialog prompts you to confirm the task.

Step 5 Click **OK**.

Viewing Storage Controller Logs

Before you begin

You must log in with admin privileges to perform this task.

Step 1 In the **Navigation** pane, click the **Storage** menu.

Step 2 On the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.

Step 3 In the **RAID Controller** area, click **Storage Log** tab and review the following information:

Name	Description
Time column	The date and time the event occurred.
Severity column	The event severity. This can be one of the following: <ul style="list-style-type: none"> • Emergency • Alert • Critical • Error • Warning • Notice • Informational • Debug
Description column	A description of the event.

Viewing SSD Smart Information for MegaRAID Controllers

You can view smart information for a solid state drive. Complete these steps:

Step 1 In the **Navigation** pane, click the **Storage** tab.

Step 2 On the **Storage** menu, click the appropriate LSI MegaRAID Controller.

Step 3 On the **Work** pane, click the **Physical Drive Info** tab.

Step 4 In the **Smart Information** area, review the following information:

Name	Description
Power Cycle Count field	Number of power cycles that the drive went through from the time it was manufactured.
Power on Hours field	Total number of hours that the drive is in the 'Power On' mode.
Percentage Life Left field	The number of write cycles remaining in a solid state drive (SSD). For instance, if an SSD is capable of 100 write cycles during its life time, and it has completed 15 writes, then the percentage of life left in the drive is 85%. Each percentage range is represented in a different color. For instance, green for 75% to 100% and red for 1 to 25%.
Wear Status in Days field	The number of days an SSD has gone through with the write cycles. SSD vendors provide a finite number of writes per day on the SSD, based on which, you can calculate the total number of years the SSD would continue to work.
Operating Temperature field	The current temperature of the drive at which the selected SSD operates at the time of selection.
Percentage Reserved Capacity Consumed field	The total capacity (out of the percentage reserved for it) consumed by the SSD.
Time of Last Refresh field	Time period since the drive was last refreshed.

Viewing NVMe Controller Details

Before you begin

- The server must be powered on.

- Step 1** In the **Navigation** pane, click the **Storage** menu.
- Step 2** In the **Storage** menu, click the appropriate NVMe controller.
- Step 3** In the **Controller** area, the **Controller Info** tab displays by default.
- Step 4** In the **Work** pane's **Health/Status** area, review the following information:

Name	Description
Composite Health field	The health of the controller. This can be one of the following: <ul style="list-style-type: none"> • Good- All the drives under the controller is in optimal state. • Severe Fault- When one or more drives under the controller is in faulty state. • N/A
Drive Count field	The number of drives configured on the controller.

Step 5 In the **Manufacturer Information** area, review the following information:

Name	Description
Vendor ID field	The vendor ID of the NVMe controller.
Product ID field	The controller product ID.
Component ID field	The component ID of the NVMe controller.
Product Revision field	The board revision number, if any.

Step 6 In the **Group PCI Info** area, review the following information:

Name	Description
Vendor ID field	The PCI vendor ID, in hexadecimal.
Device ID field	The PCI device ID, in hexadecimal.

Step 7 In the **Group Firmware Information** area, review the following information:

Name	Description
Running Firmware Images field	NVMe drive firmware version.

Step 8 In the **Group Switch Information** area, review the following information:

Name	Description
Temperature field	Temperature in degree centigrade of the switch.
Switch Status field	The current status of the switch. This can be one of the following: <ul style="list-style-type: none"> • Optimal — The controller is functioning properly. • Failed — The controller is not functioning. • Unresponsive — The controller is down.

Name	Description
Link Status field	The current status of the link. This field indicates if any of the upstream or downstream links in the switch are down. Individual drive also has a link status that can then be used to identify which drive causes the switch link status to be Link Degraded. This can be one of the following: <ul style="list-style-type: none"> • Optimal — The controller is functioning properly. • Failed — The controller is not functioning. • Unresponsive — The controller is down.
Shutdown Temperature field	This is the temperature beyond which the safe operation of switch is not guaranteed and recommends shutting down the system.

Viewing NVMe Physical Drive Details

Before you begin

- The server must be powered on.

- Step 1** In the **Navigation** pane, click the **Storage** menu.
- Step 2** In the **Storage** menu, click the appropriate NVMe controller.
- Step 3** Click **Physical Drive** tab and review the following information:

Name	Description
Physical Drives column	The list of available physical drives.
PCI Slot column	The PCI slot number in which the physical drive resides.
Managed ID column	Internal ID referenced in debug.
Product Name column	The name of the drive as assigned by the vendor.
Firmware Version column	The firmware version running on the drive.
Vendor column	The drive vendor name.
Serial Number column	Drive serial number.

