



Storage Profiles

- [Storage Profiles, on page 1](#)
- [Disk Groups and Disk Group Configuration Policies, on page 2](#)
- [RAID Levels, on page 6](#)
- [Automatic Disk Selection, on page 7](#)
- [Supported LUN Modifications, on page 8](#)
- [Unsupported LUN Modifications, on page 8](#)
- [Disk Insertion Handling, on page 9](#)
- [Virtual Drive Naming, on page 11](#)
- [LUN Dereferencing, on page 11](#)
- [Controller Constraints and Limitations, on page 12](#)
- [Storage Profiles, on page 12](#)
- [Configuring Storage Profiles, on page 31](#)

Storage Profiles

To allow flexibility in defining the number of storage disks, roles and usage of these disks, and other storage parameters, you can create and use storage profiles. A storage profile encapsulates the storage requirements for one or more service profiles. LUNs configured in a storage profile can be used as boot LUNs or data LUNs, and can be dedicated to a specific server. You can also specify a local LUN as a boot device. However, LUN resizing is not supported. The introduction of storage profiles allows you to do the following:

- Configure multiple virtual drives and select the physical drives that are used by a virtual drive. You can also configure the storage capacity of a virtual drive.
- Configure the number, type and role of disks in a disk group.
- Associate a storage profile with a service profile.

You can create a storage profile both at an org level and at a service-profile level. A service profile can have a dedicated storage profile as well as a storage profile at an org level.

Disk Groups and Disk Group Configuration Policies

Servers in a chassis can use storage that is centralized in that chassis. You can select and configure the disks to be used for storage. A logical collection of these physical disks is called a disk group. Disk groups allow you to organize local disks. The storage controller controls the creation and configuration of disk groups.

A disk group configuration policy defines how a disk group is created and configured. The policy specifies the RAID level to be used for the disk group. It also specifies either a manual or an automatic selection of disks for the disk group, and roles for disks. You can use a disk group policy to manage multiple disk groups. However, a single disk group can be managed only by one disk group policy.

A hot spare is an unused extra disk that can be used by a disk group in the case of failure of a disk in the disk group. Hot spares can be used only in disk groups that support a fault-tolerant RAID level. In addition, a disk can be allocated as a global hot spare, which means that it can be used by any disk group.

Virtual Drives

A disk group can be partitioned into virtual drives. Each virtual drive appears as an individual physical device to the Operating System.

All virtual drives in a disk group must be managed by using a single disk group policy.

Configuration States

Indicates the configuration states of a virtual drive. Virtual drives can have the following configuration states:

- Applying—Creation of the virtual drive is in progress.
- Applied—Creation of the virtual drive is complete, or virtual disk policy changes are configured and applied successfully.
- Failed to apply—Creation, deletion, or renaming of a virtual drive has failed due to errors in the underlying storage subsystem.
- Orphaned—The service profile that contained this virtual drive is deleted or the service profile is no longer associated with a storage profile.



Note Orphaned LUNs cannot be used for booting OS. Although an image can be installed on these LUNs, booting from these drives will fail. To use any specific orphaned LUN, you must reassociate the storage profile, which will return it to the “Equipped” presence state.

- Not in use—The service profile that contained this virtual drive is in the disassociated state.

Deployment States

Indicates the actions that you are performing on virtual drives. Virtual drives can have the following deployment states:

- No action—No pending work items for the virtual drive.

- **Creating**—Creation of the virtual drive is in progress.
- **Deleting**—Deletion of the virtual drive is in progress.
- **Modifying**—Modification of the virtual drive is in progress.
- **Apply-Failed**—Creation or modification of the virtual drive has failed.

Operability States

Indicates the operating condition of a virtual drive. Virtual drives can have the following operability states:

- **Optimal**—The virtual drive operating condition is good. All configured drives are online.
- **Degraded**—The virtual drive operating condition is not optimal. One of the configured drives has failed or is offline.
- **Cache-degraded**—The virtual drive has been created with a write policy of **write back** mode, but the BBU has failed, or there is no BBU.



Note This state does not occur if you select the **always write back** mode.

- **Partially degraded**—The operating condition in a RAID 6 virtual drive is not optimal. One of the configured drives has failed or is offline. RAID 6 can tolerate up to two drive failures.
- **Offline**—The virtual drive is not available to the RAID controller. This is essentially a failed state.
- **Unknown**—The state of the virtual drive is not known.

Presence States

Indicates the presence of virtual drive components. Virtual drives have the following presence states:

- **Equipped**—The virtual drive is available.
- **Mismatched**—A virtual drive deployed state is different from its configured state.
- **Missing**—Virtual drive is missing.

Configuring a Disk Group Policy

You can configure the disks in a disk group policy automatically or manually.

Procedure

-
- Step 1** In the **Navigation** pane, click **Storage**.
 - Step 2** Expand **Storage > Storage Provisioning > Storage Policies**
 - Step 3** Expand the node for the organization where you want to create the disk group policy.
 - Step 4** Right-click **Disk Group Policies** in the organization and select **Create Disk Group Policy**.
 - Step 5** In the **Create Disk Group Policy** dialog box, specify the following:

Name	Description
Name field	<p>The name of the policy</p> <p>This name can be between 1 and 16 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object has been saved.</p>
Description field	<p>A description of the policy. We recommend that you include information about where and when the policy should be used.</p> <p>Enter up to 256 characters. You can use any characters or spaces except ` (accent mark), \ (backslash), ^ (carat), " (double quote), = (equal sign), > (greater than), < (less than), or ' (single quote).</p>
RAID Level drop-down list	<p>This can be one of the following:</p> <ul style="list-style-type: none"> • RAID 0 Striped • RAID 1 Mirrored • RAID 5 Striped Parity • RAID 6 Striped Dual Parity • RAID 10 Mirrored and Striped <p>Note When you create a disk group with RAID 1 policy and configure four disks for it, a RAID 1E configuration is created internally by the storage controller.</p>

Step 6

To automatically configure the disks in a disk group policy, select **Disk Group Configuration (Automatic)** and specify the following:

Name	Description
Number of drives field	<p>Specifies the number of drives for the disk group.</p> <p>The range for drives is from 0 to 24 drives. Unspecified is the default number of drives. When you select the number of drives as Unspecified, the number of drives will be selected according to the disk selection process.</p>
Drive Type field	<p>Drive type for the disk group. You can select:</p> <ul style="list-style-type: none"> • HDD • SSD • Unspecified <p>Unspecified is the default type of drive. When you select the drive type as Unspecified, the first available drive is selected. After this drive is selected, subsequent drives will be of a compatible type. For example, if the first was SSD, all subsequent drives would be SSD.</p>
Number of Hot Spares field	<p>Number of dedicated hot spares for the disk group.</p>

Name	Description
	The range for dedicated hot spares is from 0 to 24 hot spares. Unspecified is the default number of dedicated hot spares. When you select the number of dedicated hot spares as Unspecified , the hot spares will be selected according to the disk selection process.
Min Drive Size field	Minimum drive size for the disk group. Only disks that match this criteria are available for selection. The range for minimum drive size is from 0 to 10240 GB. Unspecified is the default minimum drive size. When you select the minimum drive size as Unspecified , drives of all sizes will be available for selection.

Step 7 To manually configure the disks in a disk group policy, select **Disk Group Configuration (Manual)** and do the following:

- On the icon bar to the right of the table, click +
- In the **Create Local Disk Configuration Reference** dialog box, complete the following fields:

Name	Description
Slot field	Slot for which the local disk reference is configured.
Role field	Role of the local disk in the disk group. You can select: <ul style="list-style-type: none"> • Normal • Dedicated Hot Spare • Global Hot Spare
Span ID field	Specifies the ID of the span group to which the disk belongs. Disks belonging to a single span group can be treated as a single disk with a larger size. The values range from 0 to 8. For RAID-10, RAID-50, and RAID-60, minimum 2 spans are required and maximum 8 spans are supported. You can also set the Span ID as Unspecified , when spanning information is not required.

