



# Managing Power in Cisco UCS

---

This chapter includes the following sections:

- [Power Management in Cisco UCS](#) , page 1
- [Rack Server Power Management](#), page 1
- [Power Management Precautions](#), page 1
- [Configuring the Power Policy](#), page 2
- [Configuring the Global Cap Policy](#), page 3
- [Configuring Policy-Driven Chassis Group Power Capping](#), page 3
- [Configuring Manual Blade-Level Power Capping](#), page 8

## Power Management in Cisco UCS

You can manage power through Cisco UCS Manager by configuring any of the following features:

- Power supply redundancy for all chassis in a Cisco UCS domain
- Policy-driven chassis-level power capping
- Manual blade-level power capping

## Rack Server Power Management

Power capping is not supported for rack servers.

## Power Management Precautions

If the CIMC is reset, the power monitoring functions of Cisco UCS become briefly unavailable for as long as it takes for the CIMC to reboot. While this usually only takes 20 seconds, there is a possibility that the peak power cap could be exceeded during that time. To avoid exceeding the configured power cap in a very low power-capped environment, consider staggering the rebooting or activation of CIMCs.

# Configuring the Power Policy

## Power Policy

The power policy is a global policy that specifies 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

	Command or Action	Purpose
<b>Step 1</b>	UCS-A# <b>scope org</b> <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type <i>/</i> as the <i>org-name</i> .
<b>Step 2</b>	UCS-A /org # <b>scope psu-policy</b>	Enters PSU policy mode.
<b>Step 3</b>	UCS-A /org/psu-policy # <b>set redundancy</b> { <b>grid</b>   <b>n-plus-1</b>   <b>non-redund</b> }	<p>Specifies one of the following redundancy types:</p> <ul style="list-style-type: none"> <li>• <b>grid</b> —Provides power redundancy when two power sources are used to power the chassis. If one power source fails, the surviving power supplies on the other power circuit continue to provide power to the chassis.</li> <li>• <b>n-plus-1</b> —Balances the power load for the chassis across the number of power supplies needed to satisfy non-redundancy plus one additional power supply for redundancy. If any additional power supplies are installed, they are recognized and powered off.</li> <li>• <b>non-redund</b> —Balances the power load for the chassis evenly across all installed power supplies.</li> </ul> <p>For more information about power redundancy, see the <i>Cisco UCS 5108 Server Chassis Installation Guide</i>.</p>
<b>Step 4</b>	UCS-A /org/psu-policy # <b>commit-buffer</b>	Commits the transaction to the system configuration.

The following example configures the power policy to use grid redundancy and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # scope psu-policy
UCS-A /org/psu-policy # set redundancy grid
```

```
UCS-A /org/psu-policy* # commit-buffer
UCS-A /org/psu-policy #
```

## Configuring the Global Cap Policy

### Global Cap Policy

The global cap policy is a global policy that specifies whether policy-driven chassis group power capping or manual blade-level power capping will be applied to all servers in a chassis.

We recommend that you use the default power capping method: policy-driven chassis group power capping.



#### Important

Any change to the manual blade-level power cap configuration will result in the loss of any groups or configuration options set for policy-driven chassis group power capping.

## Configuring the Global Cap Policy

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	UCS-A# <b>scope power-cap-mgmt</b>	Enters power cap management mode.
<b>Step 2</b>	UCS-A /power-cap-mgmt # <b>set cap-policy</b> { <b>manual-blade-level-cap</b>   <b>policy-driven-chassis-group-cap</b> }	Sets the global cap policy to the specified power cap management mode.  By default, the global cap policy is set to policy driven chassis group cap.
<b>Step 3</b>	UCS-A /power-cap-mgmt # <b>commit-buffer</b>	Commits the transaction to the system configuration.

The following example sets the global cap policy to manual blade power cap and commits the transaction:

```
UCS-A# scope power-cap-mgmt
UCS-A /power-cap-mgmt # set cap-policy manual-blade-level-cap
UCS-A /power-cap-mgmt* # commit-buffer
UCS-A /power-cap-mgmt #
```

## Configuring Policy-Driven Chassis Group Power Capping

### Policy-Driven Chassis Group Power Capping

When policy-driven power chassis group power capping is selected in the global cap policy, Cisco UCS can maintain the oversubscription of servers without risking costly power failures. This is achieved through a

two-tier process. At the chassis level, Cisco UCS divides the amount of power available between members of the power group. At the blade level, the amount of power allotted to a chassis is divided between blades based on priority.

Each time a service profile is associated or disassociated, 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 in order 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 or shut down.

---

## Power Groups

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 and then set a peak power cap in AC watts for that power grouping.

Instituting power capping at the chassis level requires the following:

- IOM, CIMC, and BIOS version 1.4 or higher
- 2 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 1556 AC watts should be set for each chassis. This converts to 1400 watts of DC power, which is the minimum amount of power required to power a fully-populated chassis.

Once a chassis is added to a power group, every service profile associated with the blades in the chassis also becomes 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.

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	This message is displayed if you did not meet the minimum limit when assigning the power cap for a chassis.	Increase the power cap limit to above 1556 watts in AC (1400 watts in DC).
Insufficient power available to discover server Chassis ID/BladeID	This message is displayed when you introduce a new blade and the boot power available for blade discovery is insufficient.	You can reassign the power budget by doing one of the following: <ul style="list-style-type: none"> <li>• Remove or decommission a chassis or blade in the power group.</li> <li>• Raise the group power budget.</li> <li>• Decrease the priority associated with a blade in the group.</li> </ul>

## Creating a Power Group

### Before You Begin

Make sure the global power allocation policy is set to Policy Driven Chassis Group Cap.

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	UCS-A# <b>scope power-cap-mgmt</b>	Enters power cap management mode.
<b>Step 2</b>	UCS-A /power-cap-mgmt # <b>create power-group power-group-name</b>	Creates a power group and enters power group mode.
<b>Step 3</b>	UCS-A /power-cap-mgmt/power-group # <b>set peak {peak-num   disabled   uninitialized}</b>	Specifies the maximum peak power (in watts) available to the power group.
<b>Step 4</b>	UCS-A /power-cap-mgmt/power-group # <b>create chassis chassis-id</b>	Adds the specified chassis to the power group and enters power group chassis mode.
<b>Step 5</b>	UCS-A /power-cap-mgmt/power-group/chassis # <b>commit-buffer</b>	Commits the transaction to the system configuration.

The following example creates a power group called powergroup1, specifies the maximum peak power for the power group (10000 watts), adds chassis 1 to the group, and commits the transaction:

```
UCS-A# scope power-cap-mgmt
UCS-A /power-cap-mgmt # create power-group powergroup1
UCS-A /power-cap-mgmt/power-group* # set peak 10000
```

```
UCS-A /power-cap-mgmt/power-group* # create chassis 1
UCS-A /power-cap-mgmt/power-group/chassis* # commit-buffer
UCS-A /power-cap-mgmt/power-group/chassis #
```

## Deleting a Power Group

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	UCS-A# <b>scope power-cap-mgmt</b>	Enters power cap management mode.
<b>Step 2</b>	UCS-A /power-cap-mgmt # <b>delete power-group</b> <i>power-group-name</i>	Deletes the specified power group.
<b>Step 3</b>	UCS-A /power-cap-mgmt/power-group/chassis # <b>commit-buffer</b>	Commits the transaction to the system configuration.

The following example deletes a power group called powergroup1 and commits the transaction:

```
UCS-A# scope power-cap-mgmt
UCS-A /power-cap-mgmt # delete power-group powergroup1
UCS-A /power-cap-mgmt* # commit-buffer
UCS-A /power-cap-mgmt #
```

## 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.

For mission-critical application a special priority called no-cap is also available. Setting the priority to no-cap prevents Cisco UCS from leveraging unused power from a particular server. With this setting, the server is allocated the maximum amount of power possible for that type of server.



### Note

You must include this 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

	Command or Action	Purpose
<b>Step 1</b>	UCS-A# <b>scope org</b> <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the org-name.
<b>Step 2</b>	UCS-A /org # <b>create power-control-policy</b> <i>power-control-pol-name</i>	Creates a power control policy and enters power control policy mode.
<b>Step 3</b>	UCS-A /org/power-control-policy # <b>set priority</b> { <i>priority-num</i>   <b>no-cap</b> }	Specifies the priority for the power control policy.
<b>Step 4</b>	UCS-A /org/power-control-policy # <b>commit-buffer</b>	Commits the transaction to the system configuration.

The following example creates a power control policy called powerpolicy15, sets the priority at level 2, and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # create power-control-policy powerpolicy15
UCS-A /org/power-control policy* # set priority 2
UCS-A /org/power-control policy* # commit-buffer
UCS-A /org/power-control policy #
```

### What to Do Next

Include the power control policy in a service profile.

## Deleting a Power Control Policy

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	UCS-A# <b>scope org</b> <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the org-name.
<b>Step 2</b>	UCS-A /org # <b>delete power-control-policy</b> <i>power-control-pol-name</i>	Deletes the specified power control policy.
<b>Step 3</b>	UCS-A /org # <b>commit-buffer</b>	Commits the transaction to the system configuration.

The following example deletes a power control policy called powerpolicy15 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete power-control-policy powerpolicy15
UCS-A /org* # commit-buffer
UCS-A /org #
```

## Configuring Manual Blade-Level Power Capping

### Manual Blade-Level Power Capping

When manual blade-level power capping 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:

#### Enabled

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.

#### Disabled

No power usage limitations are imposed upon 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 manual blade-level power capping is configured 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.

#### Procedure

	Command or Action	Purpose
<b>Step 1</b>	UCS-A# <b>scope server</b> <i>chassis-id / server-id</i>	Enters chassis server mode for the specified server.



	Command or Action	Purpose
<b>Step 2</b>	UCS-A /chassis/server # <b>set power-budget committed {disabled   watts}</b>	Commits the server to one of the following power usage levels: <ul style="list-style-type: none"> <li>• <b>disabled</b> —Does not impose any power usage limitations on the server.</li> <li>• <b>watts</b> —Allows you to specify the upper level for power usage by the server. If you choose this setting, enter the maximum number of watts that the server can use. The range is 0 to 10000000 watts.</li> </ul>
<b>Step 3</b>	UCS-A /chassis/server # <b>commit-buffer</b>	Commits the transaction to the system configuration.
<b>Step 4</b>	UCS-A /chassis/server # <b>show power-budget</b>	(Optional) Displays the power usage level setting.

The following example limits the power usage for a server to 1000 watts and commits the transaction:

```
UCS-A# scope server 2/4
UCS-A /chassis/server # set power-budget committed 1000
UCS-A /chassis/server* # commit-buffer
UCS-A /chassis/server # show power-budget
Power Budget:
  Committed (W): 1100
  Oper Committed (W): Disabled

UCS-A /chassis/server #
```

## Viewing the Blade-Level Power Cap

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	UCS-A# <b>scope server chassis-id / server-id</b>	Enters chassis server mode for the specified server.
<b>Step 2</b>	UCS-A /chassis/server # <b>show stats</b>	Displays the power usage statistics collected for the server.

The following example shows the server power usage:

```
UCS-A# scope server 2/4
UCS-A /chassis/server # show stats

Mb Power Stats:
  Time Collected: 2010-04-15T21:18:04.992
  Monitored Object: sys/chassis-1/blade-2/board
  Suspect: No
```

```
Consumed Power (W): 118.285194  
Input Voltage (V): 11.948000  
Input Current (A): 9.900000  
Thresholded: Input Voltage Min
```

```
UCS-A /chassis/server #
```