Step 4 **Physical Drives** Details

Note These details of a physical drive are displayed when you expand one of the listed physical drives.

Name	Description
PCI Slot field	The PCI slot number in which the physical drive resides.
Managed ID field	Internal ID referenced in debug.
Throttle State field	The state of the throttle.
Serial Number field	Controller serial number.
Chip Temperature field	Temperature of the drive in degree centigrade. This is the max temperature read from internal sensors in the drive.
Percentage Drive Life Used field	The percentage of the drive life that is used up.
Device ID field	The PCI device ID, in hexadecimal.
Sub Device ID field	The PCI subdevice ID, in hexadecimal.
Drive Status field	Status of the drive.
Performance Level field	Indicates the performance of the drive.
Shutdown Temperature field	Temperature at which the drive shuts down.
Percentage of Total Power On Hours field	The percentage of time the drive was powered on.
Vendor ID field	The PCI vendor ID, in hexadecimal.
SubVendor ID field	The PCI subvendor ID, in hexadecimal.
LED Fault Status field	The status of the LED fault.
Controller Temperature field	Temperature of the controller in degree centigrade. This is the overall composite temperature of the NVMe subsystem ID.
Running Firmware Images field	NVMe drive firmware version.
Throttle Start Temperature field	Temperature at which the drive starts to throttle.

Starting Copyback Operation

- Step 1** In the **Navigation** pane, click the **Storage** menu.
- Step 2** On the **Storage** menu, click the appropriate LSI MegaRAID or HBA controller.
- Step 3** In the **RAID Controller** area, click the **Physical Drive Info** tab.
- Step 4** In the **Physical Drives** area, select a drive, which is in online state.
- Step 5** In the **Actions** area, click **Start Copyback**.
- Step 6** A **Start Copyback Operation** dialog box appears.

Step 7 Select the **Destination Physical Drive** to which the copyback operation needs to be done.

Step 8 Click **Start Copyback**.

Step 9 You can also perform the following copyback operations:

- **Pause Copyback** - If the drive is in copyback state, you can pause the copyback operation.
- **Resume Copyback**- The paused copyback operation can be resumed.
- **Abort Copyback** - If the drive is in copyback state, you can abort the copyback operation.

Managing the Flexible Flash Controller

Cisco Flexible Flash

On the M5 servers, Flexible Flash Controller is inserted into the mini storage module socket. The mini storage socket is inserted into the M.2 slot on the motherboard. M.2 slot also supports SATA M.2 SSD slots.



Note M.2 slot does not support NVMe in this release.

Some C-Series Rack-Mount Servers support an internal Secure Digital (SD) memory card for storage of server software tools and utilities. The SD card is hosted by the Cisco Flexible Flash storage adapter.

The SD storage is available to Cisco IMC as a single hypervisor (HV) partition configuration. Prior versions had four virtual USB drives. Three were preloaded with Cisco UCS Server Configuration Utility, Cisco drivers and Cisco Host Upgrade Utility, and the fourth as user-installed hypervisor. A single HV partition configuration is also created when you upgrade to the latest version of Cisco IMC or downgrade to the prior version, and reset the configuration.

For more information about installing and configuring the M.2 drives, see the **Storage Controller Considerations (Embedded SATA RAID Requirements)** and **Replacing an M.2 SSD in a Mini-Storage Carrier For M.2** sections in the Cisco UCS Server Installation and Service Guide for the C240 M5 servers at this URL:

<https://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-c-series-rack-servers/products-installation-guides-list.html>

For information about the Cisco software utilities and packages, see the *Cisco UCS C-Series Servers Documentation Roadmap* at this URL:

<http://www.cisco.com/go/unifiedcomputing/c-series-doc>

Card Management Feature in the Cisco Flexible Flash Controller

The Cisco Flexible Flash controller supports management of both single and two SD cards as a RAID-1 pair. With the introduction of card management, you can perform the following tasks:



- Note**
- If you want to upgrade from version 1.4(5e) to 1.5(4) or higher versions, you must first upgrade to version 1.5(2) and then upgrade to a higher version of Cisco IMC.
 - Reset the Cisco Flexible Flash controller to load the latest Flex Flash firmware after every Cisco IMC firmware upgrade.

Action	Description
Reset Cisco Flex Flash	Allows you to reset the controller.
Reset Partition Defaults	Allows you to reset the configuration in the selected slot to the default configuration.
Synchronize Card Configuration	Allows you to retain the configuration for an SD card that supports firmware version 253 and later.
Configure Operational Profile	Allows you to configure the SD cards on the selected Cisco Flexible Flash controller.

RAID Partition Enumeration

Non-RAID partitions are always enumerated from the primary card and the enumeration does not depend on the status of the primary card.

