



Power Capping and Power Management in Cisco UCS

- [Power Capping in Cisco UCS, on page 1](#)
- [Power Policy Configuration, on page 3](#)
- [Power Policy for Cisco UCS Servers, on page 3](#)
- [Configuring the Power Policy, on page 3](#)
- [Power Supply for Redundancy Method, on page 4](#)
- [Power Supply for Redundancy Method for Cisco UCSX-9508 Chassis, on page 4](#)
- [Configuring Policy Driven Chassis Group Power Capping, on page 5](#)
- [Policy Driven Chassis Group Power Capping, on page 5](#)
- [Power Control Policy, on page 5](#)
- [Power Save Mode, on page 12](#)
- [Acoustic Mode Fan Profile, on page 13](#)
- [Power Groups in UCS Manager, on page 16](#)
- [Blade Level Power Capping, on page 21](#)
- [Fan Control Policy Configuration, on page 23](#)
- [Global Power Profiling Policy Configuration, on page 25](#)
- [Global Power Allocation Policy Configuration, on page 25](#)
- [Power Sync Policy Configuration, on page 27](#)
- [Rack Server Power Management, on page 31](#)
- [UCS Mini Power Management , on page 31](#)
- [Viewing X-Fabric Module \(XFM\) Fan Status, on page 31](#)

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 M5/M6, and C480 M5/C480 M5 ML, C225 M6, and C245 M6 and Cisco UCS C220 M7 Server, Cisco UCS C240 M7 Server rack servers, UCS Mini, and mixed UCS domains.

Cisco UCS Manager supports power capping on the following:

- UCS 6300 Series Fabric Interconnects
- UCS 6324 Series Fabric Interconnects (Cisco UCS Mini)

- UCS 6400 Series Fabric Interconnects
- UCS 6500 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.



Note Cisco UCSX-9508 Chassis supports Policy Driven Chassis Group Cap.

When you choose to select Policy Driven Chassis Group Cap, Cisco UCS Manager calculates the power allotment for Cisco UCSX-9508 Chassis and when you choose to select Manual Blade Level Power Cap, Chassis Management Controller (CMC) calculates the power allotment for Cisco UCSX-9508 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.
Power Save Policy	Globally manages the chassis to maximize energy efficiency or availability.
Cisco UCSX-9508 Chassis Power Extended Policy	Manages the chassis to maximize energy efficiency or availability. Power Extended Policy is effective only when we have PSU Redundant Policy Mode. For example, the total power available can be extended when we have N+1, N+2 and Grid to PSU Redundancy modes.
Cisco UCSX-9508 Chassis Fan Control Policy	Manages you to control the fan speed to bring down server power consumption and noise levels.
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 7500 Watts) 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.
 - **N+2**—The total number of PSUs to satisfy non-redundancy, plus two 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 only when the Power Save mode is enabled. If the power to any PSU is disrupted, Cisco UCS Manager can recover without an interruption in service.

After consideration of redundancy, the effective number of PSUs are atleast two. For example, for N it is two PSUs, N+1 it is three PSUs, and for N+2 and Grid it is four PSUs.
- Note** **N+2** redundancy mode is supported only for Cisco UCSX-9508 Chassis. For all other chassis, Cisco UCS Manager treats N+2 mode as N+1 mode only.

- **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 Cisco UCSX-9508 Chassis, if there are six PSUs present in two grids each of three power sources in which Grid one is having slots 1, 2, 3 and Grid two is having slots 4, 5, 6 then Grid one is used. Upon any power loss, the Grid two will take over.

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

In addition to power supply redundancy, you can also choose to enable a Power Save Policy from the **Power Save Policy** area. For more information, see [Power Save Mode Policy, on page 12](#).

Note For Cisco UCSX-9508 Chassis, you can also choose to enable or disable chassis power extended policy from the **Cisco UCSX-9508 Chassis Power Extended Policy** area. Chassis Extended Policy works only with non-redundant policy as the extended power is derived from redundant PSU.

Step 6 Click **Save Changes**.

Power Supply for Redundancy Method

PSU Redundancy	Max Power @ 240 V
Grid	5000 Watts
N+1	7500 Watts
Non-Redundant	8280 Watts



Note This table is valid if there are four PSUs installed in the chassis.

Power Supply for Redundancy Method for Cisco UCSX-9508 Chassis

PSU Redundancy	Max Power @ 2800 W
Grid	8400 Watts
N+1	14000 Watts
N+2	11200 Watts
Non-Redundant	16800 Watts



Note This table is valid if there are six PSUs, each of 2800 watts installed in the chassis.

