

# Cisco UCS C-Series M4 Server Power Management

August 25, 2014



Sriranjan Bose, Technical Leader, Cisco Data Center Group

Chitkala Sethuraman, Software Engineer, Cisco Data Center Group

Vikram Bodireddy, Software Engineer, Cisco Data Center Group

This white paper provides an overview of the features and best practices for the new advanced power management capabilities available with Cisco UCS<sup>®</sup> C-Series M4 rack servers.

## Introduction to Power Capping

Power capping provides the ability to limit the power consumption of a system, be it a blade server or a rack-mount server, to some threshold that is less than or equal to the system's maximum rated power. The feature typically offers incremental benefits with the ability to scale a greater number of systems, as well as providing even better control over power consumption.

There are two major types of power capping. **Static power capping** involves distributing a fixed budget of power across multiple servers by calculating the individual system consumption using its rated or nameplate power rating.

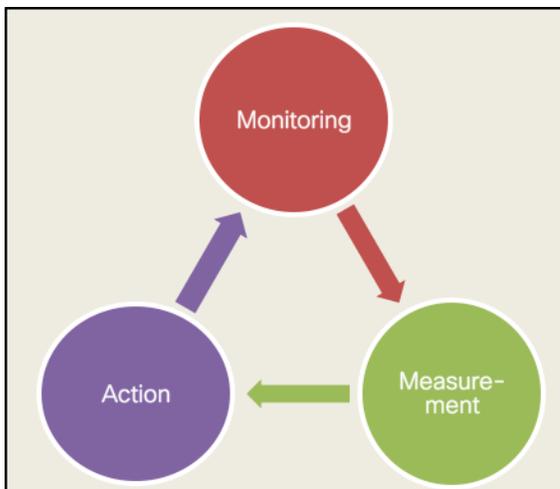
For example, if the maximum power rating of a server is given as 340 watts (W) and a rack in a data center is equipped with 10 servers, but the power available to the rack is only 3100W AC, the available power is sufficient to supply an average of 300W per rack server. In this case, each server can be capped at a maximum of 300W to avoid exceeding the capacity of the power supply. An obvious drawback of this approach is the inconsistency of power ratings across rack servers. This rating depends on the configuration of the server, affected by the types of DIMMs, CPUs, and adapters the system is equipped with. Another drawback is that the power consumption of the systems in the rack may not be peaking at the same time. So the server power consumption will be capped even when the overall power budget has not reached full capacity.<sup>5</sup>

---

Another type of capping, **dynamic power capping**, allows the power management system to allocate the total pool of power across groups of multiple systems. This type of power capping is extremely effective in situations where multiple blade servers reside within a single chassis and share the power supply. With dynamic power capping, the system as a whole can conform to a specific power budget, but power can be steered to the specific nodes that have a higher load and require additional power.

With the release of Cisco UCS C-Series M4, Cisco has extended the concept of power capping in scope and features for rack-mount servers. It provides the administrator with critical operational information, such as system utilization efficiency, while bridging the gap between static and dynamic power capping in the rack-mount form factor. The C-Series M4 generation of servers has advanced embedded systems management services that provide accurate server power ratings for efficient budgeting, as well as providing time-of-day-driven capping actions to allow the infrastructure to adapt to business needs. The solution is designed around the closed feedback loop concept of monitoring, measurement and action (Figure 1).

**Figure 1.** Closed feedback loop of monitoring, measurement and action



## Power Consumption and Its Impact on Data Center Design

Over the last decade, data center design priorities have moved from real estate and materials cost impact to infrastructure and power costs (Figure 2). The new priorities for designers are:

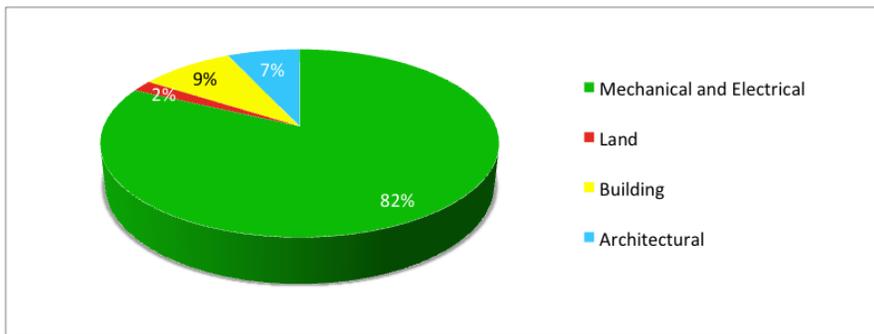
### Optimization of data center design<sup>2</sup>

- Follow best practices
- Adopt new technologies
- Optimize code
- Engineer the data center

## Right sizing

- Eliminate stranded compute resources
- Select the most efficient servers—app/watt
- Have the right amount of cores and DIMMs for applications
- Eliminate stranded power resources
- Eliminate stranded cooling resources