Following is the behavior of the RAID partition enumeration when there are two cards in the Cisco Flexible Flash controller:

Scenario	Behavior
Single card	RAID partitions are enumerated if the card is healthy, and if the mode is either Primary or Secondary-active .
Dual paired cards	RAID partitions are enumerated if one of the cards is healthy. When only one card is healthy, all read/write operations occur on this healthy card. You must use UCS SCU to synchronize the two RAID partitions.
Dual unpaired cards	If this scenario is detected when the server is restarting, then neither one of the RAID partitions is enumerated. If this scenario is detected when the server is running, when a user connects a new SD card, then the cards are not managed by the Cisco Flexible Flash controller. This does not affect the host enumeration. You must pair the cards to manage them. You can pair the cards using the Reset Partition Defaults or Synchronize Card Configuration options.

Upgrading from Single Card to Dual Card Mirroring with FlexFlash

You can upgrade from a single card mirroring to dual card mirroring with FlexFlash in one of the following methods:

- Add an empty FlexFlash card to the server, and then upgrade its firmware to the latest version.
- Upgrade the FlexFlash firmware to the latest version and then add an empty card to the server.

Prior to using either of these methods, you must keep in mind the following guidelines:

- To create RAID1 mirroring, the empty card that you want to add to the server must be of the exact size of the card that is already in the server. Identical card size is a must to set up RAID1 mirroring.
- Ensure that the card with valid data in the Hypervisor partition is marked as the primary healthy card. You can determine this state either in the Cisco IMC GUI or from the Cisco IMC CLI. To mark the state of the card as primary healthy, you can either use the **Reset Configuration** option in the Cisco IMC GUI or run the **reset-config** command in the Cisco IMC CLI. When you reset the configuration of a particular card, the secondary card is marked as secondary active unhealthy.
- In a Degraded RAID health state all read-write transactions are done on the healthy card. In this scenario, data mirroring does not occur. Data mirroring occurs only in the Healthy RAID state.
- Data mirroring is only applicable to RAID partitions. In the C-series servers, only Hypervisor partitions operate in the RAID mode.
- If you have not configured SD cards for use with prior versions, then upgrading to the latest version loads the latest 253 firmware and enumerates all four partitions to the host.

While upgrading versions of the FlexFlash, you may see the following error message:

```
Unable to communicate with Flexible Flash controller: operation ffCardsGet, status
CY_AS_ERROR_INVALID_RESPONSE"
```

In addition, the card status may be shown as **missing**. This error occurs because you accidentally switched to an alternate release or a prior version, such as 1.4(x). In this scenario, you can either revert to the latest version, or you can switch back to the FlexFlash 1.4(x) configuration. If you choose to revert to the latest Cisco IMC version, then the Cisco FlexFlash configuration remains intact. If you choose to switch back to the prior version configuration, you must reset the Flexflash configuration. In this scenario, you must be aware of the following:

- If multiple cards are present, and you revert to a prior version, then the second card cannot be discovered or managed.
- If the card type is SD253, then you must run the **reset-config** command twice from the Cisco IMC CLI - once to reload the old firmware on the controller and to migrate SD253 to SD247 type, and the second time to start the enumeration.

Configuring the Flexible Flash Controller Properties

After you upgrade to the latest version of Cisco IMC or downgrade to a prior version, and reset the configuration, the server will access HV partition only.

Before you begin

- You must log in with admin privileges to perform this task.
- Cisco Flexible Flash must be supported by your platform.



Note This task results in the host re-scanning all the virtual drives, and a loss of virtual drive connectivity. We recommend that you configure the Cisco Flexible Flash controller properties before using any virtual drives, or power down the host prior to starting this task

Configuring the Flexible Flash Controller Firmware Mode

Before you begin

- You must log in with admin privileges to perform this task.
- Cisco Flexible Flash must be supported by your platform.



Note This task results in the host re-scanning all the virtual drives, and a loss of virtual drive connectivity. We recommend that you configure the Cisco Flexible Flash controller properties before using any virtual drives, or power down the host prior to starting this task.

Step 1 In the **Navigation** pane, click the **Storage** tab.

Step 2 On the **Storage** tab, click **Cisco FlexFlash**.

Step 3 In the **Actions** area, click **Configure Firmware Mode**.

Step 4 Click **OK** in the confirmation box.

Switches the controller firmware mode from the current firmware mode to the other.

Configuring the Flexible Flash Controller Cards

Before you begin

- You must log in with admin privileges to perform this task.
- Cisco Flexible Flash must be supported by your platform.



Note This task results in the host re-scanning all the virtual drives, and a loss of virtual drive connectivity. We recommend that you configure the Cisco Flexible Flash controller properties before using any virtual drives, or power down the host prior to starting this task.