Step 8 In the **Virtual Drive Configuration** area, specify the following:

Name	Description
Strip Size (KB) field	Stripe size for a virtual drive. This can only be Platform Default .
Access Policy field	Access policy for a virtual drive. This can be one of the following: <ul style="list-style-type: none"> • Platform Default • Read Write • Read Only • Blocked
Read Policy field	Read policy for a virtual drive. This can be one of the following: <ul style="list-style-type: none"> • Platform Default

Name	Description
	<ul style="list-style-type: none"> • Read Ahead • Normal
Write Cache Policy field	Write-cache-policy for a virtual drive. This can be one of the following: <ul style="list-style-type: none"> • Platform Default • Write Through • Write Back Good Bbu • Always Write Back
IO Policy field	I/O policy for a virtual drive. This can be one of the following: <ul style="list-style-type: none"> • Platform Default • Direct • Cached
Drive Cache field	State of the drive cache. This can be one of the following: <ul style="list-style-type: none"> • Platform Default • No Change • Enable • Disable

All virtual drives in a disk group should be managed by using the same disk group policy.

Step 9 Click **OK**.

Note When you accept the virtual drive (VD) default values and associate the disk group policy to a service profile, you can modify the VD configuration after it is associated to a service profile. If you modify the VD default values from the WebBIOS to use the non-default values, a properties fault is not generated to verify the changed values.

RAID Levels

The RAID level of a disk group describes how the data is organized on the disk group for the purpose of ensuring availability, redundancy of data, and I/O performance.

The following are features provided by RAID:

- **Striping**—Segmenting data across multiple physical devices. This improves performance by increasing throughput due to simultaneous device access.

- Mirroring—Writing the same data to multiple devices to accomplish data redundancy.
- Parity—Storing of redundant data on an additional device for the purpose of error correction in the event of device failure. Parity does not provide full redundancy, but it allows for error recovery in some scenarios.
- Spanning—Allows multiple drives to function like a larger one. For example, four 20 GB drives can be combined to appear as a single 80 GB drive.

The supported RAID levels include the following:

- RAID 0 Striped—Data is striped across all disks in the array, providing fast throughput. There is no data redundancy, and all data is lost if any disk fails.
- RAID 1 Mirrored—Data is written to two disks, providing complete data redundancy if one disk fails. The maximum array size is equal to the available space on the smaller of the two drives.
- RAID 5 Striped Parity—Data is striped across all disks in the array. Part of the capacity of each disk stores parity information that can be used to reconstruct data if a disk fails. RAID 5 provides good data throughput for applications with high read request rates.

RAID 5 distributes parity data blocks among the disks that are part of a RAID-5 group and requires a minimum of three disks.

- RAID 6 Striped Dual Parity—Data is striped across all disks in the array and two sets of parity data are used to provide protection against failure of up to two physical disks. In each row of data blocks, two sets of parity data are stored.

Other than addition of a second parity block, RAID 6 is identical to RAID 5. A minimum of four disks are required for RAID 6.

- RAID 10 Mirrored and Striped—RAID 10 uses mirrored pairs of disks to provide complete data redundancy and high throughput rates through block-level striping. RAID 10 is mirroring without parity and block-level striping. A minimum of four disks are required for RAID 10.
- RAID 50 Striped Parity and Striped—Data is striped across multiple striped parity disk sets to provide high throughput and multiple disk failure tolerance.
- RAID 60 Striped Dual Parity and Striped—Data is striped across multiple striped dual parity disk sets to provide high throughput and greater disk failure tolerance.

Automatic Disk Selection

When you specify a disk group configuration, and do not specify the local disks in it, Cisco UCS Manager determines the disks to be used based on the criteria specified in the disk group configuration policy. Cisco UCS Manager can make this selection of disks in multiple ways.

When all qualifiers match for a set of disks, then disks are selected sequentially according to their slot number. Regular disks and dedicated hot spares are selected by using the lowest numbered slot.

The following is the disk selection process:

1. Iterate over all local LUNs that require the creation of a new virtual drive. Iteration is based on the following criteria, in order:
 - a. Disk type

- b. Minimum disk size from highest to lowest
 - c. Space required from highest to lowest
 - d. Disk group qualifier name, in alphabetical order
 - e. Local LUN name, in alphabetical order
2. Select regular disks depending on the minimum number of disks and minimum disk size. Disks are selected sequentially starting from the lowest numbered disk slot that satisfies the search criteria.



Note If you specify **Any** as the type of drive, the first available drive is selected. After this drive is selected, subsequent drives will be of a compatible type. For example, if the first drive was SATA, all subsequent drives would be SATA. Cisco UCS Manager Release 2.5 supports only SATA and SAS.

Cisco UCS Manager Release 2.5 does not support RAID migration.

3. Select dedicated hot spares by using the same method as normal disks. Disks are only selected if they are in an **Unconfigured Good** state.
4. If a provisioned LUN has the same disk group policy as a deployed virtual drive, then try to deploy the new virtual drive in the same disk group. Otherwise, try to find new disks for deployment.

Supported LUN Modifications

Some modifications that are made to the LUN configuration when LUNs are already deployed on an associated server are supported.

The following are the types of modifications that can be performed:

- Creation of a new virtual drive.
- Deletion of an existing virtual drive, which is in the orphaned state.
- Non-disruptive changes to an existing virtual drive. These changes can be made on an existing virtual drive without loss of data, and without performance degradation:
 - Policy changes. For example, changing the write cache policy.
 - Modification of boot parameters

The removal of a LUN will cause a warning to be displayed. Ensure that you take action to avoid loss of data.

Unsupported LUN Modifications

Some modifications to existing LUNs are not possible without destroying the original virtual drive and creating a new one. All data is lost in these types of modification, and these modifications are not supported.

Disruptive modifications to an existing virtual drive are not supported. The following are unsupported disruptive changes:

- Any supported RAID level change that can be handled through reconstruction. For example, RAID0 to RAID1.
- Increasing the size of a virtual drive through reconstruction.
- Addition and removal of disks through reconstruction.
- **Expand To Available** option is not supported for already deployed LUN.

Destructive modifications are also not supported. The following are unsupported destructive modifications:

- RAID-level changes that do not support reconstruction. For example, RAID5 to RAID1.
- Shrinking the size of a virtual drive.
- RAID-level changes that support reconstruction, but where there are other virtual drives present on the same drive group.
- Disk removal when there is not enough space left on the disk group to accommodate the virtual drive.
- Explicit change in the set of disks used by the virtual drive.

Disk Insertion Handling

When the following sequence of events takes place:

1. The LUN is created in one of the following ways:
 - a. You specify the slot specifically by using a local disk reference
 - b. The system selects the slot based on criteria specified by you
2. The LUN is successfully deployed, which means that a virtual drive is created, which uses the slot.
3. You remove a disk from the slot, possibly because the disk failed.
4. You insert a new working disk into the same slot.

The following scenarios are possible:

- [Non-Redundant Virtual Drives, on page 10](#)
- [Redundant Virtual Drives with No Hot Spare Drives, on page 10](#)
- [Redundant Virtual Drives with Hot Spare Drives, on page 10](#)
- [Replacing Hot Spare Drives, on page 10](#)
- [Inserting Physical Drives into Unused Slots, on page 11](#)

Non-Redundant Virtual Drives

For non-redundant virtual drives (RAID 0), when a physical drive is removed, the state of the virtual drive is **Inoperable**. When a new working drive is inserted, the new physical drive goes to an **Unconfigured Good** state.

