Configuring Storage Profiles

This part contains the following chapters:

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

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.

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.

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.

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.

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.

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:
   1. You specify the slot specifically by using a local disk reference
   2. 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 6
- Redundant Virtual Drives with No Hot Spare Drives, on page 6
- Redundant Virtual Drives with Hot Spare Drives, on page 6
- Replacing Hot Spare Drives, on page 7
- Inserting Physical Drives into Unused Slots, on page 7

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 UCSM to create a virtual drive, UCSM assigns a unique ID that can be used to reliably identify the virtual drive for further operations. UCSM 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, UCSM generates a unique name for the virtual drive.

You can rename virtual drives that are not referenced by any service profile or server.

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 deleted, the LUN state is changed to **Orphaned**.

Controller Constraints and Limitations

- For Cisco UCS C240, C220, C24, and C22 servers, the storage controller allows 24 virtual drives per server. For all other servers, the storage controller allows 16 virtual drives per 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`.

Configuring Storage Profiles

Configuring a Disk Group Policy

Configuring a disk group involves the following:

1. Setting the RAID level
2. Automatically or manually configuring disks in a disk group policy
3. Configuring virtual drive properties

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 the **Storage** tab.

Step 2 On the **Storage** tab, 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	The name of the policy 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.
Description field	A description of the policy. We recommend that you include information about where and when the policy should be used. 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).
RAID Level drop-down list	This can be one of the following: • **RAID 0 Striped** • **RAID 1 Mirrored** • **RAID 5 Striped Parity** • **RAID 6 Striped Dual Parity** • **RAID 10 Mirrored and Striped** • **RAID 50 Striped Parity and Striped** • **RAID 60 Striped Dual Parity and Striped**

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	Specifies the number of drives for the disk group. The range for drives is from 0 to 24 drivesfor Cisco UCS C240, C220, C24, and C22 servers. For all other servers, the limit is 16 drives per server.. **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.

Name	Description
Drive Type field	Drive type for the disk group. You can select: • **HDD** • **SSD** • **Unspecified** **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.
Number of Hot Spares field	Number of dedicated hot spares for the disk group. 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.
Number of Global Hot Spares field	Number of global hot spares for the disk group. The range for dedicated hot spares is from 0 to 24 hot spares. **Unspecified** is the default number of global hot spares. When you select the number of global 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.
Use Remaining Disks checkbox	Indicates whether the remaining disks in the disk group should be used or not. By default, this check box is not checked.

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

a) On the icon bar to the right of the table, click +

b) 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: • **Dedicated Hot Spare** • **Normal**

Name	Description
Span ID field	Span 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. You can also set the Span ID as **Unspecified** when spanning information is not required. **Unspecified** is the default Span ID of the local disk.

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: • **Platform Default** • **Read Write** • **Read Only** • **Blocked**
Read Policy field	Read policy for a virtual drive. This can be one of the following: • **Platform Default** • **Read Ahead** • **Normal**
Write Cache Policy field	Write-cache-policy for a virtual drive. This can be one of the following: • **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: • **Platform Default** • **Direct** • **Cached**

Name	Description
Drive Cache field	State of the drive cache. This can be one of the following: • **Platform Default** • **No Change** • **Enable** • **Disable**

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

If you try to associate to a server that does not support these properties, a configuration error will be generated.

Only the following storage controllers support these properties:

- LSI 6G MegaRAID SAS 9266-8i
- LSI 6G MegaRAID SAS 9271-8i
- LSI 6G MegaRAID 9265-8i
- LSI MegaRAID SAS 2208 ROMB
- LSI MegaRAID SAS 9361-8i

For the LSI MegaRAID SAS 2208 ROMB controller, these properties are supported only in the B420-M3 blade server. For the other controllers, these properties are supported in multiple rack servers.

Step 9 Click **OK**.

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.

Procedure

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

Step 2 On the **Storage** tab, expand **Storage** > **Storage Provisioning** > **Storage Profiles**

Step 3 Expand the node for the organization where you want to create the storage profile.
If the system does not include multitenancy, 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 **Storage Items** area, **Create Local LUNs** and add them to this storage profile.

Step 7 Click **OK**.

Creating a Specific Storage Profile

Procedure

Step 1 On the **Servers** tab, 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 multitenancy, 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 the Cisco UCS Manager GUI displays a confirmation dialog box, click **Yes**.

Deleting a Storage Profile

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click the **Storage** tab.	
Step 2	On the **Storage** tab, expand **Storage** > **Storage Provisioning** > **Storage Profiles**	
Step 3	Expand the node for the organization that contains the storage profile that you want to delete.	

	Command or Action	Purpose
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.	

