Power Capping and Power Management in Cisco UCS

Power Capping in Cisco UCS

You can control the maximum power consumption on a server through power capping, as well as manage the power allocation in the Cisco UCS Manager for blade servers, UCS C220 and C240 M4 rack servers, UCS Mini, and mixed UCS domains.

UCS Manager supports power capping on the following:

- UCS 6200 Series Fabric Interconnects

- UCS Mini 6324

- UCS 6300 Series Fabric Interconnects

You can use Policy Driven Chassis Group Power Cap, or Manual Blade Level Power Cap methods to allocate power that applies to all of the servers in a chassis.

Cisco UCS Manager provides the following power management policies to help you allocate power to your servers:

Power Management Policies	Description
Power Policy	Specifies the redundancy for power supplies in all chassis in a Cisco UCS domain.
Power Control Policies	Specifies the priority to calculate the initial power allocation for each blade in a chassis.
Global Power Allocation	Specifies the Policy Driven Chassis Group Power Cap or the Manual Blade Level Power Cap to apply to all servers in a chassis.
Global Power Profiling	Specifies how the power cap values of the servers are calculated. If it is enabled, the servers will be profiled during discovery through benchmarking. This policy applies when the Global Power Allocation Policy is set to Policy Driven Chassis Group Cap.

Power Policy Configuration

Power Policy for Cisco UCS Servers

The power policy is global and is inherited by all of the chassis' managed by the Cisco UCS Manager instance. You can add the power policy to a service profile to specify the redundancy for power supplies in all chassis' in the Cisco UCS domain. This policy is also known as the PSU policy.

For more information about power supply redundancy, see *Cisco UCS 5108 Server Chassis Hardware Installation Guide*.

Configuring the Power Policy

Procedure

Step 1	In the **Navigation** pane, click **Equipment**.
Step 2	Click the **Equipment** node.
Step 3	In the **Work** pane, click the **Policies** tab.
Step 4	Click the **Global Policies** subtab.
Step 5	In the **Power Policy** area, click one of the following radio buttons in the **Redundancy** field:

- **Non Redundant**—Cisco UCS Manager turns on the minimum number of power supplies (PSUs) needed and balances the load between them. If any additional PSUs are installed, Cisco UCS Manager sets them to a "turned-off" state. If the power to any PSU is disrupted, the system may experience an interruption in service until Cisco UCS Manager can activate a new PSU and rebalance the load.

  In general, a Cisco UCS chassis requires at least two PSUs for non-redundant operation. Only smaller configurations (requiring less than 2500W) can be powered by a single PSU.

- **N+1**—The total number of PSUs to satisfy non-redundancy, plus one additional PSU for redundancy, are turned on and equally share the power load for the chassis. If any additional PSUs are installed, Cisco UCS Manager sets them to a "turned-off" state. If the power to any PSU is disrupted, Cisco UCS Manager can recover without an interruption in service.

  In general, a Cisco UCS chassis requires at least three PSUs for N+1 operation.

- **Grid**—Two power sources are turned on, or the chassis requires greater than N+1 redundancy. If one source fails (which causes a loss of power to one or two PSUs), the surviving PSUs on the other power circuit continue to provide power to the chassis.

  For more information about power supply redundancy, see *Cisco UCS 5108 Server Chassis Hardware Installation Guide*.

Step 6	Click **Save Changes**.

Power Supply for Redundancy Method

PSU Redundancy	Max Power @ 220 V (Watts)	Max Power @ 110 V (Watts)
1+1 (N+1) OR 1 (N)	2500	1300
2+1 (N+1) OR 2 (N) or 2+2 (Grid)	5000	2600
3+1 (N+1) OR 3 (N)	5472	3900
4 (N)	5472	5200

Configuring Policy Driven Chassis Group Power Capping

Policy Driven Chassis Group Power Capping

When you select the Policy Driven Chassis Group Power Cap in the Global Cap Policy, Cisco UCS can maintain the over-subscription of servers without risking power failures. You can achieve over-subscription through a two-tier process. For example, at the chassis level, Cisco UCS divides the amount of power available among members of the power group, and at the blade level, the amount of power allotted to a chassis is divided among blades based on priority.