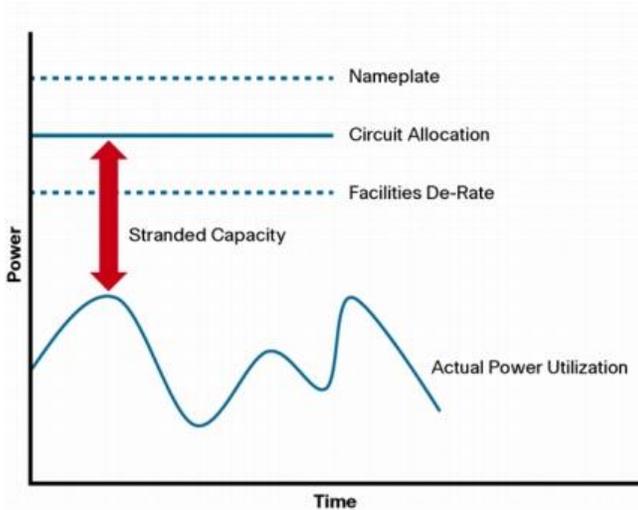
**Figure 2.** Capital Cost in the Data Center



(Source: "Incenting the Right Behaviors in the Data Center," C. Belady and G. Balakrishnan, [http://events.energetics.com/datacenters08/pdfs/Belady\\_Microsoft.pdf](http://events.energetics.com/datacenters08/pdfs/Belady_Microsoft.pdf))

The impact of these factors is further explored in the white paper "Power Management in the Cisco Unified Computing System: An Integrated Approach"<sup>5</sup>.

**Figure 3.** Data Center Power Allocation



(Source: "Power Management in the Cisco Unified Computing System: An Integrated Approach," [http://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/unified-computing/white\\_paper\\_c11-627731.html](http://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/unified-computing/white_paper_c11-627731.html))

---

Three major variables affect the optimization of power usage in a data center<sup>5</sup>:

- **Cooling and distribution system efficiency:** The cost of cooling the system is a sizable portion of the cost of the energy used to perform the workload.
- **Efficiency:** Although the overall power utilization efficiency (PUE) of a data center is outside the scope of the server design, the server design should still try to achieve maximum compute efficiency per watt of power used.
- **Capacity utilization:** It is important to have an accurate measure of the power provisioned for a system. Under-provisioning and overprovisioning can have a negative influence on data center performance. While the former can severely affect critical business operations at peak hours, the latter can lead to massive underutilization of infrastructure.

Apart from the operational efficiency aspects of power management, it is important to keep in mind the data center design aspect, such as over-current circuit protection. While compute end users and IT administrators may be concerned with the amount of power consumed per server and per chassis, the data center operations people are more sensitive to total power per circuit and the ability to protect circuit breakers from tripping.

A great deal of literature exists on the subject of power provisioning in data centers, including “Power Management in the Cisco Unified Computing System: An Integrated Approach”<sup>5</sup> and “Data Center Power and Cooling”<sup>13</sup>. A data center designer has to make multiple assumptions about the nature of the data center workload when provisioning the servers across the power buses.

Administrators and designers may assume that not all servers will draw the maximum power at the same time and also that, even at 100 percent utilization, they will not all draw the maximum power for which the system power supplies are rated. These assumptions help the designer oversubscribe the available power with respect to the number of systems deployed. Unfortunately, the nature of workloads in data centers is changing rapidly with widespread deployment of technologies such as virtualization and cloud-based storage and computing. These technologies are attempting to drive down the total cost of ownership (TCO) while maximizing the utilization of hardware infrastructure. The new workloads are making it more difficult to sustain the assumptions that allowed designers to oversubscribe their power supply infrastructure. Modern workloads are displaying increasingly unpredictable fluctuations in traffic and power load spikes<sup>12</sup>.

### Power Capping in Cisco UCS C-Series M4 Servers

Designers need to have easy access to key metrics of infrastructure utilization, in a convenient and reusable form. These metrics include power consumption patterns over periods of days; the maximum, minimum, and average trends for the entire platform; the distribution of consumption across subsystems in the servers (including CPU, memory, storage, and I/O); and the actual possible maximum power consumption of the server configuration. In rack-mount servers, which are open configuration systems with a huge matrix of possible on-board peripherals, attempting to get these numbers right has been a difficult and error-prone exercise. The penalty for miscalculating is severe—underprovisioning the power per system per rack might lead to a circuit breaker trip, eventually causing a data center outage. Overprovisioning, on the other hand, drives down the PUE of a data center. The Cisco UCS C-Series M4 servers attempt to help designers solve these problems by providing more accurate and precise telemetry and statistics.

The technology used to gather this information is built into the hardware of the systems, allowing levels of precision not possible in legacy technologies. Additionally, this solution is capable of scaling at the system level across different configurations, and at the data center level by providing an XML API for management and reporting.

---

Cisco's approach to solve this problem from the rack-mount server perspective has been to provide more control points for a nuanced configuration of hardware. By working together with Intel Corporation, we have been able to provide, in the Cisco® Integrated Management Controller (IMC), advanced telemetry and power capping profile configuration. Since each rack-mount server is equipped with its own mechanical parts such as fans and power supply units, it is necessary that the server be capable of making autonomous power capping decisions.

While having the tools to achieve efficient power capping is convenient, these tools are not enough to provide a complete, holistic solution. For the solution to be complete, it is also necessary for the data center administrator to have the means to measure and judge the policy decisions that will allow effective use of those tools. Cisco IMC provides key telemetry inputs that allow the data center administrator to make these important power budgeting decisions.

Additionally, the telemetry information is collected over a period of time, allowing the user to detect trends. Cisco IMC allows users to manage the power consumption of their servers using the time-of-day trends.

### Using the Cisco UCS C-Series M4 Power Capping Feature

The power capping feature on the Cisco UCS C-Series M4 with IMC is a powerful and potent tool that allows administrators and designers to maximize the PUE of their data centers and use it to maximum effectiveness. The use of this feature requires a methodical approach. The sections that follow provide a series of steps that will allow administrators to create an effective model for provisioning the power management solution across multiple C-Series servers.