For non-redundant virtual drives, there is no way to recover the virtual drive. You must delete the virtual drive and re-create it.

Redundant Virtual Drives with No Hot Spare Drives

For redundant virtual drives (RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, RAID 60) with no hot spare drives assigned, virtual drive mismatch, virtual drive member missing, and local disk missing faults appear until you insert a working physical drive into the same slot from which the old physical drive was removed.

If the physical drive size is greater than or equal to that of the old drive, the storage controller automatically uses the new drive for the virtual drive. The new drive goes into the **Rebuilding** state. After rebuild is complete, the virtual drive goes back into the **Online** state.

Redundant Virtual Drives with Hot Spare Drives

For redundant virtual drives (RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, RAID 60) with hot spare drives assigned, when a drive fails, or when you remove a drive, the dedicated hot spare drive, if available, goes into the **Rebuilding** state with the virtual drive in the **Degraded** state. After rebuilding is complete, that drive goes to the **Online** state.

Cisco UCSM raises a disk missing and virtual drive mismatch fault because although the virtual drive is operational, it does not match the physical configuration that Cisco UCSM expects.

if you insert a new disk in the slot with the disk missing, automatic copy back starts from the earlier hot spare disk to the newly inserted disk. After copy back, the hot spare disk is restored. In this state all faults are cleared.

If automatic copy back does not start, and the newly inserted disk remains in the **Unconfigured Good**, **JBOD**, or **Foreign Configuration** state, remove the new disk from the slot, reinsert the earlier hot spare disk into the slot, and import foreign configuration. This initiates the rebuilding process and the drive state becomes **Online**. Now, insert the new disk in the hot spare slot and mark it as hot spare to match it exactly with the information available in Cisco UCSM.

Replacing Hot Spare Drives

If a hot spare drive is replaced, the new hot spare drive will go to the **Unconfigured Good**, **Unconfigured Bad**, **JBOD**, or **Foreign Configuration** state.

Cisco UCSM will raise a virtual drive mismatch or virtual drive member mismatch fault because the hot spare drive is in a state different from the state configured in Cisco UCSM.

You must manually clear the fault. To do this, you must perform the following actions:

1. Clear the state on the newly inserted drive to **Unconfigured Good**.
2. Configure the newly inserted drive as a hot spare drive to match what is expected by Cisco UCSM.

Inserting Physical Drives into Unused Slots

If you insert new physical drives into unused slots, neither the storage controller nor Cisco UCSM will make use of the new drive even if the drive is in the **Unconfigured Good** state and there are virtual drives that are missing good physical drives.

The drive will simply go into the **Unconfigured Good** state. To make use of the new drive, you will need to modify or create LUNs to reference the newly inserted drive.

Virtual Drive Naming

When you use Cisco UCS Manager to create a virtual drive, Cisco UCS Manager assigns a unique ID that can be used to reliably identify the virtual drive for further operations. Cisco UCS Manager also provides the flexibility to provide a name to the virtual drive at the time of service profile association. Any virtual drive without a service profile or a server reference is marked as an orphan virtual drive.

In addition to a unique ID, a name is assigned to the drive. Names can be assigned in two ways:

- When configuring a virtual drive, you can explicitly assign a name that can be referenced in storage profiles.
- If you have not preprovisioned a name for the virtual drive, Cisco UCS Manager generates a unique name for the virtual drive.

You can rename an orphan virtual drive on a blade or a rack server that are not referenced by any service profile or server.



Note The renaming an orphan virtual drive is not supported for Cisco boot optimized M.2 Raid controller (UCS-M2-HWRAID).

LUN Dereferencing

A LUN is dereferenced when it is no longer used by any service profile. This can occur as part of the following scenarios:

- The LUN is no longer referenced from the storage profile
- The storage profile is no longer referenced from the service profile
- The server is disassociated from the service profile
- The server is decommissioned

When the LUN is no longer referenced, but the server is still associated, re-association occurs.

When the service profile that contained the LUN is disassociated, the LUN state is changed to **Not in use**.

When the service profile that contained the LUN is deleted, the LUN state is changed to **Orphaned**.

Controller Constraints and Limitations

- The following table provides the maximum supported virtual drives for servers:

Servers/Storage Controllers	Maximum Virtual Drives
UCSB-MRAID12G-M6	16
UCSC-C240-M5, UCSC-C480-M5	32
UCS-S3260-M5, UCSC-C3X60-M4, UCSC-C3K-M4	64
UCSC-C240-M4, UCSC-C240-M3, UCSC-C24-M3	24
UCSB-MRAID12G	16
UCS-M2-HWRAID	2
For all other servers.	18



Note

- Storage controllers support the check max feature.
- When servers with multiple storage controllers are managed by the same storage profile, the maximum virtual drives are limited to the maximum value supported by the server.

- In Cisco UCS Manager Release 2.2(4), blade servers do not support drives with a block size of 4K, but rack-mount servers support such drives. If a drive with a block size of 4K is inserted into a blade server, discovery fails and the following error message appears: Unable to get Scsi Device Information from the system.
- In Cisco UCS Manager Release 3.1(2) and later releases, RAID Controller that does not support Out of band inventory (OOB) in C240 M4 servers, display Operability as NA and Drive State as Unknown.

Storage Profiles

Creating a Storage Profile

You can create storage profile policies from the **Storage** tab in the **Navigation** pane. Additionally, you can also configure the default storage profile that is specific to a service profile from the **Servers** tab.

**Caution**

If you have a Cisco UCS blade or rack server with a default local disk configuration present in a Service Profile or Service Profile Template from an earlier release of UCS Manager and you upgrade to the 3.1 release and later releases, you can successfully create a Storage Profile with local LUNs in the same Service Profile or Service Profile Template if you change the Local Disk Configuration Default policy to **Any Configuration** instead of RAID level options in the local disk policy. The legacy LUN is thereafter part of the storage inventory.

Procedure

- Step 1** In the **Navigation** pane, click **Storage**.
- Step 2** Expand **Storage > Storage Profiles**
- Step 3** Expand the node for the organization where you want to create the storage profile.
If the system does not include multi tenancy, expand the **root** node.
- Step 4** Right-click the organization and select **Create Storage Profile**.
- Step 5** In the **Create Storage Profile** dialog box, specify the storage profile **Name**. You can provide an optional **Description** for this storage profile.
- Step 6** (Optional) In the **LUNs** area, create **Local LUNs** and add them to this storage profile.
See [Configuring Local LUNs, on page 14](#) for more information.
- Step 7** In the **LUNs** area, create **Controller Definitions** and add them to this storage profile.
See [Creating a Storage Profile PCH Controller Definition, on page 21](#) for more information.
- Step 8** In the **LUNs** area, create **Security Policy** and add them to this storage profile.
See [Creating a Local Security Policy](#) and [Creating a Remote Security Policy](#) for more information.
- Step 9** Click **OK**.

Creating a Specific Storage Profile

Procedure

- Step 1** Expand **Servers > Service Profiles**.
- Step 2** Expand the node for the organization that contains the service profile for which you want to create a specific storage profile.
If the system does not include multi tenancy, expand the **root** node.
- Step 3** Choose the service profile for which you want to create a specific storage profile.
- Step 4** In the **Work** pane, click the **Storage > LUN Configuration** tab.
- Step 5** In the **Actions** area, click **Modify Storage Profile**.