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.

Starting with Cisco UCS Manager 3.2(2), chassis dynamic power rebalance mechanism is enabled by default. The mechanism continuously monitors the power usage of the blade servers and adjusts the power allocation accordingly. Chassis dynamic power rebalance mechanism operates within the overall chassis power budget set by Cisco UCS Manager, which is calculated from the available PSU power and Group power.

For mission-critical application a special priority called **no-cap** is also available. Setting the priority to **no-cap** does not guarantee that a blade server gets maximum power all the time, however, it prioritizes the blade server over other servers during the chassis dynamic power rebalance budget allocations.



Note If all the blade servers are set with no-cap priority and all of them run high power consuming loads, then there is a chance that some of the blade servers get capped under high power usage, based on the power distribution done through dynamic balance.

Global Power Control Policy options are inherited by all the chassis managed by the Cisco UCS Manager.

Starting with Cisco UCS Manager 4.1(3), a global policy called Power Save Mode is available. It is disabled by default, meaning that all PSUs present remain active regardless of power redundancy policy selection. Enabling the policy restores the older behavior..

Starting with Cisco UCS Manager 4.1(2), the power control policy is also used for regulating fans in Cisco UCS C220 M5 and C240 M5 rack servers in acoustically-sensitive environments. The Acoustic setting for these fans is only available on these servers. On C240 SD M5 rack servers, Acoustic mode is the default mode.

Starting with Cisco UCS Manager 4.2(1), the power control policy is also used for regulating cooling in potentially high-temperature environments. This option is only available with Cisco UCS C220 M6, C240 M6, C225 M6, and C245 M6 rack servers and can be used with any fan speed option.

Starting with Cisco UCS Manager 4.3(2), a global policy called Cisco UCS X9508 Chassis Power Extended Policy. This option is only available with Cisco UCS X9508 Chassis.



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 multi tenancy, 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	<p>The name of the policy.</p> <p>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.</p>

Name	Description
Description field	<p>A description of the policy. Cisco recommends including information about where and when to use the policy.</p> <p>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).</p>

Name	Description
Fan Speed Policy drop-down	

Name	Description
	<p>Note For Cisco UCS C125 M5 Server, ensure that you select the same Fan Speed Policy for all the servers in an enclosure. Cisco UCS Manager applies the Fan Speed Policy of the server which gets associated last. Having the same Fan Speed Policy for the all the server ensures that the desired Fan Speed Policy is applied irrespective of which server is associated last.</p> <p>Fan speed is for rack servers only. This can be one of the following:</p> <ul style="list-style-type: none"> • 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. <p>Note The performance option is not supported on Cisco UCS C-Series M5 and M6 , M7 servers.</p> <ul style="list-style-type: none"> • 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. • Acoustic—The fan speed is reduced to reduce noise levels in acoustic-sensitive environments. Rather than regulating energy consumption and preventing component throttling as in other modes, the Acoustic option could result in short-term throttling to achieve a lowered noise level. Acoustic mode is available only on Cisco UCS C220 M5 Server, Cisco UCS C240 M5 Server, Cisco UCS C240 SD M5 Server, Cisco UCS C220 M6 Server, Cisco UCS C240 M6 Server , Cisco UCS C225 M6 Server, and Cisco UCS C245 M6 Server. and Cisco UCS C220 M7 Server, Cisco UCS C240 M7 Server

Name	Description
	<p>Note On C240 SD M5, C220 M6, C240 M6, C225 M6, and C245 M6 and Cisco UCS C220 M7 Server, Cisco UCS C240 M7 Server servers, Acoustic mode is the default mode. On all other platforms, Low Power mode is the default mode.</p> <ul style="list-style-type: none"> • Any—The server determines the optimal fan speed. <p>Note Cisco UCSX-9508 Chassis Fan Control Policy supports Balanced, Low Power, High Power, Max Power, and Acoustic speeds.</p>
Aggressive Cooling field	<p>Optional setting for potentially high-heat thermal environments. Enabling Aggressive Cooling draws more power, but can mitigate a potential for overheating.</p> <p>Aggressive Cooling is only supported for Cisco UCS C-Series M6 and M7 rack servers.</p> <p>The two options are:</p> <ul style="list-style-type: none"> • Disabled (default) • Enabled <p>Note The Aggressive Cooling option is independent of the fan speed setting.</p>