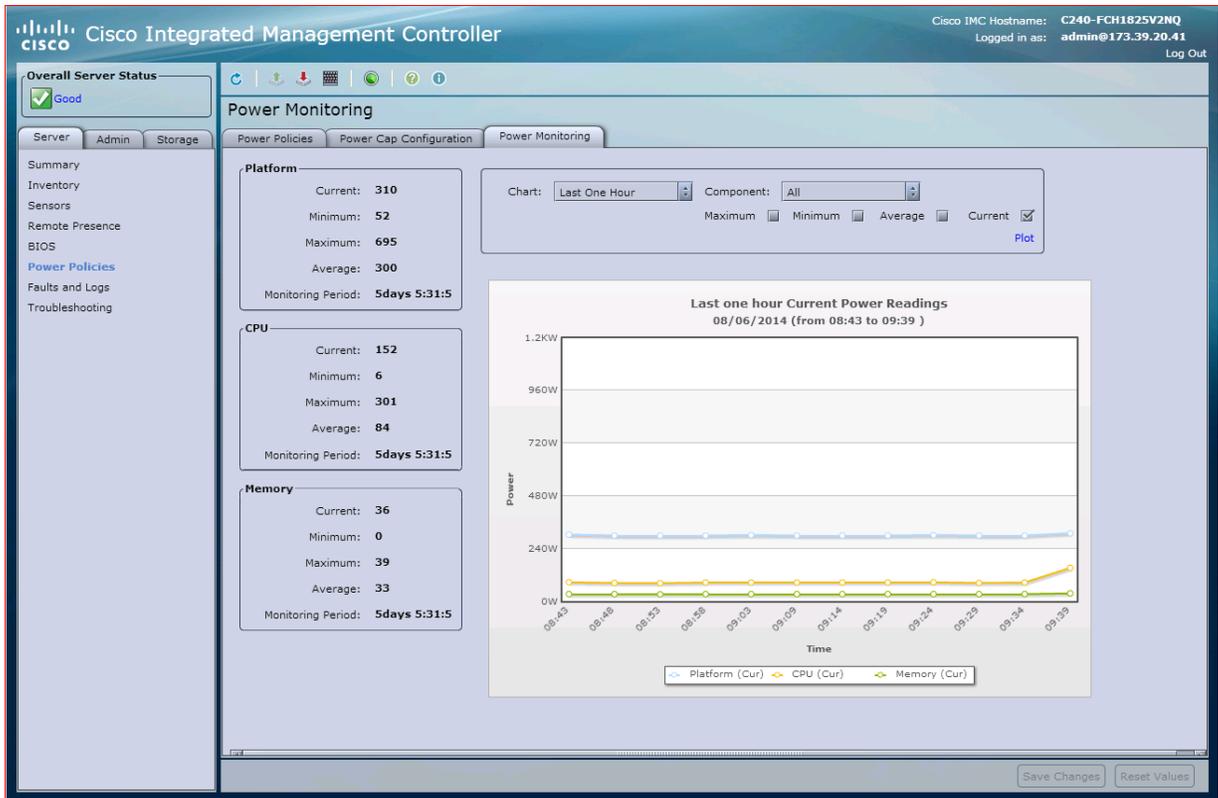
### Creating Power-Efficient Server Configurations

It is advisable to follow some general best practices while selecting the server configurations. One of the best ways to configure a server is to use the **Cisco UCS Power Calculator**.<sup>11</sup> The power calculator will help the user make sure that the server they are provisioning is correctly configured with suitable power supply units (PSUs). Adding high-wattage PSUs to low-power configurations is very inefficient and will be an extremely lossy configuration from a power perspective. Similarly, having low-wattage PSUs for high-power configurations is not useful, as the system will use the maximum possible PSU output power and might even exceed that, leading to throttled operation. Load matching the PSUs with the server configuration is an important first step in power management of rack-mount servers.

### Profiling Data Center Power Usage and Identifying Trends

The Cisco IMC Power Monitoring page in the WebUI and the command-line interface (CLI) scope provide tools for the user to view the power consumption trends of every server (Figure 4).

Figure 4. Power Monitoring in Cisco IMC



As shown in Figure 4, historical power consumption data is collected for the following subsystems:

- Core CPU operations
- Memory operations
- Overall system or platform consumption (this includes the above two subsystems and the unmanaged power consumption of the system)

The same data can be pulled from multiple servers using the XML API scripts.

Example query:

```
<configResolveClass cookie="1407295695/105e5c3d-ffee-1fee-8003-b5a11e1b0ff4" inHierarchical="false" classId="powerMonitor"/>
```

Response:

```
<configResolveClass cookie="1407295695/105e5c3d-ffee-1fee-8003-b5a11e1b0ff4" response="yes" classId="powerMonitor"> <outConfigs> <powerMonitor domain="Platform" current="301" minimum="24" maximum="692" average="384" dn="sys/rack-unit-1/pwrmonitor-Platform" ></powerMonitor><powerMonitor domain="CPU" current="140" minimum="20" maximum="332" average="185" dn="sys/rack-unit-1/pwrmonitor-CPU" ></powerMonitor><powerMonitor domain="Memory" current="20" minimum="0" maximum="22" average="19" dn="sys/rack-unit-1/pwrmonitor-Memory" ></powerMonitor></outConfigs> </configResolveClass>
```

These APIs make it possible for administrators to create script-based tools for data collection across multiple servers in the data center. This information can be collated in any form necessary for further analysis. For example, Figure 5 shows a rendering of the comma-separated values (CSVs) collected using the above API in Microsoft Excel. Data from multiple servers can be superimposed to generate an overall trend for a given set of servers under a particular PDU control or geographically colocated assets in the data center. In some cases, administrators can also profile the power consumption of servers that are hosting the same application.