Each time a service profile is associated or disassociated, Cisco UCS Manager recalculates the power allotment for each blade server within the chassis. If necessary, power from lower-priority service profiles is redistributed to higher-priority service profiles.

UCS power groups cap power in less than one second to safely protect data center circuit breakers. A blade must stay at its cap for 20 seconds before the chassis power distribution is optimized. This is intentionally carried out over a slower timescale to prevent reacting to transient spikes in demand.

Note   The system reserves enough power to boot a server in each slot, even if that slot is empty. This reserved power cannot be leveraged by servers requiring more power. Blades that fail to comply with the power cap are penalized.

Power Control Policy

Cisco UCS uses the priority set in the power control policy along with the blade type and configuration to calculate the initial power allocation for each blade within a chassis. During normal operation, the active blades within a chassis can borrow power from idle blades within the same chassis. If all blades are active and reach the power cap, service profiles with higher priority power control policies take precedence over service profiles with lower priority power control policies.

Priority is ranked on a scale of 1-10, where 1 indicates the highest priority and 10 indicates lowest priority. The default priority is 5.

Note   You must include the power control policy in a service profile and that service profile must be associated with a server for it to take effect.

Creating a Power Control Policy

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   Right-click **Power Control Policies** and choose **Create Power Control Policy**.

Step 5   In the **Create Power Control Policy** dialog box, complete the following fields:

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 is saved.
Description field	A description of the policy. Cisco recommends including information about where and when to use the policy. 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).
Fan Speed Policy drop-down	Fan speed is for rack servers only. This can be one of the following: • **Low Power**—The fan runs at the minimum speed required to keep the server cool. • **Balanced**—The fan runs faster when needed based on the heat generated by the server. When possible, the fan returns to the minimum required speed. • **Performance**—The fan is kept at the speed needed for better server performance. This draws more power but means the fan is already at speed if the server begins to heat up. • **High Power**—The fan is kept at an even higher speed that emphasizes performance over power consumption. • **Max Power**—The fan is kept at the maximum speed at all times. This option provides the most cooling and uses the most power. • **Any**—The server determines the optimal fan speed.

Name	Description
Power Capping field	What happens to a server when the demand for power within a power group exceeds the power supply. This can be one of the following: • **No Cap**—The server runs at full capacity regardless of the power requirements of the other servers in its power group. • **cap**—The server is allocated a minimum amount of power capacity based on the the server's priority relative to the other servers in its server group. If more power becomes available, Cisco UCS allows the capped servers to exceed their original allocations. It only lowers the allocations if there is a drop in the total power available to the power group. When you select **cap**, Cisco UCS Manager GUI displays the **Priority** field.
Priority field	The priority the server has within its power group when power capping is in effect. Enter an integer between 1 and 10, where 1 is the highest priority.

Step 6   Click **OK**.

What to do next

Include the policy in a service profile or service profile template.

Deleting a Power Control Policy

Procedure

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

Step 2   Expand **Servers** > **Policies** > *Organization_Name*.

Step 3   Expand the **Power Control Policies** node.

Step 4   Right-click the policy you want to delete and select **Delete**.

Step 5   If a confirmation dialog box displays, click **Yes**.

Power Groups in UCS Manager

A power group is a set of chassis that all draw power from the same power distribution unit (PDU). In Cisco UCS Manager, you can create power groups that include one or more chassis, then set a peak power cap in AC watts for that power grouping.

Implementing power capping at the chassis level requires the following:

- IOM, CIMC, and BIOS version 1.4 or higher

- Two Power Supply Units (PSUs)

The peak power cap is a static value that represents the maximum power available to all blade servers within a given power group. If you add or remove a blade from a power group, but do not manually modify the peak power value, the power group adjusts the peak power cap to accommodate the basic power-on requirements of all blades within that power group.

A minimum of 890 AC watts should be set for each chassis. This converts to 800 watts of DC power, which is the minimum amount of power required to power an empty chassis. To associate a half-width blade, the group cap needs to be set to 1475 AC watts. For a full-width blade, it needs to be set to 2060 AC watts.

After a chassis is added to a power group, all service profile associated with the blades in the chassis become part of that power group. Similarly, if you add a new blade to a chassis, that blade inherently becomes part of the chassis' power group.

Note   Creating a power group is not the same as creating a server pool. However, you can populate a server pool with members of the same power group by creating a power qualifier and adding it to server pool policy.