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 the **Storage** tab.

Step 2 On the **Storage** tab, expand **Storage** > **Storage Provisioning** > **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
Name field	Name for the new local LUN.
Size (GB) field	Size of this LUN in GB. The size can range from 1 to 10240 GB. **Note** You do not need to specify a LUN size while claiming an orphaned LUN.
Expand To Available field	Specifies that this LUN can be expanded to use the entire available disk group. For each service profile, only one LUN can use this option.
Auto Deploy field	Whether the local LUN should be automatically deployed or not.
Select Disk Group Configuration field	The disk group configuration to be applied to this local LUN.

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

Step 8 Click **OK**.

Deleting Local LUNs

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click the **Storage** tab.	
Step 2	On the **Storage** tab, expand **Storage** > **Storage Provisioning** > **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**.	

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. Creating a Service Profile with the Expert Wizard provides more information about associating a storage profile with a new service profile.

Procedure

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

Step 2 On the **Servers** tab, 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**.

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 the **Servers** tab.

Step 2 On the **Servers** tab, 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 **LUN Configuration** 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.
- **Admin State—**Specifies whether a local LUN should be deployed or not. Admin state can be **Online** or **Undeployed**.

When the local LUN is being referenced by a service profile, if the auto-deploy status is **no-auto-deploy** then the admin state will be **Undeployed** , else it will be **Online**. After the local LUN is referenced by a service profile, any change made to this local LUN's auto-deploy status is not reflected in the admin state of the LUN inherited by the service profile.

- **RAID Level—**Summary of the RAID level of the disk group used.
- **Provisioned Size (GB)—**Size, in GB, of the LUN specified in the storage profile.
- **Assigned Size (MB)—**Size, in MB, assigned by UCSM.
- **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.
- **Referenced LUN Name—**The preprovisioned virtual drive name, or UCSM-generated virtual drive name.
- **Deploy Name—**The virtual drive name after deployment.
- **ID—**LUN ID.
- **Order—**Order of LUN visibility to the server.
- **Bootable—**Whether the LUN is bootable or not.
- **LUN New Name—**New name of the LUN.
- **Drive State—**State of the virtual drive. The states are:
  - **Unknown**
  - **Optimal**
  - **Degraded**
  - **Inoperable**
  - **Partially Degraded**

Importing Foreign Configurations for a RAID Controller on a Blade Server

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click the **Equipment** tab.	
Step 2	On the **Equipment** tab, 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

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click the **Equipment** tab.	
Step 2	On the **Equipment** tab, expand **Equipment** > **Rack Mounts** > **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**.	

Configuring Local Disk Operations on a Blade Server

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click the **Equipment** tab.	
Step 2	On the **Equipment** tab, expand **Equipment** > **Chassis** > *Chassis Number* > **Servers**.	

	Command or Action	Purpose
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

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click the **Equipment** tab.	
Step 2	On the **Equipment** tab, 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.	

	Command or Action	Purpose
	• **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 Virtual Drive Operations

The following operations can be performed only on orphaned virtual drives:

- Delete an orphaned virtual drive
- Rename an orphaned virtual drive

Deleting an Orphan Virtual Drive on a Blade Server

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click the **Equipment** tab.	
Step 2	On the **Equipment** tab, expand **Equipment** > **Chassis** > ***Chassis Number*** > **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**.	

Deleting an Orphan Virtual Drive on a Rack Server

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click the **Equipment** tab.	
Step 2	On the **Equipment** tab, 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

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click the **Equipment** tab.	
Step 2	On the **Equipment** tab, 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

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click the **Equipment** tab.	
Step 2	On the **Equipment** tab, 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.

Configuring the Boot Policy for a Local Device

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click the **Servers** tab.	
Step 2	On the **Servers** tab, 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 Local LUN** to configure the boot order of the local LUN.	

	Command or Action	Purpose
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 a Local JBod Device

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click the **Servers** tab.	
Step 2	On the **Servers** tab, 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 Local JBod** 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
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, its is claimed by the service profile. If this LUN does not exist, a new LUN is created with the specified name.

Procedure

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click the **Servers** tab.	
Step 2	On the **Servers** tab, 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

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click the **Servers** tab.	
Step 2	On the **Servers** tab, 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**.	

	Command or Action	Purpose
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

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click the **Servers** tab.	
Step 2	On the **Servers** tab, 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

	Command or Action	Purpose
Step 1	In the **Navigation** pane, click the **Servers** tab.	
Step 2	On the **Servers** tab, expand **Servers** > **Service Profiles** > ***Service_Profile_Name***.	
Step 3	In the **Work** pane, click the **Storage** tab.	

	Command or Action	Purpose
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**.	