Storage Profiles

This chapter includes the following sections:

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.

Note Storage profiles on Cisco UCS rack and blade servers are supported on Cisco UCS Manager release 2.2.7 and above, and Cisco UCS Manager release 3.1.1 and above.

Because Cisco UCS M-series Modular Servers have been deprecated, storage profiles with boot orders created in Cisco UCS Central release 1.4 are not supported in Cisco UCS Central release 1.5 and later.

Storage profiles allow you to do the following:

- Configure multiple virtual drives and select the physical drives that are used by a virtual drive.
- 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.

Note LUN resizing is not supported.

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.
- 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 cache policy of Write Back Good BBU mode, but the BBU has failed, or there is no BBU.

  Note This state does not occur if you select 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.

Virtual Drive Naming

When you use Cisco UCS Central to create a virtual drive, Cisco UCS Central assigns a unique ID that can be used to reliably identify the virtual drive for further operations. Cisco UCS Central 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 Central generates a unique name for the virtual drive.

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

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. A minimum of one disk is required for RAID 0.

- 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. A minimum of two disks are required for RAID 1.
- 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. A minimum of six disks are required for RAID 50.
- 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. A minimum of eight disks are required for RAID 60.

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.

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. When decommissioning the server, the state of all the LUNs associated with the server is changed to Not in use or Orphaned. However, no action is taken to delete the actual LUNs.

Note When LUNs are orphaned, the LUNs stay in the shared storage and the content is preserved. You can reclaim the orphan LUN to retrieve the data and attach the LUN to a new service profile.

Creating or Editing a Storage Profile

Step 1 In the **Actions** bar, type **Create Storage Profile** and press Enter.

Step 2 In the **Storage Profile** dialog box, click **Basic**, then select the **Organization** in which you want to create the storage profile.

Step 3 Enter the **Name** and optional **Description**.
The name is case sensitive.

Step 4 Select the server type where you plan to apply the storage profile.

Step 5 In **Local LUNs**, click **Add** to add a new local LUN.

Step 6 Click **Basic** to add a new LUN, or **Claim Mode** to reclaim a previously orphaned LUN.
Creating a local LUN in Claim Mode will claim an orphaned LUN when the storage profile is applied to an associated service profile.

Step 7 In the **Basic** tab, do the following:

a) Enter the size in GB.
The size must be between 1 GB and 10240 GB.

b) Choose whether to enable automatic deployment for the local LUN.

c) Choose whether this LUN can be expanded to use the entire available disk group. For each service profile, only one LUN can use this option.

Step 8 In the **Disk Group** tab, select the **Disk Group Configuration Policy** that you want to apply. If you want to encrypt a disk, select the Disk Group with Security enabled.

Step 9 In **Controller Defs**, click **Add**.

Step 10 Enable configuration protection in order to prevent a service profile using this local disk policy from being associated to a server with a different physical disk configuration.
If the service profile includes a local disk policy with configuration protection enabled, and there is an attempt to associate that service profile to a server that includes disks with a different local disk configuration, the association will immediately fail with a configuration mismatch error.

Caution We recommend that you enable configuration protection to preserve any data that may exist on local disks. If disabled, any existing volume that does not match the local disk configuration policy will be deleted.

Step 11 Set the RAID level.
See RAID Levels, on page 3 for detailed information about the different levels.

Step 12 In **Security Policy**, click the **Local** tab, and enter 32 character alphanumeric character key in the **Key** field.

Attention Use the instructions in Step 12 and later only if you want to enable security for the drives. You can choose to enable a Local or Remote security policy. If you do not want to enable security for the drives, select **None** in the **Security Policy** tab and proceed.

This key is used to encrypt the data in the disks. You can create a local security policy on a new or existing Storage profile.

Step 13 (Optional) To edit or modify the **Local Security** policy, enter the current security key for the database in **Deployed Key**.

Step 14 In **Security Policy**, click the **Remote** tab, and enter the following details:

Important Before you create a Remote Security policy, you must have created a KMIP Client Certificate policy. You can create a remote policy on a new or existing storage profile.

a) Enter the **Primary Server IP Address/Hostname**. These credentials are used for accessing the KMIP server to fetch the key.
b) (Optional) Enter the **Secondary Server IP Address/Hostname**.
c) Enter the port number of the server in **Port**.
d) Enter seconds in number in **Timeout**.
e) Enter the contents of KMIP certificate in **KMIP Server Public Certificate**.
Save this certificate from the browser in base-64 format.

a) Select **Enable** to enter **Username** and **Password** to enable SED management through manual key configuration.

Step 15 Click **Create**.

Step 16 To modify a **Remote Security** policy and make it a **Local Security** policy, select the policy you want to modify and repeat the procedure in Step 14, and follow these steps:

a) Click **Local Policy** option and enter a new security key for the controller in the **Key** field.

- The key created is associated to the storage profile and is deployed under the storage controller. To verify, navigate to **Server** > **Controller**.
- After you configure the key from the KMIP server, navigate to the **Custom Attributes** column in the **Secure Key Management Console** page and add the serial number of the controller. You can obtain the serial number of the controller from the **Server** inventory.
- Verify that the **Security** field shows **Drive Security Enabled**. For more information on Enabling/Disabling security on disks, see Creating or Editing a Disk Group Configuration Policy, on page 8.
- You can check the status of the storage configuration in **Configuration Status**.

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.

Creating or Editing a Disk Group Configuration Policy

Step 1 In the **Actions** bar, type, **Create Disk Group Configuration Policy** and press Enter.

Step 2 In the **Disk Group Configuration Policy** dialog box, click **Basic** and choose the **Organization** in which you want to create the disk group configuration policy.

Step 3 Enter the **Name** and optional **Description**.
The name is case sensitive.

Step 4 Select the **Raid Level**.
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 & Striped**
- **RAID 50 Striped Parity & Striped**
- **RAID 60 Striped Dual Parity & Striped**

Step 5 In **Disk Group**, choose one of the following:

- **Auto—**Choose the **Drive Type**, type values for the drive information, and choose whether to use the remaining disks.
- **Manual—**Add a disk slot ID, and select the Span and disk role for the slots.

You can create a policy with SEDs and enter the **Disk Slot ID** in the **Configuration** details.

Step 6 In **Virtual Drive** icon, complete the fields as necessary.

Step 7 In **Security** select **Enable** or **Disable** to automatically scan the disk for the SED capable disks and pick them.

Step 8 Click **Create**.

The **Security** field in the **Controller** details displays DriveSecurityCapable, Enabled after you enable security.

Monitoring the Health of SSDs

Cisco UCS Central lets you monitor the following SSD conditions (statistics) as reported by Cisco UCS Manager:

- Wear status (displayed in days)
- Percentage of disk life remaining
- Power cycle count

- Power on hours

Step 1 Click the **Dashboard** icon and choose **Servers**.
Step 2 Choose the server that contains the SSD.
Step 3 Click the **Controller** icon.
Step 4 Choose the controller managing the SSD.
Step 5 Click the **Launch** icon and from the drop-down list choose **Statistics**.