When a chassis is removed or deleted, the chassis gets removed from the power group.

UCS Manager supports explicit and implicit power groups.

- **Explicit**: You can create a power group, add chassis' and racks, and assign a budget for the group.
- **Implicit**: Ensures that the chassis is always protected by limiting the power consumption within safe limits. By default, all chassis that are not part of an explicit power group are assigned to the default group and the appropriate caps are placed. New chassis that connect to UCS Manager are added to the default power group until you move them to a different power group.

The following table describes the error messages you might encounter while assigning power budget and working with power groups.

Error Message	Cause	Recommended Action

`Insufficient budget for power group POWERGROUP_NAME`
 	
and/or
 	
`Chassis N cannot be capped as group cap is low. Please consider raising the cap.`
 	
and/or
 	
`Admin committed insufficient for power group GROUP_NAME, using previous value N`
 	
and/or
 	
`Power cap application failed for chassis N`
	One of these messages displays if you did not meet the minimum limit when assigning the power cap for a chassis, or the power requirement increased because of the addition of blades or change of power policies.	Increase the power cap limit to the **Minimum Power Cap for Allowing Operations (W)** value displayed on the **Power Group** page for the specified power group.

`Chassis N cannot be capped as the available PSU power is not enough for the chassis and the blades. Please correct the problem by checking input power or replace the PSU`	Displays when the power budget requirement for the chassis is more than the PSU power that is available.	Check the PSU input power and redundancy policy to ensure that enough power is available for the chassis. If a PSU failed, replace the PSU.
`Power cap application failed for server N`	Displays when the server is consuming more power than allocated and cannot be capped, or the server is powered on when no power is allocated.	Do not power on un-associated servers.
`P-State lowered as consumption hit power cap for server`	Displays when the server is capped to reduce the power consumption below the allocated power.	This is an information message. If a server should not be capped, in the service profile set the value of the power control policy **Power Capping** field to **no-cap**.
`Chassis N has a mix of high-line and low-line PSU input power sources.`	This fault is raised when a chassis has a mix of high-line and low-line PSU input sources connected.	This is an unsupported configuration. All PSUs must be connected to similar power sources.

Creating a Power Group

Before you begin

Make sure that the global power allocation policy is set to **Policy Driven Chassis Group Cap** on the **Global Policies** tab.

Procedure

Step 1	In the **Navigation** pane, click **Equipment**.
Step 2	Click the **Equipment** node.
Step 3	In the **Work** pane, click the **Policies** tab.
Step 4	Click the **Power Groups** subtab.
Step 5	On the icon bar to the right of the table, click +.

If the + icon is disabled, click an entry in the table to enable it.

Step 6 On the first page of the **Create Power Group** wizard, complete the following fields:

a) Enter a unique name and description for the power group.

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 is saved.

b) Click **Next**.

Step 7 On the **Add Chassis Members** page of the **Create Power Group** wizard, do the following:

a) In the **Chassis** table, choose one or more chassis to include in the power group.

b) Click the >> button to add the chassis to the **Selected Chassis** table that displays all chassis included in the power group.

You can use the << button to remove one or more chassis from the power group.

c) Click **Next**.

Step 8 On the **Add Rack Members** page of the **Create Power Group** wizard, do the following:

a) In the **Rack Unit** table, choose one or more rack units to include in the power group.

b) Click the >> button to add the rack to the **Selected Rack Unit** table that displays all racks included in the power group.

You can use the << button to remove one or more rack units from the power group.

c) Click **Next**.

Step 9 On the **Add FEX Members** page of the **Create Power Group** wizard, do the following:

a) In the **FEX** table, choose one or more FEX to include in the power group.

b) Click the >> button to add the chassis to the **Selected FEX** table that displays all FEX included in the power group.

You can use the << button to remove one or more FEX from the power group.

c) Click **Next**.

Step 10   On the **Add FI Members** page of the **Create Power Group** wizard, do the following:

a) In the **FI** table, choose one or more FI to include in the power group.

b) Click the **>>** button to add the FI to the **Selected FI** table that displays all chassis included in the power group.

You can use the **<<** button to remove one or more FI from the power group.

c) Click **Next**.