The above data will give the data center operations manager a clear picture of power consumption trends. With this information, usage tables like Tables 1 and 2 can be generated.

**Table 1.** Power Consumption (in Watts) Averaged Across 10 Servers for Five Business Days

Time	12 a.m. to 6 a.m.	6 a.m. to 9 a.m.	9 a.m. to 12 p.m.	12 p.m. to 5 p.m.	5 p.m. to 9 p.m.	9 p.m. to 12 a.m.
Memory	4	4	6	6	6	4
CPU	90	90	200	225	230	100
System	144	144	270	295	300	152

**Table 2.** Power Consumption (in Watts) Averaged Across 10 servers on Saturday and Sunday

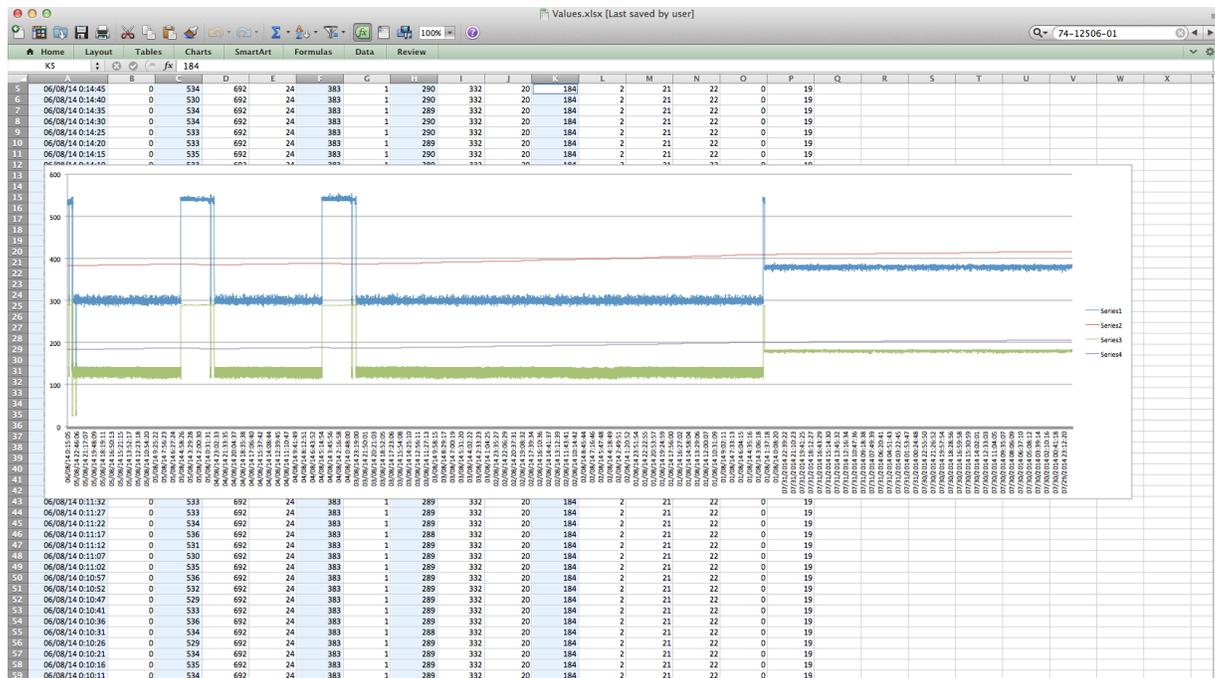
Time	12 a.m. to 6 a.m.	6 a.m. to 9 a.m.	9 a.m. to 12 p.m.	12 p.m. to 5 p.m.	5 p.m. to 9 p.m.	9 p.m. to 12 a.m.
Memory	4	4	4	4	4	4
CPU	90	90	120	110	100	90
System	144	145	178	170	154	144

As can be seen from the tables, the memory modules show very little variance in consumption. It's the CPU and overall system consumption (affected by factors such as cooling) that show the most variance. The maximum consumption happens during business hours on weekdays. Power consumption on Saturday and Sunday is very limited. Note that these distributions are based on averages. There might be individual surges that are not accounted for in the distribution. An effective power cap policy should handle those surges without really affecting the overall IT user needs for performance from the infrastructure.

As an example, using the data in the tables, the operations manager can create a policy for tight power caps on weekends and maximum consumption profiles during weekday business hours. The resulting performance sacrifice on weekends might be acceptable. These policies have to be determined by the administrators after considering the IT needs and requirements of the business unit the data center is servicing.

When generating power consumption and utilization data, it is always preferable to collect them over a group of systems that are running similar applications. This data will help the administrator derive efficient power capping policies while taking care of IT performance needs.