- Step 1** In the **Navigation** pane, click the **Storage** tab.
- Step 2** On the **Storage** tab, click **Cisco FlexFlash**.
- Step 3** In the **Actions** area, click **Configure Cards**.
Configure Cards dialog box appears.
- Step 4** In the **Configure Cards** dialog box, update the following fields:

Name	Description
Mirror radio button	<p>Enter the following:</p> <ul style="list-style-type: none"> • Mirror Partition Name field—The name that you want to assign to the partition. • Auto Sync checkbox—If selected, data from the selected primary card will sync automatically with the secondary card. <ul style="list-style-type: none"> Note <ul style="list-style-type: none"> • There must be two cards for you to choose this option. • If this option is selected, data on the secondary card is erased and overwritten by the data on the primary card. • The status of this is displayed under the Physical Driver Info tab. • Select Primary Card drop-down—Slot that you want to set as the primary card. This can be one of the following: <ul style="list-style-type: none"> • Slot1 • Slot2

Name	Description
Util radio button	<p>Select this option to configure the card in Util mode. When you configure the cards in the Util mode, the following situations occur:</p> <ul style="list-style-type: none"> • The card in the selected slot creates four partitions that has a partition each for the utilities: SCU, HUU, Drivers and one partition that can be used by the user and the card is marked healthy. • The card in the other slot, if it exists, creates a single partition and the card is marked healthy. • The card read/write error counts and read/write threshold are set to 0. • Host connectivity could be disrupted. • The configured cards will be paired. <p>Enter the following:</p> <ul style="list-style-type: none"> • User Partition Name field—The name that you want to assign to the fourth partition of the Util card. • Non Util Card Partition Name field—The name that you want to assign to the single partition on the second card, if it exists. • Select Util Card drop-down—Slot that you want to set for Util. This can be one of the following: <ul style="list-style-type: none"> • Slot1 • Slot2 • None—Applicable only when the server has one SD card.

Step 5 Click **Save**.

The cards are configured in the chosen mode.

Booting from the Flexible Flash Card

You can specify a bootable virtual drive on the Cisco Flexible Flash card that overrides the default boot priority the next time that the server is restarted, regardless of the default boot order defined for the server. The specified boot device is used only once. After the server has rebooted, this setting is ignored. You can choose a bootable virtual drive only if a Cisco Flexible Flash card is available. Otherwise, the server uses a default boot order.



Note Before you reboot the server, ensure that the virtual drive that you select is enabled on the Cisco Flexible Flash card. Go to the **Storage** tab, choose the card, and then go to the **Virtual Drive Info** subtab.

Before you begin

- You must log in with admin privileges to perform this task.
- Cisco Flexible Flash must be supported by your platform.

Step 1 In the **Navigation** pane, click the **Server** tab.

Step 2 On the **Server** tab, click **BIOS**.

Step 3 In the **Actions** area, click **Configure Boot Override Priority**.

The **Boot Override Priority** dialog box appears.

Step 4 From the **Boot Override Priority** drop-down list, choose a virtual drive to boot from.

Step 5 Click **Apply**.

Resetting the Flexible Flash Controller

In normal operation, it should not be necessary to reset the Cisco Flexible Flash. We recommend that you perform this procedure only when explicitly directed to do so by a technical support representative.



Note This operation will disrupt traffic to the virtual drives on the Cisco Flexible Flash controller.

Before you begin

- You must log in with admin privileges to perform this task.
- Cisco Flexible Flash must be supported by your platform.

Step 1 On the **Storage Adapters** pane, click **Cisco FlexFlash**.

Step 2 In the **Cisco FlexFlash** pane, click the **Controller Info** tab.

Step 3 In the **Actions** area, click **Reset FlexFlash Controller**.

Step 4 Click **OK** to confirm.

Enabling Virtual Drives

Before you begin

- You must log in with admin privileges to perform this task.
- Cisco Flexible Flash must be supported by your platform.



Note This task results in the host re-scanning all the virtual drives, and a loss of virtual drive connectivity. We recommend that you configure the Cisco Flexible Flash controller properties before using any virtual drives.

-
- Step 1** In the **Navigation** pane, click the **Storage** tab.
- Step 2** On the **Storage** tab, click **Cisco FlexFlash**.
- Step 3** Click the **Virtual Drive Info** tab.
- Step 4** In the **Virtual Drive Info** tab, click **Enable/Disable Virtual Drive(s)**.
- Step 5** In the **Enable/Disable VD(s)** dialog box, select the virtual drives that you want to enable.
- Step 6** Click **Save**.
The selected virtual drives are enabled to the host.
-

Erasing Virtual Drives

Before you begin

- You must log in with admin privileges to perform this task.
- Cisco Flexible Flash must be supported by your platform.



