



Power Management in Cisco UCS

- [Power Capping in Cisco UCS, on page 1](#)
- [Power Policy Configuration, on page 2](#)
- [Policy Driven Power Capping, on page 4](#)
- [Blade Level Power Capping, on page 10](#)
- [Global Power Profiling Policy Configuration, on page 14](#)
- [Global Power Allocation Policy, on page 15](#)
- [Power Management During Power-on Operations, on page 16](#)
- [Power Sync Policy Configuration, on page 17](#)
- [Rack Server Power Management, on page 24](#)
- [UCS Mini Power Management , on page 25](#)

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/M5/M6, and C480 M5/C480 M5 ML, C225 M6, C245 M6 rack servers, UCS Mini, and mixed UCS domains.

Cisco UCS Manager supports power capping on the following:

- UCS 6500 Series Fabric Interconnects
- UCS 6400 Series Fabric Interconnects
- UCS 6300 Series Fabric Interconnects
- UCS 6324 Series Fabric Interconnects (Cisco UCS Mini)
- UCS 6200 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.
Power Save Policy	Globally manages the chassis to maximize energy efficiency or availability.
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

	Command or Action	Purpose
Step 1	UCS-A# scope org <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> .
Step 2	UCS-A /org # scope psu-policy	Enters PSU policy mode.
Step 3	UCS-A /org/psu-policy # set redundancy { grid n-plus-1 non-redund }	Specifies one of the following redundancy types: <ul style="list-style-type: none"> • grid —Two power sources are turned on, or the chassis requires greater than N+1 redundancy. If one source fails (which

	Command or Action	Purpose
		<p>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.</p> <ul style="list-style-type: none"> • n-plus-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. • non-redund —All installed power supplies (PSUs) are turned on and the load is evenly balanced. Only smaller configurations (requiring less than 2500W) can be powered by a single PSU. <p>For more information about power redundancy, see the <i>Cisco UCS 5108 Server Chassis Installation Guide</i>.</p>
Step 4	Required: UCS-A /org/psu-policy # commit-buffer	Commits the transaction to the system configuration.

Example

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 #
```

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.

Policy Driven 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.



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

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the org-name.
Step 2	UCS-A /org # create power-control-policy <i>power-control-pol-name</i>	Creates a power control policy and enters power control policy mode.
Step 3	UCS-A /org/power-control-policy # set fanspeed { <i>any</i> balanced <i>high-power</i> <i>low-power</i> <i>max-power</i> <i>performance</i> <i>acoustic</i> }	Specifies the fan speed for the power control policy. Note The performance option is not supported on Cisco UCS C-Series M5 and M6 servers.
Step 4	UCS-A /org/power-control-policy # set priority { <i>priority-num</i> no-cap }	Specifies the priority for the power control policy.
Step 5	UCS-A /org/power-control-policy # commit-buffer	Commits the transaction to the system configuration.

Example

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.

Configuring Acoustic Mode

Acoustic Mode

Acoustic mode is a fan policy available only on Cisco UCS C220 M5 Server, C220 M6 Server, C240 M5 Server, C240 M6 Server, and C240 SD M5 Server Rack Servers and is supported from Cisco UCS Manager Release 4.1.1 onward.

The available fan policy options for these M5 and M6 servers are Acoustic, Low power, Balanced, High Power, and Max power.

On C240 SD M5 Server, C220 M6 Server, C240 M6 Server, and C245 M6 Server Acoustic mode is the default mode. On all other platforms, Low Power mode is the default mode.

The primary goal of Acoustic mode is to reduce the noise level emitted by the fans by reducing the fan speed. The standard fan policies are designed for optimal energy consumption and preventing any component throttling. Acoustic mode reduces noise but carries a higher probability of having short term throttling effects.

Acoustic Mode is independent of the power management features.

Creating an Acoustic Mode Fan Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the org-name.
Step 2	UCS-A /org # create power-control-policy <i>fan-policy-name</i>	Creates a fan control policy and enters power control policy mode. Fan policies are created through the power control interface.
Step 3	UCS-A /org/power-control-policy # set fanspeed { acoustic }	Specifies Acoustic Mode as the fan speed for the power control policy.
Step 4	UCS-A /org/power-control-policy # set priority { <i>priority-num</i> no-cap }	Specifies the priority for the fan's power control policy.

	Command or Action	Purpose
Step 5	UCS-A /org/power-control-policy # commit-buffer	Commits the transaction to the system configuration.

What to do next

Include the power control policy in a service profile.

Deleting a Power Control Policy**Procedure**

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

Example

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 #
```

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.