**Figure 5.** Rendering of Power Data over a Period of a Week Using Microsoft Excel

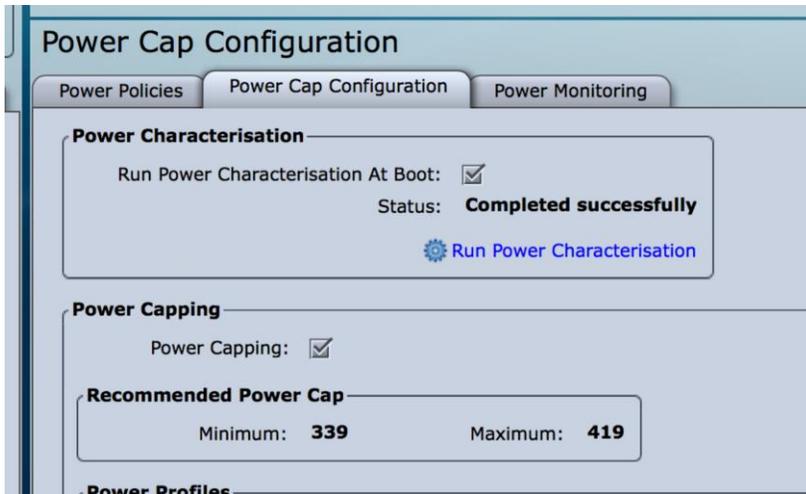


### Characterizing the System Hardware

After developing a power profile, a second important step is to obtain a power characterization of the system hardware as well as important power consumption information that will help the administrator set effective power capping limits. The Cisco IMC triggers special stress software during the host boot process that helps provide the power consumption characteristics of the CPU and memory domains, along with the entire platform. This characterization provides guideline values to the administrator when setting power cap limits for the system. It also implicitly makes sure that the power cap limits take care of the manageable and unmanageable power consumed by the system.

It is recommended that administrators keep the power characterization option enabled for all platforms (Figure 6). When this option is enabled, the system will measure the maximum and minimum power consumption of the CPU, DIMMs, and overall system at every host reboot. Alternately, if the option is kept disabled (which is the default setting), the administrator can explicitly trigger the characterization by using “Run Power Characterization.” If a server has never been characterized, the user receives a message prompting them to run the power characterization before proceeding to the power cap configuration.

**Figure 6.** Power Characterization Option and Output



The Cisco UCS C-Series M4 server power characterization values provided Cisco IMC are highly accurate, as they are derived by actually exercising the power-intensive components of server hardware such as CPUs, DIMMs, and fans. Cisco IMC also has an advanced inventory management system that allows for tracking the Cisco certified field replaceable units (FRUs) in the system, such as hard disk drives and PCIe adapters.

These values are intended to provide recommended guide points to the server administrator when setting the power capping values for different servers. The maximum recommended value indicates the maximum possible power that the given rack server configuration can consume, under any kind of load. Depending on the peripherals and types of components present on the system, this recommendation may vary from a maximum value, which is same as the rating of the power supply unit connected to the system, to a smaller subvalue. For example, a unit with 1200W power supplies may consume only 500W under the maximum possible load if it is equipped with lower-power CPUs and fewer DIMMs and drives. On the other hand, a system having 145W (or higher) CPUs along with graphics processing unit (GPU) cards may consume the full 1200W of the rated power supply output. This maximum power rating is important for efficient provisioning of the power across servers, so that stranded power capacity is not allocated to those servers that are not capable of reaching those limits.

The minimum recommended power provided by the power characterization tool provides the administrator with the minimum power value that the system can consume without drastically affecting system performance. Theoretically, the minimum power consumption of a system can be extremely low, but for a usable platform the value has to be a number that will allow the server to continue functioning, albeit with reduced performance. The administrator can set the system power cap at the minimum power consumption value under special conditions, such as when there is an outage in the primary electrical supply of the data center and it's running on emergency backup power. The critical IT infrastructure is still required to be running until the outage is resolved, so these mission-critical servers can be programmed to run at the minimum power capping values.

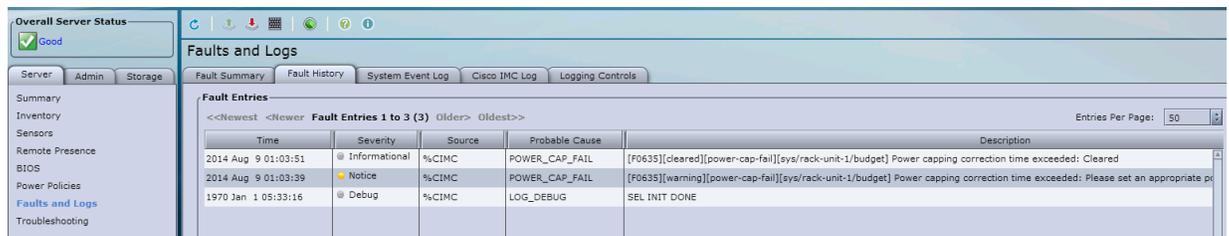
## Understanding and Applying the Findings of Monitoring and Characterization

After analyzing the power monitoring data and running the power characterization tool on the servers, we have some important pieces of information that will help in the provisioning process. We will now use this information to choose either a single policy or a combination of policies for capping the server power consumption, such that the power budget is never at risk of overshooting. To further explain this process, we will use a sample system with a nameplate power consumption of 1200W. The 1200W limitation is based on the two 1200W redundant power supplies connected to the server.

Power characterization of this server has provided a recommended minimum of 350W and a maximum of 725W. This is further validated by the data from the power monitoring of the server, in which the maximum power reading seen over a period of one week is 542W. We can also analyze the pattern of power consumption over time. For example, Table 1 shows a consumption pattern that is distributed over the day in a very specific pattern. The time from 6 a.m. to 9 p.m. appears to be high usage, with CPU and memory consumption at their peak. Performance is important during these hours. The time from 9 p.m. to 12 a.m. seems to have low CPU and memory activity, but I/O activity appears to be higher. This might be due to data backup activities happening during non-business-critical hours. The rest of the time appears to be evenly balanced between I/O and CPU operations. Table 2 indicates that on weekends there is an even distribution of load.