- Step 6** In the **Modify Storage Profile** dialog box, click the **Specific Storage Profile** tab.
- Step 7** Click **Create Specific Storage Profile**.
- Step 8** (Optional) In the **Specific Storage Profile** area, complete the **Description** field to set the description of the storage profile.
- Each service profile can have only one specific storage profile. Hence, the name of this storage profile is provided by default.
- Step 9** In the **Storage Items** area, **Create Local LUNs** and add them to this storage profile.
- Step 10** Click **OK**.
- Step 11** If a confirmation dialog box displays, click **Yes**.
-

Deleting a Storage Profile

Procedure

- Step 1** In the **Navigation** pane, click **Storage**.
- Step 2** Expand **Storage > Storage Profiles**
- Step 3** Expand the node for the organization that contains the storage profile that you want to delete.
- Step 4** Right-click the storage profile that you want to delete and select **Delete**.
- Step 5** Click **Yes** in the confirmation box that appears.
-

Local LUNs

Configuring Local LUNs

You can create local LUNs within a storage profile policy from the **Storage** tab in the **Navigation** pane. Additionally, you can also create local LUNs within the default storage profile that is specific to a service profile from the **Servers** tab.

Procedure

- Step 1** In the **Navigation** pane, click **Storage**.
- Step 2** Expand **Storage > Storage Profiles**
- Step 3** Expand the node for the organization that contains the storage profile within which you want to create a local LUN.
- Step 4** In the **Work** pane, click the **General** tab.
- Step 5** In the **Actions** area, click **Create Local LUN**.
- Step 6** In the Create Local LUN dialog box, complete the following fields:

Name	Description
Create Local LUN option	(Appears when you create a local LUN) Selected by default when you create a local LUN.
Prepare Claim Local LUN option	(Appears when you create a local LUN) Select when you want to claim an orphan LUN.
Name field	<p>The name of the local LUN.</p> <p>This name can be between 1 and 10 alphanumeric characters. You cannot use spaces or any special characters other than - (hyphen), _ (underscore), : (colon), and . (period), and you cannot change this name after the object has been saved.</p> <p>Note If the name given in Prepare Claim Local LUN is different from the name to be claimed, this LUN name and the Virtual drive name appearing in the LUN properties are different.</p>
Size (GB) field	<p>Size of this LUN in GB.</p> <p>Note You do not need to specify a LUN size while claiming an orphaned LUN.</p>
Fractional Size (MB) field	The fractional size of this LUN in MB.
Auto Deploy radio buttons	<p>Whether the local LUN should be automatically deployed or not. This can be one of the following:</p> <ul style="list-style-type: none"> • Auto Deploy —Automatically deploys the local LUN. • No Auto Deploy —Does not automatically deploy the local LUN.
Expand To Available checkbox	<p>(Only available for rack and blade servers) Specifies that this LUN can be expanded to use the entire available disk group.</p> <p>For each service profile, only one LUN can use this option.</p> <p>Expand To Available option is not supported for already deployed LUN.</p>
Select Disk Group Configuration drop-down list	Choose the disk group configuration to be applied to this local LUN from the drop-down list.
Create Disk Group Policy link	Displays the Create Disk Group Policy dialog box to create a new disk group.

Step 7 (Optional) Click **Create Disk Group Policy** to create a new disk group policy for this local LUN.

Step 8 Click **OK**.

Displaying Details of All Local LUNs Inherited By a Service Profile

Storage profiles can be defined under org and as a dedicated storage profile under service profile. Thus, a service profile inherits local LUNs from both possible storage profiles. It can have a maximum of 2 such local LUNs. You can display the details of all local LUNs inherited by a service profile by using the following command:

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization that contains the service profile that you want to display.
- Step 4** Choose the service profile whose inherited local LUNs you want to display.
- Step 5** In the **Work** pane, click the **Storage** tab.
- Step 6** Click the **Storage Profiles** subtab, and then click the **Local LUNs** tab.

Displays the following detailed information about all the local LUNs inherited by the specified service profile:

- **Name**—LUN name in the storage profile.
 - **RAID Level**—Summary of the RAID level of the disk group used.
 - **Size (MB)**—Size, in MB, of the LUN specified in the storage profile.
 - **Config State**—State of LUN configuration. The states can be one of the following:
 - **Applying**—Admin state is online, the LUN is associated with a server, and the virtual drive is being created.
 - **Applied**—Admin state is online, the LUN is associated with a server, and the virtual drive is created.
 - **Apply Failed**—Admin stage is online, the LUN is associated with a server, but the virtual drive creation failed.
 - **Not Applied**—The LUN is not associated with a server, or the LUN is associated with a service profile, but admin state is undeployed.
 - **Deploy Name**—The virtual drive name after deployment.
 - **LUN ID**—LUN ID.
 - **Drive State**—State of the virtual drive. The states are:
 - **Unknown**
 - **Optimal**
 - **Degraded**
 - **Inoperable**
 - **Partially Degraded**
-

Deleting Local LUNs

Procedure

-
- Step 1** In the **Navigation** pane, click **Storage**.
- Step 2** Expand **Storage > Storage Profiles**
- Step 3** Expand the node for the organization that contains the storage profile from which you want to delete a local LUN.
- Step 4** Expand **Local LUNs** for the storage profile that you want and select the LUN that you want to delete.
- Step 5** Right-click the LUN that you want to delete and select **Delete**.
A confirmation dialog box appears.
- Step 6** Click **Yes**.
-

Configuring Aero Controllers

Autoconfiguration Mode for Aero Storage Controllers

Cisco UCS C220M6/C240M6 C-series M6 servers support Aero PCIe SAS316-port storage controllers for Direct Attached Storage. Aero controllers support an Autoconfiguration mode in which the state of a newly inserted disk is automatically moved to the Unconfigured-good state.

Because of this, you can choose whether or not to use Autoconfiguration by creating a Storage Profile and associating it with the server. The default is that the automatic configuration feature is disabled, which retains the drive state when the server is rebooted.

If Autoconfiguration is used, you must select a drive state from one of the following:

- unconfigured-good
- JBOD
- raid0-writeback

This is because the Aero firmware changes the behavior of systemPD to EPD-PT. EPD-PT is internally a RAID0 volume without any drive DDF metadata. The controller stores the metadata for identifying it as a RAID0 volume. The EPD-PT drives are considered as JBOD drives so the drive status is reported as JBOD and online.

The table below shows the behavior of Autoconfiguration in different scenarios.

Autoconfig Mode	Reboot/OCR	Hotplug	User Action
unconfigured-good (OFF)	<ul style="list-style-type: none"> • All UG drives remain UG. • All previously configured EPD remain EPD. 	<ul style="list-style-type: none"> • Inserted drive remains UG • EPD from a different server remains UG on this controller. 	Disabling Autoconfig has no impact on the existing configuration

Autoconfig Mode	Reboot/OCR	Hotplug	User Action
jbod (EPD-PT)	<ul style="list-style-type: none"> All UG are converted to EPD. Any UG will remain UG until next boot. During boot, UG is converted to UPD. 	Inserted device converted to EPD.	All UG drives on the controller while running Autoconfig will be converted to EPD.
raid0-writeback (RAID0 WB)	<ul style="list-style-type: none"> All UG converted to RAID0 WB. Any UG will remain UG across control boot. 	Inserted device converted to RAID0 WB.	All UG drives on the controller while running Autoconfig will be converted to RAID0 WB.

Selecting EPD-PT (JBOD) as the default configuration does not retain the unconfigured-good state across host reboot. The drive state can be retained by disabling the automatic configuration feature. If the Autoconfig option is used, the default automatic configuration will always mark a drive as unconfigured-good.

When Autoconfig is selected, then the drive is configured to the desired drive state, the JBOD and unconfigured drives will set the drive state accordingly on the next controller boot or OCR,

The following table shows sample use cases for different Autoconfig scenarios.