Note This task results in the host re-scanning all the virtual drives, and a loss of virtual drive connectivity. We recommend that you configure the Cisco Flexible Flash controller properties before using any virtual drives.

-
- Step 1** In the **Navigation** pane, click the **Storage** tab.
- Step 2** On the **Storage** tab, click **Cisco FlexFlash**.
- Step 3** Click the **Virtual Drive Info** tab.
- Step 4** In the **Virtual Drive Info** tab, click **Erase Virtual Drive(s)**.
- Step 5** In the **Erase Virtual Drive(s)** dialog box, select the virtual drives that you want to erase.
- Step 6** Click **Save**.
Data on the selected virtual drives is erased.
-

Syncing Virtual Drives

Before you begin

- You must log in with admin privileges to perform this task.
- Cisco Flexible Flash must be supported by your platform.

- Cards must be in mirror mode.



Note This task results in the host re-scanning all the virtual drives, and a loss of virtual drive connectivity. We recommend that you configure the Cisco Flexible Flash controller properties before using any virtual drives.

-
- Step 1** In the **Navigation** pane, click the **Storage** tab.
- Step 2** On the **Storage** tab, click **Cisco FlexFlash**.
- Step 3** Click the **Virtual Drive Info** tab.
- Step 4** In the **Virtual Drive Info** tab, click **Sync Virtual Drive**.
- Step 5** Click **OK** in the confirmation dialog box.
Syncs the virtual drive hypervisor with the primary card.
-

Adding an ISO Image Configuration

Before you begin

- You must log in with admin privileges to perform this task.
- Cisco Flexible Flash must be supported by your platform.
- The cards must be configured in Util mode.



Note This task results in the host re-scanning all the virtual drives, and a loss of virtual drive connectivity. We recommend that you configure the Cisco Flexible Flash controller properties before using any virtual drives, or power down the host prior to starting this task.

-
- Step 1** In the **Navigation** pane, click the **Storage** tab.
- Step 2** On the **Storage** tab, click **Cisco FlexFlash**.
- Step 3** Click the **Virtual Drive Info** tab.
- Step 4** In the **Virtual Drive Info** tab, select the virtual drive for which you want to add an image, click **Add Image**.
- Step 5** In the **Add Image** dialog box, update the following fields:

Name	Description
Volume field	The identity of the image mounted for mapping. This can be one of the following: <ul style="list-style-type: none"> • SCU • HUU • Drivers

Name	Description
Mount Type drop-down list	<p>The type of mapping. This can be one of the following:</p> <ul style="list-style-type: none"> • NFS—Network File System. • CIFS—Common Internet File System.
Remote Share field	<p>The URL of the image to be mapped. The format depends on the selected Mount Type:</p> <ul style="list-style-type: none"> • NFS—Use <code>serverip:/share path</code>. • CIFS—Use <code>//serverip/share path</code>.
Remote File field	<p>The name and location of the .iso file in the remote share. Following are the example of remote share files:</p> <ul style="list-style-type: none"> • NFS — <code>/softwares/ucs-cxx-scu-3.1.9.iso</code> • CIFS — <code>/softwares/ucs-cxx-scu-3.1.9.iso</code>
Mount Options field	<p>Industry-standard mount options entered in a comma separated list. The options vary depending on the selected Mount Type.</p> <p>If you are using NFS, leave the field blank or enter one or more of the following:</p> <ul style="list-style-type: none"> • <code>ro</code> • <code>rw</code> • <code>noexec</code> • <code>noexec</code> • <code>soft</code> • <code>port=VALUE</code> • <code>timeo=VALUE</code> • <code>retry=VALUE</code> <p>If you are using CIFS, leave the field blank or enter one or more of the following:</p> <ul style="list-style-type: none"> • <code>soft</code> • <code>nounix</code> • <code>noserverino</code>
User Name field	The username for the specified Mount Type , if required.
Password field	The password for the selected username, if required.

Step 6 Click **Save**.

Updating an ISO Image

Before you begin

- You must log in with admin privileges to perform this task.
- Cisco Flexible Flash must be supported by your platform.
- This task is available only when the cards are configured in **Util** mode.



Note This task results in the host re-scanning all the virtual drives, and a loss of virtual drive connectivity. We recommend that you configure the Cisco Flexible Flash controller properties before using any virtual drives, or power down the host prior to starting this task.

Step 1 In the **Navigation** pane, click the **Storage** tab.

Step 2 On the **Storage** tab, click **Cisco FlexFlash**.

Step 3 Click the **Virtual Drive Info** tab.

Step 4 In the **Virtual Drive Info** tab, select the virtual drive on which you want to update the image, click **Update Image**.

Note SCU and HUU update may take up to an hour and the drivers update may take up to five hours.

