



Storage Profiles

This chapter includes the following sections:

- [Storage Profiles, page 1](#)
- [Disk Groups and Disk Group Configuration Policies, page 7](#)

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.

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 a Storage Profile

Procedure

	Command or Action	Purpose
Step 1	UCSC# connect policy-mgr	Enters policy manager mode.
Step 2	UCSC(policy-mgr)# scope org org-name	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 3	UCSC(policy-mgr) /org # create storage-profile profile-name	Creates the specified storage profile and enters organization storage profile mode.
Step 4	UCSC(policy-mgr) /org/storage-profile # create local-lun name	Creates a local LUN with the specified name, and enters local LUN mode.
Step 5	UCSC(policy-mgr) /org/storage-profile/local-lun # set auto-deploy	Enables automatic deployment for the local LUN. Use the set no-auto-deploy command to disable.
Step 6	UCSC(policy-mgr) /org/storage-profile/local-lun # set disk-policy-name policy-name	Specifies the disk group configuration policy that you want to use.
Step 7	UCSC(policy-mgr) /org/storage-profile/local-lun # set order {order_number not-applicable}	For Cisco UCS M-Series servers, you can select an order for your local LUNs. LUN order is not supported for Cisco UCSB-Series and C-Series servers.
Step 8	UCSC(policy-mgr) /org/storage-profile/local-lun # exit	Set the size in GB or select unspecified to expand the LUN to use the entire available disk group. For each service profile, only one LUN can use this option.
Step 9	UCSC(policy-mgr) /org/storage-profile/local-lun # exit	Returns to organization storage profile mode.
Step 10	UCSC(policy-mgr) /org/storage-profile # create controller-def name	Creates the specified controller definition and enters controller def mode.
Step 11	UCSC(policy-mgr) /org/storage-profile/controller-def # enter controller-mode-config	Enters controller configuration mode.
Step 12	UCSC(policy-mgr) /org/storage-profile/controller-def/controller-mode-config # set protect-config {yes no}	Choose whether to 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.

	Command or Action	Purpose
		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.
Step 13	UCSC(policy-mgr) /org/storage-profile/controller-def/controller-mode-config # set raid-mode {any-configuration no-local-storage no-raid raid-0-striped raid-1-mirrored raid-5-striped-parity raid-50-striped-parity-and-striped raid-6-striped-dual-parity raid-60-striped-dual-parity-and-striped raid-10-mirrored-and-striped}	Set the raid level.
Step 14	UCSC(policy-mgr) /org/storage-profile/controller-def/controller-mode-config # commit-buffer	Commits the transaction to the system configuration.

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 a Disk Group Configuration Policy

Procedure

	Command or Action	Purpose
Step 1	UCSC# connect policy-mgr	Enters policy manager mode.

	Command or Action	Purpose
Step 2	UCSC(policy-mgr)# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the <i>org-name</i> .
Step 3	UCSC(policy-mgr) /org # create disk-group-config-policy <i>policy-name</i>	Creates the specified disk group configuration policy and enters disk group configuration policy mode.
Step 4	UCSC(policy-mgr) /org/disk-group-config-policy # set raid-level { raid-0-striped raid-1-mirrored raid-10-mirrored-and-striped raid-5-striped-parity raid-50-striped-parity-and-striped raid-6-striped-dual-parity raid-60-striped-dual-parity-and-striped }	Specifies the RAID level for the disk group configuration policy.
Step 5	UCSC(policy-mgr) /org/disk-group-config-policy # create disk-group-qual	Creates a disk group qualification policy and enters disk group qualification mode.
Step 6	UCSC(policy-mgr) /org/disk-group-config-policy/disk-group-qual # set drive-type { hdd ssd unspecified }	Specifies the drive type for the disk group. Note If you specify unspecified 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 was SSD, all subsequent drives would be SSD.
Step 7	UCSC(policy-mgr) /org/disk-group-config-policy/disk-group-qual # set drive-size { <i>drive_size</i> unspecified }	Specifies the minimum drive size for the disk group. Only disks that match this criteria will be available for selection. The range for minimum drive size is from 0 to 10240 GB. If you set the minimum drive size as unspecified, drives of all sizes will be available for selection.
Step 8	UCSC(policy-mgr) /org/disk-group-config-policy/disk-group-qual # set num-ded-hot-spares { <i>number_of_spare</i> unspecified }	Specifies the number of dedicated hot spares for the disk group. The range for dedicated hot spares is from 0 to 24 hot spares. If you set the number of dedicated hot spares as unspecified, the hot spares will be selected according to the disk selection process.
Step 9	UCSC(policy-mgr) /org/disk-group-config-policy/disk-group-qual # set num-drives { <i>number_of_drives</i> unspecified }	Specifies the number of drives for the disk group. The range for drives is from 0 to 24 drives for Cisco UCS C240, C220, C24, and C22 servers. For all other servers, the limit is 16

	Command or Action	Purpose
		drives per server. If you set the number of drives as unspecified, the number of drives will be selected according to the disk selection process.
Step 10	UCSC(policy-mgr) /org/disk-group-config-policy/disk-group-qual # exit	Returns to disk group configuration policy mode.
Step 11	UCSC(policy-mgr) /org/disk-group-config-policy # create local-disk-config-ref <i>slot_number</i>	Creates a local disk configuration reference for the specified slot and enters local disk configuration reference mode.
Step 12	UCSC(policy-mgr) /org/disk-group-config-policy/local-disk-config-ref # set role { dedicated-hot-spare global-hot-spare normal }	Specifies the role of the local disk in the disk group.
Step 13	UCSC(policy-mgr) /org/disk-group-config-policy/local-disk-config-ref # set span-id { <i>span-id</i> unspecified }	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. You can also set the span ID as unspecified when spanning information is not required.
Step 14	UCSC(policy-mgr) /org/disk-group-config-policy/local-disk-config-ref # exit	Returns to disk group configuration policy mode.
Step 15	UCSC(policy-mgr) /org/disk-group-config-policy # create virtual-drive-def	Creates a virtual drive definition and enters the virtual drive definition mode.
Step 16	UCSC(policy-mgr) /org/disk-group-config-policy/virtual-drive-def # set access-policy { blocked platform-default read-only read-write }	Specifies the access policy.
Step 17	UCSC(policy-mgr) /org/disk-group-config-policy/virtual-drive-def # set drive-cache { disable enable no-change platform-default }	Specifies the state of the drive cache.
Step 18	UCSC(policy-mgr) /org/disk-group-config-policy/virtual-drive-def # set io-policy { cached direct platform-default }	Specifies the I/O policy.
Step 19	UCSC(policy-mgr) /org/disk-group-config-policy/virtual-drive-def # set read-policy { normal platform-default read-ahead }	Specifies the read policy.

	Command or Action	Purpose
Step 20	UCSC(policy-mgr) /org/disk-group-config-policy/virtual-drive-def # set stripe-size { 1042kb 128kb 16kb 256kb 32kb 512kb 64kb 8kb platform-default }	Specifies the strip size.
Step 21	UCSC(policy-mgr) /org/disk-group-config-policy/virtual-drive-def # set write-cache-policy { always-write-back platform-default write-back-good-bbu write-through }	Specifies the write-cache-policy.
Step 22	UCSC(policy-mgr) /org/disk-group-config-policy/virtual-drive-def # commit-buffer	Commits the transaction to the system configuration.