Use Case Scenario	Autoconfig Option
Using the server for JBOD Only (for example: Hyper converged, Hadoop data node etc)	jbod
Using the server for RAID volume (for example: SAP HANA database)	unconfigured-good
Using the server for Mixed JBOD and RAID volume	unconfigured-good
Using the server for per drive R0 WB (for example: Hadoop data node)	raid0-writeback

Creating an Autoconfiguration Storage Profile

The Autoconfiguration (Auto Config) mode option for storage is only available on Cisco UCS M6 servers with Aero controllers.

Procedure

-
- Step 1** In the **Navigation** pane, click **Storage**.
 - Step 2** Go to **Profiles**. Expand the root node.
 - Step 3** Right click on **Storage**.

- Step 4** In the **Create Storage Profile** menu, name the profile. The menu will come up with Auto Config Mode marked as **Unspecified**.
- Step 5** To enable the Auto Config Mode option with a specific state to be retained on reboot, de-select **Unspecified**, then choose the desired state: Unconfigured Good, JBOD, or RAID 0. The state selection will be pushed to the BMC when the system is rebooted.
- If Auto Config is left as **Unspecified**, it will retain whatever state was configured prior to reboot.
- Note** Service profile association will fail if no Aero controllers are present.
- Step 6** Click **OK**.
-

SPDM Authentication

The Security Protocol and Data Model (SPDM) is used by the BMC for authentication with the storage controller. It requires that the storage controller firmware is secure booted as well as having a Broadcom certificate chain installed in the slot0. During a firmware update, the Broadcom firmware will retain the older measurements for the storage firmware until the OCR or host reboots. If device authentication fails, the firmware will allow only inventory related commands and no set operations can be performed.

PCH Controller Definitions

PCH SSD Controller Definition

Cisco UCS Manager Platform Controller Hub (PCH) Solid State Drive (SSD) Controller Definition provides a local storage configuration in storage profiles where you can configure all the disks in a single RAID or in a JBOD disk array.



Note You must re-acknowledge the server if you remove or insert disks managed by the PCH Controller.

The PCH Controller Definition configuration provides the following features:

- Ability to configure a single LUN RAID across two internal SSDs connected to the onboard PCH controller
- A way to configure the controller in two modes: AHCI (JBOD) and SWRAID (RAID).
- Ability to configure the PCH storage device in an Embedded Local LUN and Embedded Local Disk boot policy so precision control for boot order is achieved even with the presence of other bootable local storage devices in the server. Do not use the Local LUN or the Local JBOD options to boot from PCH disks
- Scrub policy support for the internal SSD drives. This is applicable only for the SWRAID mode. This does not apply for the AHCI and NORAIID of PCH Controller modes. *See the UCS Manager Server Management Guide.*
- Firmware upgrade support for the internal SSD drives.
 - For M4 and below servers, disk firmware upgrade is supported only when the PCH Controller is in SWRAID mode. It is not supported for AHCI mode.

You can configure PCH controller SSDs in a storage profile policy. You can enable or disable protect configuration which saves the LUN configuration even after a service profile disassociation. You choose a controller mode. The PCH controller configuration supports only these two RAID options: RAID0 and RAID1. Use No RAID configuration option for AHCI mode where all the disks connected to the controller configured as JBOD disks. The configuration deployment happens as part of the storage profile association to a service profile process.

Cisco UCS Manager supports PCH managed internal SSDs in the following M4 servers:

- UCSC-C240-M4L
- UCSC-C240-M4SX

For the PCH Controller Definition configuration in a Cisco UCS Manager boot policy two new devices exist to select: PCH LUN and PCH Disk. **EmbeddedLocalLun** represents the boot device in SWRAID mode and **EmbeddedLocalDisk** represent the boot devices in AHCI mode.

The system uses the same scrub policy is used to scrub supported SSDs. If the scrub is Yes, configured LUNs are destroyed as part of disassociation or re-discovery. If the scrub is No, configured LUNs are saved during disassociation and re-discovery.

Cisco UCS Manager supports firmware upgrade for the internal SSDs only when the PCH Controller is in SWRAID mode. It is not supported in AHCI mode.

FCH Controller Configuration

Fusion Controller Hub (FCH) SSD Controller Definition provides a local storage configuration in storage profiles for AMD based Cisco UCS C125 M5 Server. For AMD processor based servers, the PCH controller is referred to as FCH controller. The controller type remains as PCH in the Cisco UCS Manager GUI.

The FCH Controller works the same as PCH Controller except for the following differences:

- FCH is only in AHCI (JBOD) mode.



Note Cisco UCS Manager GUI shows the RAID support as **RAID0**, **RAID1**, however, Cisco UCS C125 M5 Server supports only AHCI mode.



Note You must re-acknowledge the server if you remove or insert disks managed by the PCH Controller.

- There are two FCH controllers:
 - First PCH controller manages SATA disks in the front panel (in the absence of a separate PCIe storage controller)
 - Second PCH controller manages the M.2 SSDs



Note For Cisco UCS C125 M5 Server, the PCH IDs are 3 and 4.



Note Further information and procedures related to PCH controller in this document are applicable to both Intel and AMD based servers.

Creating a Storage Profile PCH Controller Definition

The PCH Controller Definition provides a storage configuration in Storage Profiles where you can configure internal SSDs connected to a PCH controller. You create a name for the controller definition, specify whether you want the storage profile to retain the configuration even if the storage profile is disassociated from the service profile, and chose the RAID level to indicate the controller mode.

Procedure

- Step 1** In the **Navigation** pane, click **Storage > Storage Profiles**.
- Step 2** Choose the storage profile where you want to create the controller definition.
- Step 3** Click the **Controller Definitions** tab and then click **Add** at the bottom of the panel or right-click the storage profile and select **Create Controller Definition**.
- Step 4** In **Create Controller Definition** dialog box, configure the following information:

Name	Description
Name field	<p>The name of the storage controller.</p> <p>Note Once you save a PCH Controller Definition, you cannot modify the name from the General Tab Properties area.</p> <p>Enter up to 16 characters. You can use any alphanumeric characters. Special characters and spaces are not supported.</p>
Protect Configuration check box	<p>If checked, the storage profile retains the configuration even if the storage profile is disassociated from the service profile.</p> <p>Note If you disassociate the storage profile from a service profile with this option enabled, and then associate it with a new service profile that includes a local disk configuration policy with different properties, the server returns a configuration mismatch error and the association fails.</p>

Name	Description
RAID Level drop-down list	

Name	Description
	<p>This can be one of the following disk policy modes:</p> <ul style="list-style-type: none"> • Disable Local Storage—(Supported for PCH SSD Controller Definition) This disk policy mode is to disable the SATA AHCI Controller. This mode can be set only when disks are not present under the SATA AHCI controller. To re-enable this controller and to bring the controller back to its default value (AHCI), you can select No RAID or No Local Storage mode. • No Local Storage—(Supported for PCH SSD Controller Definition) For a diskless server or a SAN only configuration. If you select this option, you cannot associate any service profile which uses this policy with a server that has a local disk. • RAID 0 Striped—(Supported for PCH SSD Controller Definition) Data is striped across all disks in the array, providing fast throughput. There is no data redundancy, and all data is lost if any disk fails. • RAID 1 Mirrored—(Supported for PCH SSD Controller Definition) Data is written to two disks, providing complete data redundancy if one disk fails. The maximum array size is equal to the available space on the smaller of the two drives. • Any Configuration—(Supported for PCH SSD Controller Definition) For a server configuration that carries forward the local disk configuration without any changes. • No RAID—(Supported for PCH SSD Controller Definition) All the disks can be used individually without interdependency similar to JBOD disks. If you choose No RAID and you apply this policy to a server that already has an operating system with RAID storage configured, the system does not remove the disk contents. Therefore, there may be no visible differences on the server after you apply the No RAID mode. This can lead to a mismatch between the RAID configuration in the policy and the actual disk configuration shown in the Inventory > Storage tab for the server. To make sure that any previous RAID configuration information is removed from a disk, apply a scrub policy that removes all disk information after you apply the No RAID configuration mode. • RAID 5 Striped Parity—(Not supported for PCH SSD Controller Definition) Data is striped across all disks in the array. Part of the capacity of each disk stores parity information that can be used to reconstruct data if a disk fails. RAID 5 provides good data throughput for applications with high read request rates. • RAID 6 Striped Dual Parity—(Not supported for PCH SSD Controller Definition) Data is striped across all disks in the array and two parity disks are used to provide protection against the failure of up to two physical disks. In each row of data blocks, two sets of parity data are stored.