Name	Description
Power Capping field	<p>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:</p> <ul style="list-style-type: none"> • No Cap—The server runs at full capacity regardless of the power requirements of the other servers in its power group. <p>Note For Cisco UCS C-Series M5 and M6 and M7 servers, if you select No Cap in this field, ensure that you do not select Performance for Fan Speed Policy field. Associating a service profile with a server fails if you select Performance for fan speed policy, and No Cap for the power capping.</p> <ul style="list-style-type: none"> • 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. <p>When you select cap, Cisco UCS Manager GUI displays the Priority field.</p>
Priority field	<p>The priority the server has within its power group when power capping is in effect.</p> <p>Enter an integer between 1 and 10, where 1 is the highest priority.</p>

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 Save Mode

Power Save Mode Policy

Power Save Mode is a configurable chassis policy that allows you to non-disruptively apply bias for either energy efficiency (when enabled) or availability (when disabled). By default, the Power Save Policy will be disabled. Disabling the Power Save Mode policy allows all PSUs present to remain active, regardless of the Power Redundancy setting. Enabling the Power Save Policy sets the PSUs to active as per the power redundancy policy.



Note Today, when the requested power budget is less than the available power capacity, the additional PSU capacity is placed in Power Save Mode automatically. This increases the efficiency of active PSUs and minimizes energy wasted for conversion losses. However, there are a few use cases, where this default behavior can result in an outage:

1. Lightly loaded chassis that only requires 2X PSU to support the requested power policy (Grid) and the customer did NOT follow the installation guide recommendation regarding PSU input power connections. In this scenario, the chassis has both active PSUs connected to one feed, and the other two PSUs in Power Save mode connected to another feed. If the feed connected to the active PSUs is lost, the entire chassis will experience a service interruption.
2. A heavily loaded chassis that requires a 3X PSUs to support the requested power policy (N+1), and the customer's rack provides the chassis with dual feed. In this scenario, 3X PSUs are active and 1X PSU is placed in Power Save mode. If the feed connected to two of the active PSUs is lost (planned or unplanned), the customer could experience an outage if the load is greater than the remaining active PSU can support.

A Power Save mode policy can help avoid an outage situation.

The policy is global and is inherited by all chassis managed by the Cisco UCS Manager.

Creating a Power Save Policy

Use this process to create a global Power Save 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 Save Policy** area, check the **Enable** checkbox to enable the Global Power Control Policy.
 - Step 6** Click **Save Changes**.
-

Acoustic Mode Fan Profile

Acoustic Mode Fan Profile

The Acoustic Mode fan profile is available on Cisco UCS C220 M5 Server, C240 M5, C240 SD M5, C220 M6, C240 M6, C225 M6, and C245 M6 , C220 M7 and C240 M7 rack servers.

Setting up an Acoustic Mode fan policy lets you reduce the noise level on M5 and M6 rack servers. The higher capacity fans on M5 and M6, M7 servers increase the cooling capacity, but also create more acoustic noise. The standard fan profiles for M5 and M6, M7 servers (Low Power, Balanced, High Power, and Max Power) are designed to regulate the server to optimize energy consumption. The primary goal of these fan profiles is to prevent throttling of CPU's and peripherals.

The goal of Acoustic Mode is reducing fan speed to reduce noise levels in acoustic-sensitive environments. Power capping has no effect when Acoustic Mode is selected.

Acoustic Mode is supported from Cisco UCS Manager 4.1.1 onward. Acoustic Mode is the default mode for C240 SD M5, C220 M6, C240 M6, C225 M6, and C245 M6, C220 M7 and C240 M7 servers and is automatically selected in the GUI for this server. For all other M5 and M6, M7 servers, the default mode is Low Power.

Configuring Acoustic Mode

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 multi tenancy, expand the **root** node.
- Step 4** Right-click **Power Control Policies** and choose **Create Power Control Policy**. Although these steps use the Power Control menu, you are creating a Fan Policy, which is administered through these menus.
- 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).