Unmapping an ISO Image

Before you begin

- You must log in with admin privileges to perform this task.
- Cisco Flexible Flash must be supported by your platform.



Note This task results in the host re-scanning all the virtual drives, and a loss of virtual drive connectivity. We recommend that you configure the Cisco Flexible Flash controller properties before using any virtual drives, or power down the host prior to starting this task.

Step 1 In the **Navigation** pane, click the **Storage** tab.

Step 2 On the **Storage** tab, click **Cisco FlexFlash**.

Step 3 Click the **Virtual Drive Info** tab.

Step 4 In the **Virtual Drive Info** tab, select the virtual drive for which you want to un map the image, click **Unmap Image**.

Resetting the Cisco Flexible Flash Card Configuration

When you reset the configuration of the slots in the Cisco Flexible Flash card, the following situations occur:

- The card in the selected slot is marked as primary healthy.
- The card in the other slot is marked as secondary-active unhealthy.
- One RAID partition is created.
- The card read/write error counts and read/write threshold are set to 0.
- Host connectivity could be disrupted.

If you upgrade to the latest version and select reset configuration option, a single hypervisor (HV) partition is created, and the existing four partition configurations are erased. This may also result in data loss. You can retrieve the lost data only if you have not done any data writes into HV partition, and downgrade to prior version.

Before you begin

You must log in with admin privileges to perform this task.

Step 1 On the **Storage Adapters** pane, click **Cisco FlexFlash**.

Step 2 In the **Cisco FlexFlash** pane, click the **Controller Info** tab.

Step 3 In the **Actions** area, click **Reset Partition Defaults**.

Step 4 In the **Reset Partition Defaults** dialog box, update the following fields:

Name	Description
Slot radio button	Select the slot for which you want to mark the card as primary healthy. The card in the other slot, if any, is marked as secondary-active unhealthy.
Reset Partition Defaults button	Resets the configuration of the selected slot.
Cancel button	Closes the dialog box without making any changes.

Step 5 Click **Yes**.

Retaining Configuration of the Cisco Flexible Flash Cards

You can retain the configuration for an FlexFlash that supports firmware version 253 and later card in the following situations:

- There are two unpaired FlexFlash
- The server is operating from a single FlexFlash, and an unpaired FlexFlash is in the other slot.

- One FlexFlash supports firmware version 253, and the other FlexFlash is unpartitioned.

When you retain the configuration, the following situations occur:

- The configuration for the FlexFlash in the selected slot is copied to the other card.
- The card in the selected slot is marked as primary healthy.
- The card in the secondary slot is marked as secondary-active unhealthy.

Before you begin

- You must log in with admin privileges to perform this task.

Step 1 On the **Storage Adapters** pane, click **Cisco FlexFlash**.

Step 2 In the **Cisco FlexFlash** pane, click the **Controller Info** tab.

Step 3 In the **Actions** area, click **Synchronize Card Configuration**.

Step 4 In the **Synchronize Card Configuration** dialog box, update the following fields:

Name	Description
Slot radio button	Select the slot for which you want the configuration retained. The configuration is copied from the selected slot to the card in the other slot, and the card in the selected slot is marked as primary healthy.
Synchronize Card Configuration button	Copies the configuration from the selected card only if the selected card is of type SD253 and has single HV configuration.
Cancel button	Closes the dialog box without making any changes.

Step 5 Click **Yes**.

Cisco Boot Optimized M.2 Raid Controller

Viewing Cisco Boot Optimized M.2 Raid Controller Details

Before you begin

The server must be powered on.

Step 1 In the **Navigation** pane, click the **Storage** menu.

Step 2 In the **Storage** menu, click the appropriate M.2 Raid controller.

Step 3 In the **Controller** area, the **Controller Info** tab displays by default.

Step 4 In the Work pane's **Health/Status** area, review the following information:

Name	Description
Composite Health field	The combined health of the controller, the attached drives, and the battery backup unit. This can be one of the following: <ul style="list-style-type: none"> • Good • Moderate Fault • Severe Fault • N/A
Controller Status field	The current status of the controller. This can be one of the following: <ul style="list-style-type: none"> • Optimal — The controller is functioning properly. • Failed — The controller is not functioning.

Step 5 In the **Firmware Versions** area, review the following information:

Name	Description
Product Name field	The name of the Cisco Boot optimized M.2 Raid controller.
Product PID field	The product PID of the Cisco Boot optimized M.2 Raid controller.
Serial Number field	The serial number of the Cisco Boot optimized M.2 Raid controller.
Firmware Package Build field	The active firmware package version number. For the firmware component version numbers, see the Running Firmware Images area.