For managing such a workload profile, we can come up with the following scheme. We enable the standard power profile with a capping value of 170W. We also make sure that we select an action to be performed when the power capping fails. Typically this is selected as “Alert,” but the user can also select “Alert and Shut Down.” If the latter is selected, the system will log a fault in the **Fault History** and shut down the host. If the former option is selected, only a fault will be logged so that users can later review the performance of power capping on the server. In some cases, it is advisable to allow throttling of the CPU and memory bandwidth to maintain a safe option to cap power consumption. Enabling the throttling options usually means that the user is preferring to sacrifice system performance in order to achieve a certain level of power consumption. When the system fails to maintain the power cap, and the “Alert” option is selected, they will be see a message like the one in Figure 7 in the Fault History tab of WebUI.

**Figure 7.** Cisco IMC Fault History Showing Power Capping Exception

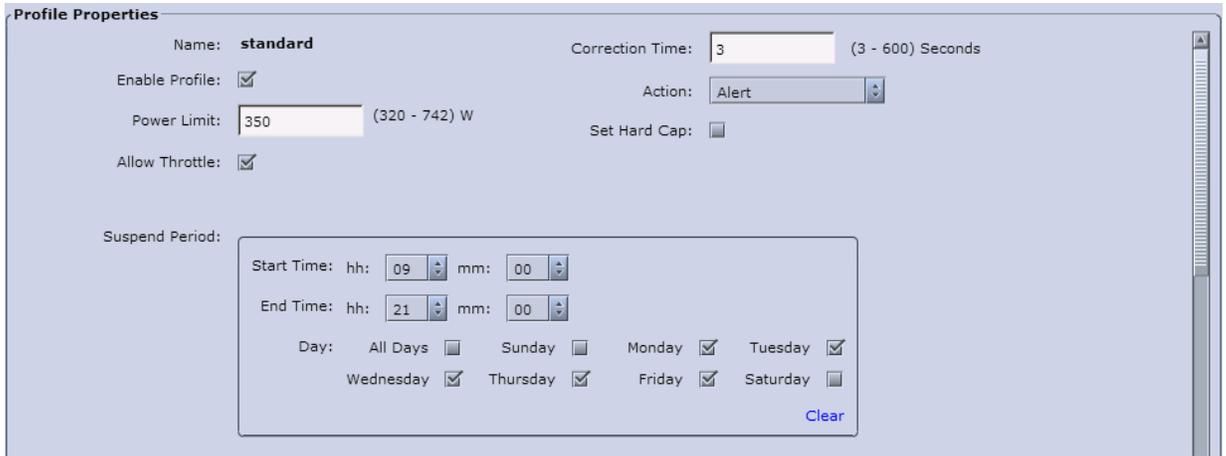


The screenshot shows the Cisco IMC WebUI interface. The top left corner displays 'Overall Server Status' as 'Good'. The main area is titled 'Faults and Logs' and contains a 'Fault History' tab. A table of fault entries is visible, showing three entries:

Time	Severity	Source	Probable Cause	Description
2014 Aug 9 01:03:51	Informational	%CIMC	POWER_CAP_FAIL	[F0635][cleared][power-cap-fail][sys/rack-unit-1/budget] Power capping correction time exceeded: Cleared
2014 Aug 9 01:03:39	Notice	%CIMC	POWER_CAP_FAIL	[F0635][warning][power-cap-fail][sys/rack-unit-1/budget] Power capping correction time exceeded: Please set an appropriate p
1970 Jan 1 05:33:16	Debug	%CIMC	LOG_DEBUG	SEL INIT DONE

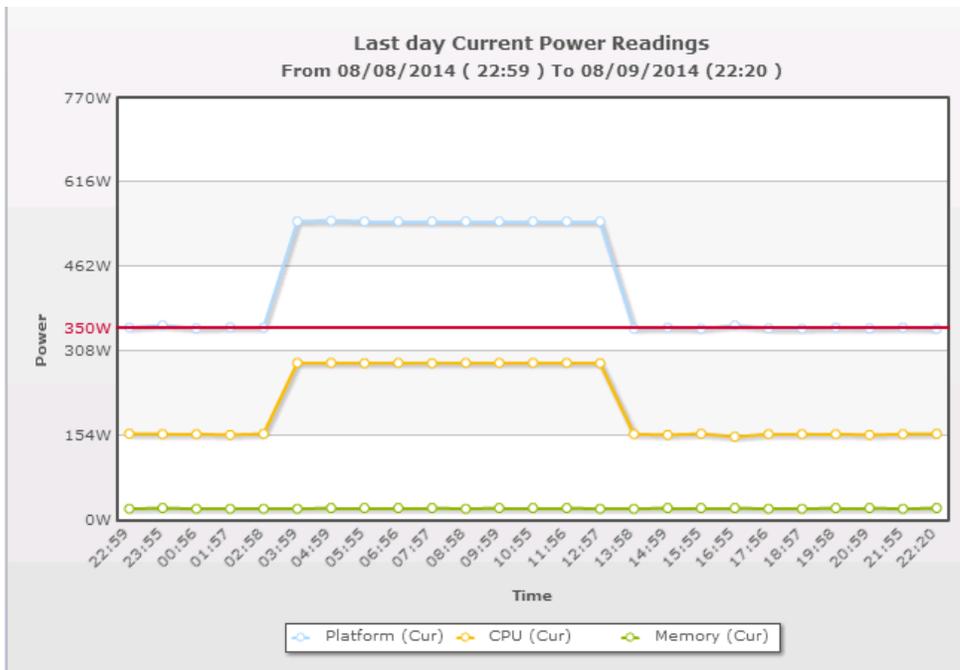
We still need the servers to operate at maximum performance during peak business hours, which appear to be 9 a.m. to 9 p.m. on weekdays. So we select the suspend period for the standard profile as 9:00 to 21:00 for all weekdays (Figure 8).