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

Ensure that the global power allocation policy is set to Policy Driven Chassis Group Cap.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope power-cap-mgmt	Enters power cap management mode.
Step 2	UCS-A /power-cap-mgmt # create power-group <i>power-group-name</i>	Creates a power group and enters power group mode.
Step 3	UCS-A /power-cap-mgmt/power-group # set peak { <i>peak-num</i> disabled uninitialized }	Specifies the maximum peak power (in watts) available to the power group.
Step 4	UCS-A /power-cap-mgmt/power-group # create chassis <i>chassis-id</i>	Adds the specified chassis to the power group and enters power group chassis mode.
Step 5	UCS-A /power-cap-mgmt/power-group # create rack <i>rack-id</i>	Adds the specified rack to the power group.
Step 6	UCS-A /power-cap-mgmt/power-group # create fex <i>fex-id</i>	Adds the specified FEX to the power group.

	Command or Action	Purpose
Step 7	UCS-A /power-cap-mgmt/power-group # create fi <i>fi-id</i>	Adds the specified FI to the power group.
Step 8	UCS-A /power-cap-mgmt/power-group/chassis # commit-buffer	Commits the transaction to the system configuration.

Example

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

Example

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 #
```

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 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

Ensure that the global power allocation policy is set to Manual Blade Level Cap.

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>chassis-id / server-id</i>	Enters chassis server mode for the specified server.
Step 2	UCS-A /chassis/server # set power-budget committed { unbounded <i>watts</i> }	Commits the server to one of the following power usage levels: <ul style="list-style-type: none"> • unbounded —Does not impose any power usage limitations on the server. • <i>watts</i> —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.
Step 3	UCS-A /chassis/server # commit-buffer	Commits the transaction to the system configuration.

	Command or Action	Purpose
Step 4	UCS-A /chassis/server # show power-budget	(Optional) Displays the power usage level setting.

Example

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

```
UCS-A# scope server 1/7
UCS-A /chassis/server # show power-budget

Budget:
  AdminCommitted (W)
  -----
  139
UCS-A /chassis/server # set power-budget committed unbounded
UCS-A /chassis/server* # commit-buffer
UCS-A /chassis/server # show power-budget

Budget:
  AdminCommitted (W)
  -----
  Unbounded

UCS-A /chassis/server # set power-budget committed 1000
UCS-A /chassis/server* # commit-buffer
UCS-A /chassis/server # show power-budget

Budget:
  AdminCommitted (W)
  -----
  1000
UCS-A /chassis/server #
```

Configuring a Chassis Level Fan Policy

Configuring Fan Speed for Power Management

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. The two options are:

- **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. (Default.)
- **Low Power**—The fan runs at the minimum speed that is required to keep the server cool.

The new option takes effect when the new selection is saved. Use **Low Power** to save on system power.

Configuring the Global Fan Control 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 **Fan Control Policy** area, click one of the following radio buttons:
- **Balanced**—This is the default option.
 - **Low Power**
- Step 6** Click **Save Changes**.
-

Viewing Server Statistics

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope server <i>chassis-id / server-id</i>	Enters chassis server mode for the specified server.
Step 2	UCS-A /chassis/server # show stats	Displays the following server statistics: <ul style="list-style-type: none"> • Ethernet Port Error • Ethernet Port Multicast • Ethernet Port • Virtual Interface • Motherboard Power • PC Ie Fatal Completion Error • PC Ie Fatal Protocol Error • PC Ie Fatal Receiving Error • PC Ie Fatal Error • Memory Error • DIMM Env • CPU Env

Example

The following example shows the section on motherboard power usage statistics:

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

Motherboard Power Statistics:
Time Collected: 2016-07-11T20:51:24.722
Monitored Object: sys/chassis-1/blade-1/board/power-stats
Suspect: No
Consumed Power (W): 126.000000
Input Voltage (V): 11.859000
Input Current (A): 10.624842
Thresholded: 0

UCS-A /chassis/server #
```

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

	Command or Action	Purpose
Step 1	UCS-A# scope power-cap-mgmt	Enters power cap management mode.
Step 2	UCS-A /power-cap-mgmt # set profile-policy {no yes}	Enables or disables the global power profiling policy.
Step 3	UCS-A /power-cap-mgmt # commit-buffer	Commits the transaction to the system configuration.