Step 6 In the **PCI Info** area, review the following information:

Name	Description
PCI Slot field	The name of the PCIe slot in which the controller is located.
Vendor ID field	The PCI vendor ID, in hexadecimal.
Device ID field	The PCI device ID, in hexadecimal.
SubVendor ID field	The PCI subvendor ID, in hexadecimal.
SubDevice ID field	The PCI subdevice ID, in hexadecimal.

Step 7 In the **Manufacturing Data** area, review the following information:

Name	Description
Manufactured Date field	The manufacturing date of the Cisco Boot optimized M.2 Raid controller, in the format yyyy-mm-dd.
Revision No field	The board revision number, if any.

Step 8 In the **Next Patrol Read Schedule** area, review the following information:

Name	Description
PR State field	The patrol read state of the M.2 Raid controllers. It can be one of the following: <ul style="list-style-type: none"> • ready • stopped • active Default state: N/A .
PR Schedule Mode field	The patrol read schedule mode of the M.2 Raid controllers. The PR Schedule Mode is manual by default.

Step 9 In the **Running Firmware Images** area, review the following information:

Name	Description
BIOS Version field	The BIOS option PROM version number.
Firmware Version field	The active firmware version number.
Boot Block Version field	The boot block version number.

Step 10 In the **Virtual Drive Count** area, review the following information:

Name	Description
Virtual Drive Count field	The number of virtual drives configured on the controller.
Degraded Drive Count field	The number of virtual drives in a degraded state on the controller.
Offline Drive Count field	The number of virtual drives that have failed on the controller.

Step 11 In the **Physical Drive Count** area, review the following information:

Name	Description
Disk Present Count field	The number of physical drives present on the controller.

Name	Description
Critical Disk Count field	The number of physical drives in a critical state on the controller.
Failed Disk Count field	The number of physical drives that have failed on the controller.

Step 12 In the **Capabilities** area, review the following information:

Name	Description
RAID Levels Supported field	The RAID levels supported by the controller. Raid 1 —Simple mirroring.

Step 13 In the **HW Configuration** area, review the following information:

Name	Description
Number of Backend Ports field	The number of SATA ports on the controller.

Viewing Physical Drive Info for Cisco Boot Optimized M.2 Raid Controller

Before you begin

The server must be powered on.

Step 1 In the **Navigation** pane, click the **Storage** menu.

Step 2 In the **Storage** menu, click the appropriate M.2 raid controller.

Step 3 Click **Physical Drive Info** tab and in the **Physical Drives** area, review the following information:

Name	Description
Controller column	The name of the PCIe slot in which the controller drive is located.
Physical Drive Number column	The physical drive number.

Name	Description
Status column	<p>The physical drive status. This can be one of the following:</p> <ul style="list-style-type: none"> • JBOD—The drive is in JBOD mode. • Failed—The drive was in use, but it has failed. • Offline—The drive is offline and cannot be accessed. • Online—The drive is in use as part of a drive group. • Predicted Failure—The drive is marked as due to fail soon by the controller. • Rebuild—The drive is currently being rebuilt. • Unknown—The drive status is unknown.
State column	<p>The state of the physical drive. This can be one of the following:</p> <ul style="list-style-type: none"> • JBOD— The physical drive is in JBOD state. • Online— The physical drive is in use as part of a drive group. • Failed— The physical drive is in failed state. • Rebuild— The physical drive is in rebuild state.
Health column	<p>The health of the physical drive. This can be one of the following:</p> <ul style="list-style-type: none"> • Good • Moderate Fault • Severe Fault <p>The health field includes both text and a color-coded icon. If the color-coded icon is:</p> <ul style="list-style-type: none"> • Green — It indicates normal operation. • Yellow — It is an informational message. • Red — It indicates warning, critical, and non-recoverable errors.
Drive Firmware column	The firmware version running on the drive.
Model column	The drive vendor name.
Type column	Whether the drive is a hard drive (HDD) or a solid state drive (SSD).

Step 4 In the **Health/Status** area, review the following information:

Name	Description
Health	The health of the physical drive. This can be one of the following: <ul style="list-style-type: none"> • Good • Moderate Fault • Severe Fault
State	The state of the physical drive. This can be one of the following: <ul style="list-style-type: none"> • JBOD— The physical drive is in JBOD state. • Online— The physical drive is in use as part of a drive group. • Failed— The physical drive is in failed state. • Rebuild— The physical drive is in rebuild state.
Status	The physical drive status. This can be one of the following: <ul style="list-style-type: none"> • JBOD—The drive is in JBOD mode. • Failed—The drive was in use, but it has failed. • Offline—The drive is offline and cannot be accessed. • Online—The drive is in use as part of a drive group. • Predicted Failure—The drive is marked as due to fail soon by the controller. • Rebuild—The drive is currently being rebuilt. • Unknown—The drive status is unknown.
Fault	If this field displays true , the drive is in the failed state.
Online	If this field displays true , the drive is in the online state.