**Figure 8.** Standard Power Profile Settings



This setting will ensure that system power consumption is not capped during these critical business hours. Figure 9 shows a sample power consumption profile with policy suspend periods enabled. The correction time is the time within which the system will be brought within the power cap setting. A minimum of 3 seconds is recommended for this value.

**Figure 9.** Policy Suspend Periods in Action



In some cases, administrators may want to control the power consumption of the system at a more macro level. In such cases the advanced profiles are useful, as they provide the ability to perform power capping at the CPU, memory, and platform level (Figure 10). This option is more useful for advanced users who have extremely skewed system resources consumption characteristics. Another example is when users prefer memory bandwidth throttling to take precedence over CPU performance reduction. For most users of power capping, the standard profile should suffice.

**Figure 10.** Advanced Power Profile Settings

advanced	No		alert	10	No	Yes
<b>Profile Properties</b> Name: <b>advanced</b> Correction Time: <input type="text" value="10"/> (3 - 600) Seconds Enable Profile: <input type="checkbox"/> Action: <input type="text" value="Alert"/> Allow Throttle: <input checked="" type="checkbox"/> Set Hard Cap: <input type="checkbox"/> <b>Domain Specific Power Limit</b> To enable domain specific power capping specify the individual power limits below. CPU: <input type="text"/> (54 - 222) W      Memory: <input type="text"/> (4 - 9) W      Platform: <input type="text"/> (210 - 532) W						

It is extremely important for the administrator to carefully plan out the deployment of this feature and have a clear understanding of the characteristics of the applications running on the servers. These settings can have a severe impact on application performance. As a best practice, it is recommended that the administrator roll out these settings on a small set of servers in the infrastructure and observe and measure the results before doing a widespread rollout.

### Fail-Safe

The Cisco UCS C-Series M4 power capping solution is an extremely powerful tool for managing the power consumption and performance efficiency of servers in the data center. It is a complex solution that has an impact at the very lowest levels of the hardware. As such, as in any complex system, there can be unforeseen environmental failures. In such cases it is important that there be fail-safe procedures in place. The **Advanced Profile** settings contain multiple options that provide a degree of fail-safe security to the administrator.

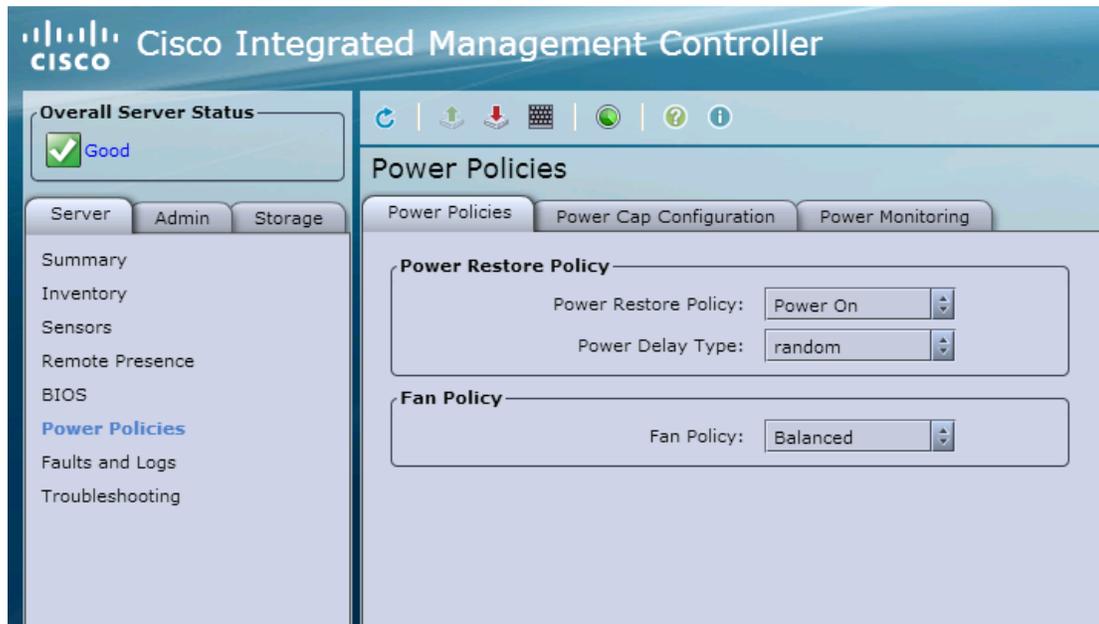
**Safe throttle level:** This option provides a mechanism for the power capping function to fall back to a fail-safe mode when there is an internal fault that prevents its regular functioning. The timeout period in seconds is the time that the system will wait for the internal fault conditions to last before kicking in the fail-safe procedures. The fail-safe procedures involve throttling the memory, CPU, and overall platform operations to some percentage of maximum as specified by the user. These procedures stay in place as long as the fault conditions persist, and the power capping function goes back to normal mode of operation when the source of failure is removed. The internal fault can happen due to catastrophic hardware failure or for some more mundane reason, such as rebooting the Cisco IMC after a firmware update.