Name	Description
	<ul style="list-style-type: none"> • RAID 10 Mirrored and Striped—(Not supported for PCH SSD Controller Definition) RAID 10 uses mirrored pairs of disks to provide complete data redundancy and high throughput rates. • RAID 50 Striped Parity and Striped—(Not supported for PCH SSD Controller Definition) Data is striped across multiple striped parity disk sets to provide high throughput and multiple disk failure tolerance. • RAID 60 Striped Dual Parity and Striped—(Not supported for PCH SSD Controller Definition) Data is striped across multiple striped dual parity disk sets to provide high throughput and greater disk failure tolerance. <p>Note Some Cisco UCS servers require a license for certain RAID configuration options. When Cisco UCS Manager associates a service profile containing this local disk policy with a server, Cisco UCS Manager verifies that the selected RAID option is properly licensed. If there are issues, Cisco UCS Manager displays a configuration error during the service profile association.</p> <p>For RAID license information for a specific Cisco UCS server, see the Hardware Installation Guide for that server.</p>

Step 5

Click OK.

The new PCH Controller Definition appears in the navigation pane.

What to do next

For specific operating system software RAID driver installation procedures, see:

- [Installing LSI MegaSR Drivers for Windows and Linux](#) section in the [Cisco UCS C240 M4 Server Installation and Service Guide](#)

Modifying a Service Profile PCH Controller Definition

Before you begin

If you want to modify RAID level from **RAID 0 Striped** or **RAID 1 Mirrored** to **NO RAID**, then perform the following steps before starting the procedure:

1. Ensure that you have a scrub policy in the associated service profile. Refer *Creating a Service Profile with the Expert Wizard* in *Cisco UCS Manager Server Management Guide*.
2. Disassociate the server from the service profile. Refer *Disassociating a Service Profile from a Server or Server Pool* in *Cisco UCS Manager Server Management Guide*.

Procedure

- Step 1** In the **Navigation** pane, click the **Storage** tab.
- Step 2** Expand **Storage Profiles** to the specific storage profile name that you want.
- Step 3** Expand **Controller Definitions** and click the specific controller definition that you want.
- Step 4** On the **General** tab, modify the following information:

Name	Description
Name field	<p>The name of the storage controller.</p> <p>Note Once you save a PCH Controller Definition, you cannot modify the name from the General Tab Properties area.</p> <p>Enter up to 16 characters. You can use any alphanumeric characters. Special characters and spaces are not supported.</p>
Protect Configuration check box	<p>If checked, the storage profile retains the configuration even if the storage profile is disassociated from the service profile.</p> <p>Note If you disassociate the storage profile from a service profile with this option enabled, and then associate it with a new service profile that includes a local disk configuration policy with different properties, the server returns a configuration mismatch error and the association fails.</p>

Name	Description
RAID Level drop-down list	

Name	Description
	<p>This can be one of the following disk policy modes:</p> <ul style="list-style-type: none"> • Disable Local Storage—(Supported for PCH SSD Controller Definition) This disk policy mode is to disable the SATA AHCI Controller. This mode can be set only when disks are not present under the SATA AHCI controller. To re-enable this controller and to bring the controller back to its default value (AHCI), you can select No RAID or No Local Storage mode. • No Local Storage—(Supported for PCH SSD Controller Definition) For a diskless server or a SAN only configuration. If you select this option, you cannot associate any service profile which uses this policy with a server that has a local disk. • RAID 0 Striped—(Supported for PCH SSD Controller Definition) Data is striped across all disks in the array, providing fast throughput. There is no data redundancy, and all data is lost if any disk fails. • RAID 1 Mirrored—(Supported for PCH SSD Controller Definition) Data is written to two disks, providing complete data redundancy if one disk fails. The maximum array size is equal to the available space on the smaller of the two drives. • Any Configuration—(Supported for PCH SSD Controller Definition) For a server configuration that carries forward the local disk configuration without any changes. • No RAID—(Supported for PCH SSD Controller Definition) All the disks can be used individually without interdependency similar to JBOD disks. If you choose No RAID and you apply this policy to a server that already has an operating system with RAID storage configured, the system does not remove the disk contents. Therefore, there may be no visible differences on the server after you apply the No RAID mode. This can lead to a mismatch between the RAID configuration in the policy and the actual disk configuration shown in the Inventory > Storage tab for the server. To make sure that any previous RAID configuration information is removed from a disk, apply a scrub policy that removes all disk information after you apply the No RAID configuration mode. • RAID 5 Striped Parity—(Not supported for PCH SSD Controller Definition) Data is striped across all disks in the array. Part of the capacity of each disk stores parity information that can be used to reconstruct data if a disk fails. RAID 5 provides good data throughput for applications with high read request rates. • RAID 6 Striped Dual Parity—(Not supported for PCH SSD Controller Definition) Data is striped across all disks in the array and two parity disks are used to provide protection against the failure of up to two physical disks. In each row of data blocks, two sets of parity data are stored.

Name	Description
	<ul style="list-style-type: none"> • RAID 10 Mirrored and Striped—(Not supported for PCH SSD Controller Definition) RAID 10 uses mirrored pairs of disks to provide complete data redundancy and high throughput rates. • RAID 50 Striped Parity and Striped—(Not supported for PCH SSD Controller Definition) Data is striped across multiple striped parity disk sets to provide high throughput and multiple disk failure tolerance. • RAID 60 Striped Dual Parity and Striped—(Not supported for PCH SSD Controller Definition) Data is striped across multiple striped dual parity disk sets to provide high throughput and greater disk failure tolerance. <p>Note Some Cisco UCS servers require a license for certain RAID configuration options. When Cisco UCS Manager associates a service profile containing this local disk policy with a server, Cisco UCS Manager verifies that the selected RAID option is properly licensed. If there are issues, Cisco UCS Manager displays a configuration error during the service profile association.</p> <p>For RAID license information for a specific Cisco UCS server, see the Hardware Installation Guide for that server.</p>

Step 5 Click OK.

The system displays whether it saved the modified PCH Controller Definition successfully.

What to do next

If you had disassociated the server from the service profile to modify RAID level from **RAID 0 Striped** or **RAID 1 Mirrored** to **NO RAID**, then perform the following steps:

1. For any M4 server, re-acknowledge the server. Refer *Reacknowledging a Rack-Mount Server* in *Cisco UCS Manager Server Management Guide*.
2. Associate the service profile to the server. Refer *Associating a Service Profile with a Server or Server Pool* in *Cisco UCS Manager Server Management Guide*.

Deleting a Storage Profile PCH Controller Definition

Procedure

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

Step 2 Expand **Storage Profiles**.