Step 5 In the **Smart Information** area, review the following information:

Name	Description
Power Cycle Count	Number of power cycles that were performed on the drive from the time it was manufactured. Power Cycle Count this field.
Power On Hours	Total number of hours that the drive is in the 'Power On' mode.

Name	Description
Percentage Life Left	The number of write cycles remaining in a M.2 drive. For instance, if an M.2 drive is capable of 100 write cycles during its life time, and it has completed 15 writes, then the percentage of life left in the drive is 85%. Each percentage range is represented in a different color. For instance, green for 75% to 100% and red for 1 to 25%.
Wear Status in Days	The number of days an M.2 drive has gone through with the write cycles.
Operating Temperature (°C)	The current temperature of the drive at which the selected M.2 drive operates at the time of selection.
Percentage Reserved Capacity Consumed	The total capacity (out of the percentage reserved for it) consumed by the M.2 drive.
Time of Last Refresh	Time period since the drive was last refreshed.

Step 6 In the **Operation Status** area, review the following information:

Name	Description
Operation	The current operation that is in progress on the drive. This can be one of the following: <ul style="list-style-type: none"> • Rebuild in progress • Patrol read in progress
Progress in %	The progress of the operation, in percentage.

Step 7 In the **Inquiry Data** area, review the following information:

Name	Description
Product ID	The drive product ID. This field generally displays the drive model number.
Vendor	The vendor for the drive.
Drive Firmware	The active firmware version on the drive.
Drive Serial Number	The serial number for the drive.

Step 8 In the **General** area, review the following information:

Name	Description
Slot Number	The slot number in which the physical drive resides.
Raw Size	The capacity of drive in megabytes, including the space used for formatting.

Name	Description
Negotiated Link Speed	The speed of the link between the drive and the controller.
Media Type	The drive type is a solid state drive (SSD).
Interface Type	The drive interface type.
Enclosure Association	Indicates whether the drives are associated to the controller directly or not. This can be: Direct attached —drives are directly attached to the controller.

Viewing the Virtual Drive Info for Cisco Boot Optimized M.2 Raid Controller

Step 1 In the **Navigation** pane, click the **Storage** menu.

Step 2 In the **Storage** menu, click the M.2 Raid Controller.

Step 3 Select the **Virtual Drive Info** tab and in the **Virtual Drives** area, review the following information:

Name	Description
Virtual Drive Number column	The number of the virtual drive.
Name column	The name of the virtual drive.
Status column	The virtual drive state. This can be one of the following: <ul style="list-style-type: none"> • Partially Degraded—Virtual drive or physical drive rebuild is in progress. • Degraded—One or more spans of the drive has no redundancy. • Off Line—The drive is not visible to the host. • Unknown—The virtual drive state is unknown. • Optimal—The drive has full redundancy.

Name	Description
Health column	<p>The health of the virtual drive. This can be one of the following:</p> <ul style="list-style-type: none"> • Good • Moderate Fault • Severe Fault <p>The health field includes both text and a color-coded icon. If the color-coded icon is:</p> <ul style="list-style-type: none"> • Green — It indicates normal operation. • Yellow — It is an informational message. • Red — It indicates warning, critical, and non-recoverable errors.
Size column	The capacity of the drive in megabytes.
RAID Level column	<p>The RAID level on the virtual drive.</p> <p>Raid 1—Simple mirroring.</p>

Step 4 In the **General** area, review the following information:

Name	Description
Name	The name of the virtual drive.
Strip Size	The size of each strip, in KB.

Step 5 In the **Physical Drives** area, review the following information:

Name	Description
Physical Drive Number	The physical drive number.
State	<p>The state of the physical drive. This can be one of the following:</p> <ul style="list-style-type: none"> • JBOD— The physical drive is in JBOD state. • Online— The physical drive is in use as part of a drive group. • Failed— The physical drive is in failed state. • Rebuild— The physical drive is in rebuild state.

Name	Description
Status	<p>The physical drive status. This can be one of the following:</p> <ul style="list-style-type: none"> • JBOD—The drive is in JBOD mode. • Failed—The drive was in use, but it has failed. • Offline—The drive is offline and cannot be accessed. • Online—The drive is in use as part of a drive group. • Predicted Failure—The drive is marked as due to fail soon by the controller. • Rebuild—The drive is currently being rebuilt. • Unknown—The drive status is unknown.

Step 6 In the **Operation Status** area, review the following information:

Name	Description
Operation	<p>The current operation that is in progress on the drive. This can be:</p> <p>Rebuild in progress</p>
Progress in %	The progress of the operation, in percentage.

