

Cisco Unified Computing System Site Planning Guide: Data Center Power and Cooling



This document provides a technical overview of power, space, and cooling considerations required for successful deployment of IT equipment in the data center. Topics are introduced with a high-level conceptual discussion and then discussed in the context of Cisco[®] products.

The Cisco Unified Computing System[™] (Cisco UCS[®]) product line works with industry-standard rack and power solutions that are generally available for the data center. Cisco also offers racks and power distribution units (PDUs) that have been tested with Cisco UCS and selected Cisco Nexus® products.

This document is intended to inform those tasked with physical deployment of IT equipment in the data center. It does not discuss equipment configuration or deployment from the viewpoint of a system administrator.

Contents

Data Center Thermal Considerations	3
Data Center Temperature and Humidity Guidelines	3
Best Practices	4
Hot-Aisle and Cold-Aisle Layout	5
Populating the Rack	6
Containment Solutions	6
Cable Management	7
Relationship between Heat and Power	8
Energy Savings in Cisco's Facilities	8
Cisco Rack Solutions	8
Cisco Rack Options and Descriptions	9
Multi-Rack Deployment Solutions	9
Data Center Power Considerations	9
Overview	
Power Planning	
Gather the IT Equipment Power Requirements	
Gather the Facility Power and Cooling Parameters	14
Design the PDU Solution	
Cisco RP Series Power Distribution Unit (PDU)	
Cisco RP Series Basic PDUs	
Cisco RP Series Metered Input PDUs	
Cisco RP Series PDU Input Plug Types	
For More Information	17
Appendix: Sample Designs	
Example 1: Medium Deployment (Rack and Blade Server)	
Example 2: Large Deployment (Blade Server)	

Data Center Thermal Considerations

Cooling is a major cost factor in data centers. If cooling is implemented poorly, the power required to cool a data center can match or exceed the power used to run the IT equipment itself. Cooling also is often the limiting factor in data center capacity (heat removal can be a bigger problem than getting power to the equipment).

Data Center Temperature and Humidity Guidelines

The <u>American Society of Heating</u>, <u>Refrigeration</u>, and <u>Air Conditioning Engineers (ASHRAE) Technical Committee</u> <u>9.9</u> has created a widely accepted set of guidelines for optimal temperature and humidity set points in the data center. These guidelines specify both a required and an allowable range of temperature and humidity. ASHRAE 2015 thermal guidelines are presented in the <u>2016 ASHRAE Data Center Power Equipment Thermal Guidelines</u> <u>and Best Practices</u>. Figure 1 illustrates these guidelines.

Figure 1. ASHRAE and NEBS Temperature and Humidity Limits



Although the ASHRAE guidelines define multiple classes with different operating ranges, the recommended operating range is the same for each class. The recommended temperature and humidy are shown in Table 1.

Table 1.	ASHRAE Class A1 to A4 Recommended Temperature and Relative Humidity Range	je
----------	---	----

Property	Recommended Value
Lower limit temperature	64.4°F [18°C]
Upper limit temperature	80.6°F [27°C]
Lower limit humidity	40% relative humidity and 41.9°F (5.5°C) dew point
Upper limit humidity	60% relative humidity and 59°F (15°C) dew point

These temperatures describe the IT equipment inlet air temperature. However, there are several locations in the data center where the environment can be measured and controlled, as shown in Figure 2. These points include:

- Server inlet (point 1)
- Server exhaust (point 2)
- Floor tile supply temperature (point 3)
- Heating, Ventilation, and Air Conditioning (HVAC) unit return air temperature (point 4)
- Computer room air conditioning unit supply temperature (point 5)

Figure 2. Example of a Data Center Air Flow Diagram



Typically, data center HVAC units are controlled based on return air temperature. Setting the HVAC unit return air temperatures to match the ASHRAE requirements will result in very low server inlet temperatures, because HVAC return temperatures are closer to server exhaust temperatures than inlet temperatures.

The lower the air supply temperature in the data center, the greater the cooling costs. In essence, the air conditioning system in the data center is a refrigeration system. The cooling system moves heat generated in the cool data center into the outside ambient environment. The power requirements for cooling a data center depend on the amount of heat being removed (the amount of IT equipment in the data center) and the temperature delta between the data center and the outside air.

The rack arrangement on the data center raised floor can also have a significant impact on cooling-related energy costs and capacity, as summarized in the next section.

Best Practices

Although this document is not intended to be a complete guide for data center design, it presents some basic principles and best practices for data center airflow management.

Hot-Aisle and Cold-Aisle Layout

The hot-aisle and cold-aisle layout in the data center has become standard (Figure 3). Arranging the racks into rows of hot and cold aisles minimizes the mixing of air in the data center. If warm air is allowed to mix with the server inlet air, the air supplied by the air conditioning system must be at an even colder temperature to compensate. As described earlier, lower supply-air temperatures cause increased energy use by the chiller and limit the cooling efficiency of the data center by creating hot spots.

Figure 3. Hot-Aisle and Cold-Aisle Layout



In contrast, not using segregated hot and cold aisles results in server inlet air mixing. Air must be supplied from the floor tile at a lower temperature to meet the server inlet requirements, as shown in Figure 4.



Figure 4. Server Inlet Air Mixing

Populating the Rack

Racks should be populated with the heaviest and most power-dense equipment at the bottom. Placing heavy equipment at the bottom helps lower the rack's center of mass and helps reduce the risk of tipping. Power-dense equipment also tends to draw more air. In the typical data center, in which air is supplied through perforated floor tiles, placing power-dense equipment near the bottom of the rack gives that equipment the best access to the coldest air.