Example

The following example shows how to enable the global power profile policy and commit the transaction:

```
UCS-A# scope power-cap-mgmt
UCS-A /power-cap-mgmt # set profile-policy yes
UCS-A /power-cap-mgmt* # commit-buffer
UCS-A /power-cap-mgmt #
```

Global Power Allocation Policy

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

	Command or Action	Purpose
Step 1	UCS-A# scope power-cap-mgmt	Enters power cap management mode.
Step 2	UCS-A /power-cap-mgmt # set cap-policy { manual-blade-level-cap policy-driven-chassis-group-cap }	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.
Step 3	UCS-A /power-cap-mgmt # commit-buffer	Commits the transaction to the system configuration.

Example

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 #
```

Viewing the Power Cap Values for Servers

Procedure

	Command or Action	Purpose
Step 1	UCS-A# scope power-cap-mgmt	Enters power cap management mode.
Step 2	UCS-A /power-cap-mgmt # show power-measured	Displays the minimum and maximum power cap values.

Example

The following example shows how to display the minimum and maximum power cap values:

```
UCS-A# scope power-cap-mgmt
UCS-A /power-cap-mgmt # show power-measured
```

```
Measured Power:
  Device Id (W)  Minimum power (W)  Maximum power (W)  OperMethod
-----
  blade  1/1    234                353                Pnuos
```

```
UCS-A /power-cap-mgmt #
```

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

Displaying the Global Power Sync Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the org-name.
Step 2	UCS-A/org # scope power-sync-policy default	Enters the global power sync policy mode.
Step 3	UCS-A /org/power/-sync-policy # show { detail expand detail expand }	Displays the global power sync policy information.

Example

The following example displays the global (default) power sync policy:

```
UCS-A # scope org
UCS-A /org # scope power-sync-policy default-sync
UCS-A /org/power-sync-policy # show expand
```

```
Power Sync Policy:
  Name                Power Sync Option
  -----
  default             Default Sync
```

```
UCS-A /org/power-sync-policy # show detail expand
```

```
Power Sync Policy:
  Full Name: org-root/power-sync-default
  Name: default
  Description:
  Power Sync Option: Default Sync
  Policy Owner: Local
```

```
UCS-A /org/power-sync-policy #
```

Setting Global Policy Reference for a Service Profile

To refer the global power sync policy in a service profile, use the following commands in service profile mode:

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the org-name.
Step 2	UCS-A/org # scope service-profile <i>service-profile-name</i>	Enters the service profile mode for the specified service profile. The name of the service profile can be a minimum of two characters and a maximum up to 32 characters.
Step 3	UCS-A /org/service-profile # set power-sync-policy default	Specifies the global power sync policy that can be referenced in the service profile. You can also change the policy reference from the default to other power sync policies using this command.
Step 4	UCS-A /org/service-profile* # commit-buffer	Commits the transaction to the system configuration.

Example

The following example sets the reference to the global power sync policy for use in the service profile.

```
UCS-A # scope org
UCS-A/org # scope service-profile spnew
UCS-A/org/service-profile # set power-sync-policy default
UCS-A/org/service-profile* # commit-buffer
```

Creating a Power Sync Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the org-name.
Step 2	UCS-A /org # create power-sync-policy <i>power-sync-pol-name</i>	Creates a power sync policy and enters power sync policy mode. The power sync policy name can be up to 16 characters.

	Command or Action	Purpose
Step 3	(Optional) UCS-A /org/power-sync-policy* # set descr <i>optionall-description</i>	Specifies the description of the power-sync-policy. You can also modify the description using the descr keyword.
Step 4	UCS-A /org/power-sync-policy* # set sync-option { always-sync default-sync initial-only-sync }	Specifies the power synchronization option to the physical server. You can also modify the power synchronization option using the sync-option keyword. This can be one of the following: <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 always synchronizes the desired power state to the physical server even if the physical server power state is on and the 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 5	UCS-A /org/power-sync-policy* # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates a power sync policy called newSyncPolicy, sets the default sync-option, and commits the transaction to the system configuration:

```
UCS-A # scope org
UCS-A /org # create power-sync-policy newSyncPolicy
UCS-A /org/power-sync-policy* # set decsr newSyncPolicy
UCS-A /org/power-sync-policy* # set sync-option default-sync
UCS-A /org/power-sync-policy* # commit-buffer
UCS-A /org/power-sync-policy #
```

What to do next

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

Deleting a Power Sync Policy

Procedure

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

Example

The following example deletes the power sync policy called spnew and commits the transaction to the system:

```
UCS-A # scope org
UCS-A /org # delete power-sync-policy spnew
UCS-A /org # commit-buffer
```

Displaying All Power Sync Policies

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the org-name.
Step 2	UCS-A /org # show power-sync-policy { detail expand detail expand }	Displays the default, local, and other power sync policies.

Example

The following example displays power sync policies that are defined:

```
UCS-A # scope org
UCS-A /org # show power-sync-policy expand
Power Sync Policy:
  Name                               Power Sync Option
  -----
```

```

default                Default Sync
policy-1               Default Sync

UCS-A /org # show power-sync-policy detail expand
Power Sync Policy:
  Full Name: org-root/power-sync-default
  Name: default
  Description:
  Power Sync Option: Default Sync
  Policy Owner: Local

  Full Name: org-root/power-sync-policy-1
  Name: policy-1
  Description:
  Power Sync Option: Default Sync
  Policy Owner: Local

UCS-A /org #

```

Creating a Local Policy

To create a local power sync policy that you want to use by any service profile, create a power sync definition for the power sync policy.

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the org-name.
Step 2	UCS-A /org # scope service-profile <i>service-profile-name</i>	Enters the service profile mode for the specified service profile. The name of the service profile can be a minimum of two characters and a maximum up to 32 characters.
Step 3	UCS-A /org/service-profile # create power-sync-definition	Enters the power sync definition mode. You can create a power sync policy definition that you defined for the power sync policy.
Step 4	(Optional) UCS-A /org/service-profile/power-sync-definition* # set descr <i>optional-description</i>	Specifies the description of the power-sync-policy. You can also change the description using the descr keyword.
Step 5	UCS-A /org/service-profile/power-sync-definition* # set sync-option { always-sync default-sync initial-only-sync }	Specifies the power synchronization option to the physical server. You can also change the power synchronization option using the sync-option keyword.
Step 6	UCS-A /org/service-profile/power-sync-definition* # commit-buffer	Commits the transaction to the system configuration.

Example

The following example creates a local policy using the policy sync definition, sets the sync-option, and commits the transaction to the system configuration:

```
UCS-A # scope org
UCS-A/org # scope service-profile spnew
UCS-A/org/service-profile # create power-sync-definition
UCS-A/org/service-profile/power-sync-definition* # set decsr spnew
UCS-A/org/service-profile/power-sync-definition* # set sync-option default-sync
UCS-A/org/service-profile/power-sync-definition* # commit-buffer
```

Showing a Local Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the org-name.
Step 2	UCS-A/org # scope service-profile <i>service-profile-name</i>	Enters the service profile mode for the specified service profile. The name of the service profile can be a minimum of two characters and a maximum up to 32 characters.
Step 3	(Optional) UCS-A /org/service-profile # show power-sync-policy { detail expand detail expand }	Displays the local policy in the power-sync-policy mode.
Step 4	UCS-A /org/service-profile # show power-sync-definition { detail expand detail expand }	Displays the local policy for the specified service policy in the power-sync-definition mode. Note If you do not have a definition for the power sync policy, you can still use the command, but you cannot see anything displayed.

Example

The following example displays the local policy in use by the service profile spnew:

```
UCS-A # scope org
UCS-A/org # scope service-profile spnew
UCS-A/org/service-profile # show power-sync-definition expand
```

```
Power Sync Definition:
  Name                Power Sync Option
  -----
  spnew                Always Sync
```

```
UCS-A/org/service-profile # show power-sync-definition detail expand

Power Sync Definition:
  Full Name: org-root/ls-sp2/power-sync-def
  Name: spnew
  Description: optional description
  Power Sync Option: Always Sync
  Policy Owner: Local

UCS-A/org/service-profile #
```

Deleting a Local Policy

Procedure

	Command or Action	Purpose
Step 1	UCS-A # scope org <i>org-name</i>	Enters organization mode for the specified organization. To enter the root organization mode, type / as the org-name.
Step 2	UCS-A/org # scope service-profile <i>service-profile-name</i>	Enters the service profile mode for the specified service profile. The name of the service profile can be a minimum of two characters and a maximum up to 32 characters.
Step 3	UCS-A /org/service-profile # delete power-sync-definition	Enters the power sync definition mode. You can delete a power sync policy definition that you defined for the power sync policy.
Step 4	UCS-A /org/service-profile* # commit-buffer	Commits the transaction to the system configuration.

Example

The following example deletes the local policy in use by the service profile.

```
UCS-A # scope org
UCS-A/org # scope service-profile spnew
UCS-A/org/service-profile # delete power-sync-definition
UCS-A/org/service-profile* # commit-buffer
```

Rack Server Power Management

Power capping is supported for following rack servers:

- Cisco UCS C220 M4 Server
- Cisco UCS C240 M4 Server
- 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.