Name	Description
<p>Fan Speed Policy drop-down</p>	<p>Fan speed is for C-Series Rack servers only. Acoustic mode is a fan policy available only on Cisco UCS C220 M5, C240 M5, C240 SD M5, C220 M6, C240 M6 , C220 M7 and C240 M7 , and C245 M6 Rack Servers.</p> <p>Fan speed can be one of the following:</p> <ul style="list-style-type: none"> • Acoustic—The fan speed is reduced to reduce noise levels in acoustic-sensitive environments. The Acoustic option can result in short-term throttling to achieve a lowered noise level. <p>Note For Cisco UCS C-Series M5 and M6 , M7 servers using Acoustic Mode, cap in the Power Capping field is automatically selected. Acoustic Mode is the default fan speed policy for C240 SD M5, C220 M6, C240 M6, C220 M7 and C240 M7, C225 M6, and C245 M6 Rack Servers.</p> <ul style="list-style-type: none"> • 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. <p>Note Performance Mode is not available on M5 and M6, M7 servers.</p>

Name	Description
Power Capping field	<p>Power Capping occurs when the demand for power within a power group exceeds the power supply. For Cisco UCS C-Series M5 and M6 , M7 servers using Acoustic Mode, cap in the Power Capping field is automatically selected.</p> <p>Note Acoustic Mode is the default fan speed policy for C240 SD M5, C220 M6, C240 M6, C220 M7 and C240 M7, C225 M6, and C245 M6 Rack Servers and will automatically be selected, along with the cap option.</p> <ul style="list-style-type: none"> • No Cap —Allows you to set a priority for the server power throttling when Acoustic mode is selected. • cap—The server is allocated an amount of power capacity based on Acoustic Mode's need for power throttling and the server's priority relative to the other servers in its server group. <p>When cap is selected, Cisco UCS Manager GUI displays the Priority field.</p>
Priority field	<p>The priority the server has within its power group when power capping is in effect.</p> <p>Enter an integer between 1 and 10, where 1 is the highest priority. For Acoustic Mode, the default is 5.</p>

Step 6 Click OK.

What to do next

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

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
<pre>Insufficient budget for power group POWERSGROUP_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</pre>	<p>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.</p>	<p>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.</p>

Error Message	Cause	Recommended Action
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.



Note Beginning with release 4.3(2b), Cisco UCS Manager supports Power Group creation for Cisco UCSX-9508 Chassis.

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
Input Power(W) field	The maximum peak power (in watts) available to the power group. Enter an integer between 0 and 10000000.
Recommended value for Input Power field	The recommended range of input power values for all the members of the power group.

- 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.



Note Cisco UCSX-9508 Chassis does not support Manual Blade Level Power Cap. When you choose to select Manual Blade Level Power Cap, Chassis Management Controller (CMC) calculates the power allotment for Cisco UCSX-9508 Chassis.

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 1300 watts.



Note B480 M5 systems using 256GB DIMMs must have a manual blade level cap at 1300W.

- **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:
- Click the **Expand** icon to the right of the heading to display the fields.
 - Complete the following fields:

Name	Description
Admin Status field	Whether this server is power capped. This can be one of the following: <ul style="list-style-type: none"> • Unbounded—The server is not power capped under any circumstances. • Enabled—The Cisco UCS Manager GUI displays the Watts field. <p>Note Manual blade level power capping will limit the power consumption of a single system, regardless of available power in the chassis.</p>
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.
-

Fan Control Policy Configuration

Fan Control Policy

Fan Control Policies enable you to control the fan speed to bring down server power consumption and noise levels. With the introduction of Fan Control policies, you can determine the right fan speed for the server, based on the components in the server.

Globally managing the fan speed can help in power management by applying a single policy for all B-series server fans in an enclosure, based on general cooling needs. Set the fan speed on a per-chassis basis in the Global Policies.

Fan Control policy options include:

- **Balanced**—The fan runs at a faster speed when needed, based on the heat generated by the server. When possible, the fan returns to the minimum required speed. This is the default option.
- **Low Power**—The fan runs at the minimum speed that is required to keep the server cool.

Fan Control Policy for Cisco UCSX-9508 Chassis

Fan Control Policy enables you to control the fan speed to bring down server power consumption and noise levels of Cisco UCSX-9508 Chassis. With the introduction of Fan Control policies, you can determine the right fan speed for the server, based on the components in the server.

Globally managing the fan speed can help in power management by applying a single policy for all B-series and X-series server fans in an enclosure, based on general cooling needs. For X-series servers, set the fan speed on a per-chassis basis in the Global Policies.

Fan Control policy options include:

- **Balanced**—The fan runs at a faster speed when needed, based on the heat generated by the server. When possible, the fan returns to the minimum required speed. This is the default option.
- **Low Power**—The fan runs at the minimum speed that is required to keep the server cool.
- **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.
- **Acoustic**—The fan speed is reduced to reduce noise levels in acoustic-sensitive environments. Rather than regulating energy consumption and preventing component throttling as in other modes, the **Acoustic** option could result in short-term throttling to achieve a lowered noise level.