Step 11   On the **Power Group Attributes** page of the **Create Power Group** wizard, do the following:

a) Complete the following fields:

Name	Description
Power Cap field	The maximum peak power (in watts) available to the power group. Enter an integer between 0 and 10000000.
Enable Dynamic Reallocation field	This can be one of the following: • **Chassis**—Cisco UCS monitors power usage and changes the blade allocations as required to maximize power utilization. • **None**—Blade allocations are not adjusted dynamically.

b) Click **Finish**.

Adding a Chassis to a Power Group

Procedure

Step 1   In the **Navigation** pane, click **Equipment**.

Step 2   Click the **Equipment** node.

Step 3   In the **Work** pane, click the **Power Groups** tab.

Step 4   Right-click the power group to which you want to add a chassis and choose **Add Chassis Members**.

Step 5   In the **Add Chassis Members** dialog box, do the following:

a) In the **Chassis** table, choose one or more chassis to include in the power group.

b) Click the **>>** button to add the chassis to the **Selected Chassis** table that displays all chassis included in the power group.

You can use the **<<** button to remove one or more chassis from the power group.

c) Click **OK**.

Removing a Chassis from a Power Group

Procedure

Step 1	In the **Navigation** pane, click **Equipment**.
Step 2	Click the **Equipment** node.
Step 3	In the **Work** pane, click the **Power Groups** tab.
Step 4	Expand the power group from which you want to remove a chassis.
Step 5	Right-click the chassis that you want to remove from the power group and choose **Delete**.
Step 6	If a confirmation dialog box displays, click **Yes**.

Deleting a Power Group

Procedure

Step 1	In the **Navigation** pane, click **Equipment**.
Step 2	Click the **Equipment** node.
Step 3	In the **Work** pane, click the **Power Groups** tab.
Step 4	Right-click the power group that you want to delete and choose **Delete**.
Step 5	If a confirmation dialog box displays, click **Yes**.

Blade Level Power Capping

Manual Blade Level Power Cap

When manual blade-level power cap is configured in the global cap policy, you can set a power cap for each blade server in a Cisco UCS domain.

The following configuration options are available:

- **Watts**—You can specify the maximum amount of power that the server can consume at one time. This maximum can be any amount between 0 watts and 1100 watts.

- **Unbounded**—No power usage limitations are imposed on the server. The server can use as much power as it requires.

If the server encounters a spike in power usage that meets or exceeds the maximum configured for the server, Cisco UCS Manager does not disconnect or shut down the server. Instead, Cisco UCS Manager reduces the power that is made available to the server. This reduction can slow down the server, including a reduction in CPU speed.

Note If you configure the manual blade-level power cap using **Equipment** > **Policies** > **Global Policies** > **Global Power Allocation Policy**, the priority set in the Power Control Policy is no longer relevant.

Setting the Blade-Level Power Cap for a Server

Before you begin

Make sure the global power allocation policy is set to **Manual Blade Level Cap** on the **Global Policies** tab.

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 set the power budget.

Step 4   In the **Work** pane, click the **General** tab.

Step 5   In the **Power Budget** area, do the following:

a) Click the **Expand** icon to the right of the heading to display the fields.

b) Complete the following fields:

Name	Description
Admin Status field	Whether this server is power capped. This can be one of the following: • **Unbounded**—The server is not power capped under any circumstances. • **Enabled**—The Cisco UCS Manager GUI displays the **Watts** field. **Note** Power capping goes into effect only if there is insufficient power available to the chassis to meet the demand. If there is sufficient power, the server can use as many watts as it requires.
Watts field	The maximum number of watts that the server can use if there is not enough power to the chassis to meet the demand. The value range is from 0 and 10000000.

Step 6   Click **Save Changes**.

Viewing the Blade-Level Power Cap

Procedure

Step 1   In the **Navigation** pane, click **Equipment**.

Step 2   Expand **Equipment** > **Chassis**.

Step 3   Choose the chassis for which you want to view the server power usage.

Step 4   Do one of the following:

- To view the power usage for all servers in the chassis, click the **Power** tab in the **Work** pane.
- To view the power usage for one server in the chassis, expand the chassis and click the server. Then click the **Power** tab in the **Work** pane.

Step 5   If necessary, expand the **Motherboards** node to view the power counters.