- Step 3** Expand **PCH Controller Definitions**.
- Step 4** In the **Navigation** pane, click the specific **Controller Definition** that you want to delete.
- Step 5** In the **General** tab **Actions** area, click **Delete**.
- Step 6** Confirm whether you want to delete the definition.
- The system displays whether it deleted the definition successfully. If not, see [PCH Controller Definition Configuration Troubleshooting](#), on page 29
- Step 7** If successfully deleted, click OK.
-

PCH Controller Definition Configuration Troubleshooting

PCH Controller Definition Creation

Unsuccessful PCH Controller Definition configuration exists under the following situations:

- You try to configure a Controller definition for an unsupported server model
- You try to use the legacy local disk configuration policy and also configures the PCH storage in storage profile
- You try to configure same controller using storage profile controller definition and also by using storage profile Local LUN configuration interface
- If the **Protect Configuration** checkbox is ON and you configured the RAID Type differently than the deployed configuration in SWRAID mode.
- If the **Protect Configuration** checkbox is ON and the RAID Type does not match the present controller mode.



Warning

Any configuration change in the PCH storage configuration (like Controller mode change, RAID level change or controller qualifier change) for an already associated server triggers a PNUOS boot to happen causing a down time for the host OS.

Boot Policy

A configuration error occurs for any of the following cases:

- You select PCH Disk in boot policy but the primary or secondary target path slot number did not match with any of the inventoried internal SSD slot numbers.
- You select both PCH LUN and PCH Disk at the same time in the boot policy.

Firmware

For an incompatible software combination, there will not be any configuration error at the time of association. However, the storage configuration for the PCH SSD controller might fail or might not be deployed during association if you do not use the supported software combinations. Also, booting from the PCH SSD controller internal SSD might fail at the end of association for an incompatible software combination.

Replacing a Faulty M.2 Disk

Perform this procedure to replace a faulty M.2 disk.

Before you begin

Ensure that the SWRAID controller definition is configured and the replacement disk formatted empty drive.

Procedure

- Step 1** Gracefully power down the server.
 - Step 2** Physically remove the faulty M.2 drive. Use the **Serial Number** and **Disk Slot** to identify the faulty disk.
 - Step 3** Insert the replacement M.2 drive.
 - Step 4** Power on the server.
 - Step 5** Wait for the disk to rebuild and then re-acknowledge the server.
- Note** SWRAID rebuild may take anywhere between 35 to 75 minutes depending on the disk size, disk speed, OS content, and other parameters.
- AHCI is a NORAIID configuration and hence rebuild is not applicable.
- Note** After replacing the faulty M.2 drive, the operability state and drive-state of the drive in other slot change to Degraded and Rebuilding. To bring back the drive to normal state, decommission and recommit the blade.
-

Associating a Storage Profile with an Existing Service Profile

You can associate a storage profile with an existing service profile or a new service profile. See [Creating a Service Profile with the Expert Wizard](#).

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles**.
- Step 3** Expand the node for the organization that contains the service profile that you want to associate with a storage profile.
- Step 4** Choose the service profile that you want to associate with a storage profile.
- Step 5** In the **Work** pane, click the **Storage** tab.
- Step 6** Click the **LUN Configuration** subtab.
- Step 7** In the **Actions** area, click **Modify Storage Profile**. The **Modify Storage Profile** dialog box appears.
- Step 8** Click the **Storage Profile Policy** tab.
- Step 9** To associate an existing storage profile with this service profile, select the storage profile that you want to associate from the **Storage Profile** drop-down list, and click **OK**. The details of the storage profile appear in the **Storage Items** area.

- Step 10** To create a new storage profile and associate it with this service profile, click **Create Storage Profile**, complete the required fields, and click **OK**. [Creating a Storage Profile, on page 12](#) provides more information on creating a new storage profile.
- Step 11** (Optional) To dissociate the service profile from a storage profile, select **No Storage Profile** from the **Storage Profile** drop-down list, and click **OK**.
-

Configuring Storage Profiles

Importing Foreign Configurations for a RAID Controller on a Blade Server

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment > Chassis > Chassis Number > Servers**.
- Step 3** Choose the server of the RAID controller for which you want to import foreign configurations.
- Step 4** In the **Work** pane, click the **Inventory** tab and then the **Storage** subtab.
- Step 5** Click the **Controller** subtab.
- Step 6** In the **Actions** area, click **Import Foreign Configuration**.
-

Importing Foreign Configurations for a RAID Controller on a Rack Server

Configuring Local Disk Operations on a Blade Server

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment > Chassis > Chassis Number > Servers**.
- Step 3** Choose the server for which you want to configure local disk operations.
- Step 4** In the **Work** pane, click the **Inventory** tab and then the **Storage** subtab.
- Step 5** Click the **Disks** subtab.
- Step 6** Right-click the disk that you want and select one of the following operations:
- **Clear Foreign Configuration State**—Clears any foreign configuration that exists in a local disk when it is introduced into a new configuration.
 - **Set Unconfigured Good**—Specifies that the local disk can be configured.
 - **Set Prepare For Removal**—Specifies that the local disk is marked for removal from the chassis.
 - **Set Undo Prepare For Removal**—Specifies that the local disk is no longer marked for removal from the chassis.

- **Mark as Dedicated Hot Spare**—Specifies the local disk as a dedicated hot spare. You can select the virtual drive from the available drives.
 - **Remove Hot Spare**—Specifies that the local disk is no longer a hot spare.
 - **Set JBOD to Unconfigured Good**—Specifies that the new local disk can be configured after being marked as Unconfigured Good.
-

Configuring Local Disk Operations on a Rack Server

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
 - Step 2** Expand **Equipment** > **Rack Mounts** > **Servers**.
 - Step 3** Choose the server for which you want to configure local disk operations.
 - Step 4** In the **Work** pane, click the **Inventory** tab and then the **Storage** subtab.
 - Step 5** Click the **Disks** subtab.
 - Step 6** Right-click the disk that you want and select one of the following operations:
 - **Clear Foreign Configuration State**—Clears any foreign configuration that exists in a local disk when it is introduced into a new configuration.
 - **Set Unconfigured Good**—Specifies that the local disk can be configured.
 - **Set Prepare For Removal**—Specifies that the local disk is marked for removal.
 - **Set Undo Prepare For Removal**—Specifies that the local disk is no longer marked for removal.
 - **Mark as Dedicated Hot Spare**—Specifies the local disk as a dedicated hot spare. You can select the virtual drive from the available drives.
 - **Remove Hot Spare**—Specifies that the local disk is no longer a hot spare.
 - **Set JBOD to Unconfigured Good**—Specifies that the new local disk can be configured after being marked as **Unconfigured Good**.
-

Configuring Local Disk Operations

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Chassis** > **Chassis Number**
- Step 3** In the **Work** pane, click the **Storage** tab.
- Step 4** Click the **Disks** subtab.
- Step 5** Right-click the disk that you want and select one of the following operations:
 - **Clear Foreign Configuration State**—Clears any foreign configuration that exists in a local disk when it is introduced into a new configuration.
 - **Set Unconfigured Good**—Specifies that the local disk can be configured.