Creating a Fan Control Policy

You can create a Fan Control Policy and define the right fan control setting based on the server configuration and server components.

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** In the **Fan Control Policy** area, click one of the following radio buttons to define the fan control setting:
- **Balanced**—This setting can cool almost any server configuration. This is the default option.
 - **Low Power**—This setting is ideal for minimal configuration servers.
- Step 5** Click **Save Changes**.
-

Creating a Fan Control Policy for Cisco UCSX-9508 Chassis

You can create a Fan Control Policy for Cisco UCSX-9508 Chassis and define the right fan control setting based on the server configuration and server components.

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** Go to **Policies > Global Policies > Cisco UCSX-9508 Chassis Fan Control Policy**. Click one of the following radio buttons to define the fan control setting:
- **Balanced**—The fan runs at a faster speed when needed, based on the heat generated by the server. When possible, the fan returns to the minimum required speed. This is the default option.
 - **Low Power**—The fan runs at the minimum speed that is required to keep the server cool.
 - **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.
 - **Acoustic**—The fan speed is reduced to reduce noise levels in acoustic-sensitive environments. Rather than regulating energy consumption and preventing component throttling as in other modes, the **Acoustic** option could result in short-term throttling to achieve a lowered noise level.
- Step 5** Click **Save Changes**.
-

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.

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.



Note Cisco UCSX-9508 Chassis supports Policy Driven Chassis Group Cap only.

When you choose to select Policy Driven Chassis Group Cap, Cisco UCS Manager calculates the power allotment for Cisco UCS X9508 chassis and when you choose to select Manual Blade Level Power Cap, Chassis Management Controller (CMC) calculates the power allotment for Cisco UCSX-9508 Chassis.



Note For Cisco UCSX-9508 Chassis **Allocated (W)** and **Measured Max. (W)** will not match. The max allocated values are used to calculate the chassis-level power limit and Intelligent Fabric Modules (IFM) allocates the power based on the power limit.

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.

Note Cisco UCS X9508 chassis supports **Policy Driven Chassis Group Cap**.

When you choose to select Policy Driven Chassis Group Cap, Cisco UCS Manager calculates the power allotment for Cisco UCS X9508 chassis and when you choose to select Manual Blade Level Power Cap, Chassis Management Controller (CMC) calculates the power allotment for Cisco UCS X9508 chassis.

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 multi tenancy, 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	<p>The name of the policy.</p> <p>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.</p>
Description field	<p>A description of the policy. Cisco recommends including information about where and when to use the policy.</p> <p>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).</p>
Sync-Option field	<p>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:</p> <ul style="list-style-type: none"> • 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 multi tenancy, 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 M5 Server
- Cisco UCS C240 M5 Server
- Cisco UCS C240 SD M5 Server
- Cisco UCS C480 M5 Server
- Cisco UCS C480 M5 ML Server
- Cisco UCS C220 M6 Server
- Cisco UCS C240 M6 Server
- Cisco UCS C225 M6 Server
- Cisco UCS C245 M6 Server

Power capping is not supported for Cisco UCS C125 M5 Servers.

UCS Mini Power Management

You can manage power of the blade servers in the Cisco UCS 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 Cisco UCS 6324 Fabric Interconnect. You can manage 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 Cisco UCS Mini 6324.

Viewing X-Fabric Module (XFM) Fan Status

This procedure is applicable only for Cisco UCS X9508 Server Chassis equipped with XFM.

Procedure

- Step 1** In the **Navigation** pane, click **Equipment**.
- Step 2** Expand **Equipment** > **Chassis** > **Chassis Number** > **XFM Modules** > **XFM Module Number** > **Fans** > **Fan Module Number** > **Equipment** > **Chassis** > **Chassis Number** > **IO Modules**.
- Step 3** Choose the Fan number for which you want to view.
- Step 4** In the Work pane, click the **General** tab.

The overall status for this fan appear. The fields in this tab are:

Column	Description
Name column	A navigation tree that allows you to view a particular component and its subcomponents. You can right-click a component to view any actions available for that component.
Operability column	A brief description of the operating state of the component. If a fan is inoperable, it can be replaced with a new fan module. Contact Cisco technical support for more information.
Performance column	A brief description of the performance state of the component.
Power column	A brief description of the power state of the component.
Temperature column	A description of the temperature state of the component.