Global Power Profiling Policy Configuration

Global Power Profiling Policy

The Global Power Profiling Policy specifies how power allocation is applied to all of the servers in a chassis. The policy applies when you set the Global Power Allocation Policy to **Policy Driven Chassis Group Cap**. You can set the Global Power Profiling Policy to one of the following:

- **Disabled**—The minimum and maximum power cap values of the blades are calculated based on the static power consumption values of each of the components.
- **Enabled**—The minimum and maximum power cap values of the blades are measured as part of the server discovery. These values are similar to the actual power consumption of the blades.

Note   After enabling the Global Power Profiling Policy, you must re-acknowledge the blades to obtain the minimum and maximum power cap.

Important   Power profiling is not supported in Cisco UCS B460 M4 blades.

Configuring the Global Power Profile Policy

Procedure

Step 1	In the **Navigation** pane, click **Equipment**.
Step 2	Click the **Equipment** node.
Step 3	In the **Work** pane, click the **Policies** tab.
Step 4	Click the **Global Policies** subtab.
Step 5	In the **Global Power Profiling Policy** area, check the **Profile Power** checkbox to enable the Global Power Profiling Policy.
Step 6	Click **Save Changes**.

Global Power Allocation Policy Configuration

Global Power Allocation Policy

The Global Power Allocation Policy allows you to specify the Policy Driven Chassis Group Power Cap or Manual Blade-level Power Cap power allocation method applied to servers in a chassis.

Cisco recommends using the default Policy Driven Chassis Group Power Cap power allocation method.

> ☞
>
> **Important** Any change to the Manual Blade level Power Cap configuration results in the loss of any groups or configuration options set for the Policy Driven Chassis Group Power Cap.

Configuring the Global Power Allocation Policy

Procedure

Step 1	In the **Navigation** pane, click **Equipment**.
Step 2	Click the **Equipment** node.
Step 3	In the **Work** pane, click the **Policies** tab.
Step 4	Click the **Global Policies** subtab.
Step 5	In the **Global Power Allocation Policy** area, click one of the following radio buttons in the **Allocation Method** field to determine the power cap management mode used in the Cisco UCS domain:

- **Manual Blade Level Cap**—Power allocation is configured on each individual blade server in all chassis. If you select this option, you cannot create power groups.

- **Policy Driven Chassis Group Cap**—Power allocation is configured at the chassis level through power control policies included in the associated service profiles. If you select this option, you can also create power groups that contain one or more chassis in the Cisco UCS domain.

By default, power allocation is done for each chassis through a power control policy.

Step 6   Click **Save Changes**.

Power Management During Power-on Operations

Boot Staggering during Power on

Cisco UCS Manager attempts to boot as many blades as possible based on the amount of available power. If the power required to boot a blade is not available, Cisco UCS Manager staggers the boot in the Finite State Machine (FSM) CheckPowerAvailability stage, and raises the following fault on the blade: Insufficient power available to power-on server x/y.

When the required power becomes available, the FSM proceeds with blade power on. After a blade powers off, the allocated power budget is reclaimed.

Note   When the power budget that was allocated to the blade is reclaimed, the allocated power displays as 0 Watts.

Limitation

If you power on a blade outside of the Cisco UCS Manager and if there is not enough power available for allocation, the following fault is raised:

```
Power cap application failed for server x/y
```

Power Allocation during Service Profile Association

The power allocated to a blade during service profile association depends on the Power Control Policy used, and the power that is available from the power group. After the power is allocated to a server during a successful service profile association, the blade is guaranteed the minimum power cap. If the Power Control Policy priority is set to no-cap, a blade is allocated a potential maximum power cap, which might exceed the measured maximum power cap that displays.

Note   If the priority of an associated blade is changed to no-cap, and is not able to allocate the maximum power cap, you might see one of the following faults:

- `PSU-insufficient`—There is not enough available power for the PSU.

- `Group-cap-insufficient`—The group cap value is not sufficient for the blade.

Power Sync Policy Configuration

Power Sync Policy

Cisco UCS Manager includes a global (default) power sync policy to address power synchronization issues between the associated service profiles and the servers. You can use the power sync policy to synchronize the power state when the power state of the service profile differs from the actual power state of the server. The policy allows you to control when to synchronize the power state on the associated service profiles for the servers. The power sync policy does not affect other power-related policies.