- **Set Prepare For Removal**—Specifies that the local disk is marked for removal from the chassis.
 - **Set Undo Prepare For Removal**—Specifies that the local disk is no longer marked for removal from the chassis.
 - **Mark as Dedicated Hot Spare**—Specifies the local disk as a dedicated hot spare. You can select the virtual drive from the available drives.
 - **Remove Hot Spare**—Specifies that the local disk is no longer a hot spare.
 - **Set JBOD to Unconfigured Good**—Specifies that the new local disk can be configured after being marked as Unconfigured Good.
-

Deleting an Orphan Virtual Drive

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
 - Step 2** Expand **Equipment** > **Chassis** > **Chassis Number**
 - Step 3** In the **Work** pane, click the **Storage** tab.
 - Step 4** Click the **LUNs** subtab.
 - Step 5** Right-click the virtual drive that you want and select **Delete Orphaned LUN**.
A confirmation dialog box appears.
 - Step 6** Click **Yes**.
-

Deleting an Orphan Virtual Drive on a Rack Server

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
 - Step 2** Expand **Equipment** > **Rack Mounts** > **Servers**.
 - Step 3** Choose the server for which you want to delete an orphan virtual drive.
 - Step 4** In the **Work** pane, click the **Inventory** tab and then the **Storage** subtab.
 - Step 5** Click the **LUNs** subtab.
 - Step 6** Right-click the virtual drive that you want and select **Delete Orphaned LUN**.
A confirmation dialog box appears.
 - Step 7** Click **Yes**.
-

Renaming an Orphan Virtual Drive on a Blade Server

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
 - Step 2** Expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.
 - Step 3** Choose the server for which you want to rename an orphan virtual drive.
 - Step 4** In the **Work** pane, click the **Inventory** tab and then the **Storage** subtab.
 - Step 5** Click the **LUNs** subtab.
 - Step 6** Right-click the virtual drive that you want and select **Rename Referenced LUN**.
 - Step 7** In the **Rename Referenced LUN** dialog box that appears, enter the new **LUN Name**.
 - Step 8** Click **OK**.
-

Renaming an Orphan Virtual Drive on a Rack Server

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
 - Step 2** Expand **Equipment** > **Rack Mounts** > **Servers**.
 - Step 3** Choose the server for which you want to rename an orphan virtual drive.
 - Step 4** In the **Work** pane, click the **Inventory** tab and then the **Storage** subtab.
 - Step 5** Click the **LUNs** subtab.
 - Step 6** Right-click the virtual drive that you want and select **Rename Referenced LUN**.
 - Step 7** In the **Rename Referenced LUN** dialog box that appears, enter the new **LUN Name**.
 - Step 8** Click **OK**.
-

Boot Policy for Local Storage

You can specify the primary boot device for a storage controller as a local LUN or a JBOD disk. Each storage controller can have one primary boot device. However, in a storage profile, you can set only one device as the primary boot LUN.



Note Cisco UCS S3260 M3 server nodes do not support **Any** as an option for the **Local LUN Image Path** while adding a local LUN to the boot policy through the Cisco UCS Manager GUI. In the Cisco UCS Manager CLI, Cisco UCS S3260 M3 server nodes do not support the **local-any** command option.

Configuring the Boot Policy for an Embedded Local LUN

**Note**

- For a Cisco UCS S3260 M3 server node, before you migrate the boot policy from a local LUN or JBOD to an embedded local LUN or disk, you must delete the local storage configuration, wait for the association to complete and finally add the new local storage configuration. This will enable the boot operation from the PCH disks or LUNs.
- Specify one bootable LUN as either primary or secondary boot device. If you specify the bootable LUN as both primary and secondary boot devices, the boot policy will result in the service profile configuration error.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the root node.
- Step 4** Select the boot policy that you want to configure.
- Step 5** In the **Work** pane, click the **General** tab.
- Step 6** Click the down arrows to expand the **Local Devices** area.
- Step 7** Click **Add Embedded Local LUN** to configure the boot order of the local LUN.
- Step 8** To configure the local LUN as the primary boot device, select **Primary**.
- Step 9** In the **LUN Name** field, enter the name of the LUN to be configured as the primary boot device.
- Step 10** Click **OK**.

Configuring the Boot Policy for an Embedded Local Disk

**Note**

For UCSC-C125 server, if there is no separate PCIe storage controller, then do not configure boot policy for embedded local disk. Instead, use **Add Local Disk** option.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Policies**.
- Step 3** Expand the node for the organization where you want to create the policy.
If the system does not include multitenancy, expand the root node.
- Step 4** Select the boot policy that you want to configure.

- Step 5** In the **Work** pane, click the **General** tab.
- Step 6** Click the down arrows to expand the **Local Devices** area.
- Step 7** Click **Add Embedded Local Disk** to configure the local JBOD device as the primary boot device.

BOD is supported only on the following servers:

- Cisco UCS B200 M3 blade server
- Cisco UCS B260 M4 blade server
- Cisco UCS B460 M4 blade server
- Cisco UCS B200 M4 blade server
- Cisco UCS C220 M4 rack-mount server
- Cisco UCS C240 M4 rack-mount server
- Cisco UCS C460 M4 rack-mount server
- All Cisco UCS M5 servers
- Cisco UCS S3260 M4 and M5 servers

- Step 8** In the **Disk Slot Number** field, enter the slot number of the JBOD disk to be configured as the primary boot device.
- Step 9** Click **OK**.
-

Local LUN Operations in a Service Profile

Preprovisioning a LUN Name

Preprovisioning a LUN name can be done only when the admin state of the LUN is **Undeployed**. If this LUN name exists and the LUN is orphaned, it is claimed by the service profile. If this LUN does not exist, a new LUN is created with the specified name.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
- Step 2** Expand **Servers > Service Profiles > Service_Profile_Name**.
- Step 3** In the **Work** pane, click the **Storage** tab.
- Step 4** Click the **LUN Configuration** tab.
- Step 5** In the **Local LUNs** subtab, right-click the LUN for which you want to preprovision a LUN name and select **Pre-Provision LUN Name**.
- Step 6** In the **Set Pre-Provision LUN Name** dialog box, enter the LUN name.
- Step 7** Click **OK**.
-

Claiming an Orphan LUN

Claiming an orphan LUN can be done only when the admin state of the LUN is **Undeployed**. You can explicitly change the admin state of the LUN to **Undeployed** for claiming an orphan LUN.

If the LUN name is empty, set a LUN name before claiming it.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand **Servers > Service Profiles > Service_Profile_Name**.
 - Step 3** In the **Work** pane, click the **Storage** tab.
 - Step 4** Click the **LUN Configuration** tab.
 - Step 5** In the **Local LUNs** subtab, right-click the LUN that you want to claim and select **Claim Orphan LUN**.
 - Step 6** In the **Claim Orphan LUN** dialog box that appears, select an orphaned LUN.
 - Step 7** Right-click the LUN and select **Set Admin State**.
 - Step 8** In the **Set Admin State** dialog box that appears, select **Undeployed** to undeploy a LUN and claim ownership.
 - Step 9** Click **OK**.
-

Deploying and Undeploying a LUN

You can deploy or undeploy a LUN. If the admin state of a local LUN is **Undeployed**, the reference of that LUN is removed and the LUN is not deployed.

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.
 - Step 2** Expand **Servers > Service Profiles > Service_Profile_Name**.
 - Step 3** In the **Work** pane, click the **Storage** tab.
 - Step 4** Click the **LUN Configuration** tab.
 - Step 5** In the **Local LUNs** subtab, right-click the LUN that you want to deploy or undeploy and select **Set Admin State**.
 - Step 6** In the **Set Admin State** dialog box that appears, select **Online** to deploy a LUN or **Undeployed** to undeploy a LUN.
 - Step 7** Click **OK**.
-

Renaming a Service Profile Referenced LUN

Procedure

- Step 1** In the **Navigation** pane, click **Servers**.

- Step 2** Expand **Servers** > **Service Profiles** > *Service_Profile_Name*.
- Step 3** In the **Work** pane, click the **Storage** tab.
- Step 4** Click the **LUN Configuration** tab.
- Step 5** In the **Local LUNs** subtab, right-click the LUN for which you want to rename the referenced LUN, and select **Rename Referenced LUN**.
- Step 6** In the **Rename Referenced LUN** dialog box that appears, enter the new name of the referenced LUN.
- Step 7** Click **OK**.
-