**Ambient temperature-based power capping:** This option allows the power capping function to react to catastrophic failures in the server operational environment. If sufficient cooling is not happening in the data center, the system can react to an ambient temperature trigger set by an administrator in this field. The system will immediately throttle the system power consumption to the limit set by the administrator, so that the strained cooling resources of the data center can be relieved.

**Set hard cap:** The hard capping option can be used by the administrator when there is very limited tolerance available in the power budget provisioning. In some cases, the power budget might be extremely tightly subscribed, which means that under no condition should any of the systems in the budget group exceed the capping constraints. In the event of such an exception, the PDU might trip, bringing down the entire infrastructure. In such cases, the hard cap values ensure that the system never exceeds the configured power cap value under any circumstances.

**Power capping at boot time:** Historically, the power consumed by servers when they are booting has been a source of concern for administrators. The Cisco IMC power capping functionality takes effect within 2 minutes of host power-on and ensures that the platform power consumption never exceeds the minimum recommended power cap value. The additional power restore policy in Cisco IMC further allows the administrator to ensure that systems boot up in a staggered fashion (Figure 11). These settings ensure that there is no sudden spike in power consumption when many servers are booted up at the same time.

**Figure 11.** Random Delay Power-On Policy



---

## Evaluating the Performance of the Infrastructure with Power Management

It is important for the data center operations manager as well as the server administrator to understand the effectiveness of the power management policies applied across the servers. The ability of Cisco IMC to report server utilization at the platform level, along with CPU and memory domain level reporting, is extremely useful in this regard. The following sample XML query can be used to monitor server utilization:

```
<configResolveClass cookie="1407579000/0b85c8b0-0030-1030-8002-b5a11e1b0ff4" inHierarchical="false"
classId="serverUtilization"/>
```

The following is a sample response to the XML query:

```
<configResolveClass cookie="1407579000/0b85c8b0-0030-1030-8002-b5a11e1b0ff4" response="yes"
classId="serverUtilization"> <outConfigs> <serverUtilization dn="sys/rack-unit-1/utilization" overallUtilization="86"
cpuUtilization="100" memoryUtilization="68" ioUtilization="0" ></serverUtilization></outConfigs>
</configResolveClass>
```

This data can be collected over a period of time, before and after applying the power management policies. The two sets of data can then be compared to measure the effectiveness of the policies.

---

## For More Information

1. Cisco Unified Computing System™: <http://www.cisco.com/go/ucs>
2. Cisco Unified Computing System (UCS) Manager: <http://www.cisco.com/en/US/products/ps10281/index.html>
3. Cisco on Cisco IT: <http://www.cisco.com/web/about/ciscoitwork/index.html>
4. The Energy-Efficient Data Center—Cisco Systems:  
<http://www.cisco.com/en/US/netsol/ns980/index.html> - ~products
5. Power Management in the Cisco Unified Computing System: An Integrated Approach:  
[http://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/unified-computing/white\\_paper\\_c11-627731.html](http://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/unified-computing/white_paper_c11-627731.html)
6. Understanding Cisco Unified Computing System Service Profiles:  
[http://www.cisco.com/en/US/prod/collateral/ps10265/ps10281/white\\_paper\\_c11-590518.pdf](http://www.cisco.com/en/US/prod/collateral/ps10265/ps10281/white_paper_c11-590518.pdf)
7. Incenting the Right Behaviors in the Data Center, C. Belady and G. Balakrishnan:  
[http://events.energetics.com/datacenters08/pdfs/Belady\\_Microsoft.pdf](http://events.energetics.com/datacenters08/pdfs/Belady_Microsoft.pdf)
8. RAPL OS kernel support: <https://lwn.net/Articles/545745/>
9. Intel® 64 and IA-32 Architectures Software Developer's Manual Combined, Volumes 3A, 3B, and 3C: System Programming Guide: <http://www.intel.com/content/dam/www/public/us/en/documents/manuals/64-ia-32-architectures-software-developer-system-programming-manual-325384.pdf>
10. Intel Power Governor: <https://software.intel.com/en-us/articles/intel-power-governor>
11. Cisco UCS Power Calculator:  
<https://mainstayadvisor.com/Go/Cisco/Cisco-UCS-Power-Calculator-results.aspx>
12. Experimental Study on the Energy Consumption in IaaS Cloud Environments, Alexandra Carpen-Amarie, Anne-Cécile Orgerie, and Christine Morin: <http://hal.inria.fr/hal-00905118>
13. Data Center Power and Cooling: [http://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/unified-computing/white\\_paper\\_c11-680202.pdf](http://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/unified-computing/white_paper_c11-680202.pdf)



---

**Americas Headquarters**  
Cisco Systems, Inc.  
San Jose, CA

**Asia Pacific Headquarters**  
Cisco Systems (USA) Pte. Ltd.  
Singapore

**Europe Headquarters**  
Cisco Systems International BV Amsterdam,  
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at [www.cisco.com/go/offices](http://www.cisco.com/go/offices).

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: [www.cisco.com/go/trademarks](http://www.cisco.com/go/trademarks). Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)