The power synchronization policy applies to all the service profiles by default. You cannot delete the default power sync policy, but you can edit the default policy. You can create your own power sync policies and apply them to the service profiles. You can also create a power sync policy that is specific to a service profile and it always takes precedence over the default policy.

Cisco UCS Manager creates a fault on the associated service profile when the power sync policy referenced in the service profile does not exist. Cisco UCS Manager automatically clears the fault once you create a power sync policy for the specified service profile or change the reference to an existing policy in the service profile.

Power Synchronization Behavior

Cisco UCS Manager synchronizes the power state only when the actual power state of the server is OFF. The current power synchronization behavior is based on the actual power state and the preferred power state after shallow association occurs.

For example, the following events trigger shallow association:

- Fabric Interconnects(FI) and IOM disconnected.
- IOM reset
- FI power loss or reboot
- Chassis reacknowledgment
- Chassis power loss
- Service profile change

The following table describes the current power synchronization behavior:

Event	Preferred Power State	Actual Power State Before Event	Actual Power State After Event
Shallow Association	ON	OFF	ON
Shallow Association	OFF	OFF	OFF
Shallow Association	ON	ON	ON
Shallow Association	OFF	ON	ON

Creating a Power Sync Policy

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 Right-click **Power Sync Policies** and choose **Create Power Sync Policy**.

Step 5 In the **Create Power Sync Policy** dialog box, complete the following fields:

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 is saved.
Description field	A description of the policy. Cisco recommends including information about where and when to use the policy. 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).

Name	Description
Sync-Option field	The options that allow you to synchronize the desired power state of the associated service profile to the physical server. This can be one of the following: • **Default Sync**—After the initial server association, any configuration change or management connectivity changes that you perform trigger a server reassociation. This option synchronizes the desired power state to the physical server if the physical server power state is off and the desired power state is on. This is the default behavior. • **Always Sync**—When the initial server association or the server reassociation occurs, this option synchronizes the desired power state to the physical power state, even if the physical server power state is on and desired power state is off. • **Initial Only Sync**—This option only synchronizes the power to a server when a service profile is associated to the server for the first time, or when the server is re-commissioned. When you set this option, resetting the power state from the physical server side does not affect the desired power state on the service profile.

Step 6   Click **OK**.

What to do next

Include the policy in a service profile or service profile template.

Changing a Power Sync Policy

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   Choose a service profile policy from the **root** node.

Step 5   In the **Work** pane, click the **Policies** tab.

Step 6   Click the **Change Power Sync Policy** from the **Actions** area.

The information displayed depends on what you choose in the **Select the Power Sync Policy** drop-down list. You can choose:

- • **No Power Sync Policy**—If you choose this option, Cisco UCS Manager GUI does not display any other information. When you choose this option, Cisco UCS Manager implicitly uses the default power sync policy. Cisco UCS Manager searches for the default power sync policy under service profile organizations. If the policy is not found, then it uses the default power sync policy under root.

- • **Use an Existing Power Sync Policy**—if you want to select a global policy. Cisco UCS Manager GUI displays the **Power Sync Policy** drop-down list that enables you to choose an existing policy.

- • **Create a Local Power Sync Policy**—if you want to create a power sync policy that can only be accessed by this service profile. You can also create a power sync policy by using the **Create Power Sync Policy** link from the Power Sync Policy area.

Deleting a Power Sync Policy

Procedure

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

Step 2   Expand **Servers** > **Policies** > *Organization_Name*.

Step 3   Expand the **Power Sync Policies** node.

Step 4   Right-click the policy you want to delete and select **Delete**.

Step 5   If a confirmation dialog box displays, click **Yes**.

Rack Server Power Management

Power capping is supported for following rack servers:

- • Cisco UCS C220 M4 Server

- • Cisco UCS C240 M4 Server

UCS Mini Power Management

You can manage power of the blade servers in 6324 Fabric Interconnect (FI), which is used for remote offices and branch sites, and for limited server deployments. UCS Manager supports Dual Line Power Supply Unit and 110V when used with the 6324 Fabric Interconnect. You can manager how you want to allocate power when using 110V power supplies, because they might not provide enough power for a fully loaded chassis. Dual power supplies is standard for both AC and DC-48V on the UCS Mini 6324.