Unoccupied space in the rack can also cause hot air to penetrate back into the cold aisle. Blanking panels are a simple measure that can be used to prevent this problem, as shown in Figure 5.



Figure 5. Using Blanking Panels to Prevent Airflow Short-Circuiting and Bypass

In summary, populate racks from the bottom up and fill any gaps between hardware or at the top of the rack with blanking panels.

Containment Solutions

An effective extension of the hot- and cold-aisle concept is airflow containment. Figure 6 depicts hot-aisle containment. Containment provides complete segregation of the hot and cold air streams, which has the benefit of reducing energy use in the HVAC system by allowing the temperature of the cold air output to be raised. Because there is no mixing of air, there is no need to set the air temperature lower to compensate. This approach increases the temperature of the air returning to the HVAC system, which improves the efficiency of the HVAC system.

For hot-aisle containment, care should be taken to not create pressure in the hot aisle. IT systems are designed so that they have a near zero pressure difference between their air intake and exhaust. Backpressure in the hot aisle can cause fans to work harder in the system.



Figure 6. Hot-Aisle Airflow Containment Example

Cable Management

To the greatest extent possible, airflow obstructions should be removed from the intake and exhaust openings of the equipment mounted in the chassis. Lack of sufficient airflow may result in increased equipment fan power consumption to compensate for increased airflow impedance. If a rack door is installed, it should be perforated and should be at least 65 percent open. Solid doors, made of glass or any other material, inevitably result in airflow problems and should be avoided. Please consult the hardware installation guide for specific equipment requirements.

Proper cable management is critical to reducing airflow blockage. Cisco UCS significantly reduces the number of cables required. However, it is still important to properly dress the cables to provide the best airflow (Figure 7).

Figure 7. Cisco UCS Power and Network Cabling



Relationship between Heat and Power

All power that is consumed by IT equipment is converted to heat. Though power is typically reported in watts (W) and heat is typically reported in British Thermal Units (BTUs) per hour (BTU/hr), these units are in fact interchangeable. Although power is almost always reported in watts, heat load is commonly reported in watts or BTU/hr. The conversion from watts to BTU/hr is 1W = 3.412 BTU/hr. So, for example, a server that consumes 100W produces approximately 341.2 BTU/hr of heat energy.

Energy Savings in Cisco's Facilities

To carefully study the effects of best practices to promote energy efficiency, Cisco underwent a data center efficiency study in the Cisco research and development laboratories. As part of this study, the following best practices were applied:

- Redundant power was disabled where possible
- · Power savings programs were used
- Computational fluid dynamics (CFD) modeling was used
- Virtualization was applied
- Blanking panels were used
- The floor grilles were rearranged
- The chilled water temperature was raised from 44°F to 48°F (7°C to 9°C)

This study demonstrated major improvements in data center power and cooling efficiency. Even though an increase in hardware installations caused the IT load to increase slightly (from 1719 to 1761 kilowatts [kW]), the overhead for cooling the data center dropped (from 801 to 697 kW). The overall power usage effectiveness (PUE) dropped from 1.48 to 1.36. The payback from the proof of concept was 6 to 12 months. The ideas from this pilot project are being applied to all Cisco facilities and are projected to save US\$2 million per year.

Cisco Rack Solutions

The Cisco R42612 Rack is an industry standard EIA-310-D rack that is optimized for the Cisco UCS. It supports all Cisco UCS blade, rack, and storage servers. Standard and expansion racks are available for single-rack or multiple-rack deployments.

The Cisco R42612 Rack provides the highest level of reliability, structural integrity, and security for mission-critical environments. The design of the 42RU, 19-inch rack delivers power, cooling, and cable management features along with the strength and stability required in today's rack enclosures. See Table 2 for specifications.

Features of the Cisco R42612 Rack include the following:

- Front and rear doors are perforated to maximize air flow. Locks are included for added security. The doors also feature tool-less removal for convenient servicing. Split rear doors minimize the clearance required at the rear of the rack, and the adjustable front door can swing from right to left or from left to right.
- Lightweight, two-piece side panels are easy to install and remove, and locks provide extra security.
- The ventilated top canopy has cable access cutouts and a large cable access opening with a brush insert that helps ensure proper airflow.
- Integrated PDU trays enable tool-less mounting for the 0RU Cisco RP Series PDUs.
- Front and back stabilizer brackets are included with the racks.
- A joining kit is available to connect adjacent expansion racks within a row.
- Casters permit mobility if needed.
- RU markings on all rack rails simplify equipment installation.

Cisco Rack Options and Descriptions

Optimizing airflow through the Cisco rack involves eliminating bypass air and helping ensure proper cable management. The Cisco rack solution has options to help. Proper care should be taken to eliminate mixing of hot and cold air. Similarly, proper cable management should be used to reduce airflow impedance.

	Cisco R42612 Rack	Standard	Expansion
	Dimensions (H x W x D)	79.25 x 23.50 x 49.84 in.	79.25 x 23.50 x 49.84 in.
		(2013 x 597 x 1266 mm)	(2013 x 597 x 1266 mm)
	Dimensions (H x W x D) with	84.25 x 32 x 54.84 in.	84.25 x 32 x 54.84 in.
10 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1	раскаділд	(2140 x 813 x 1393 mm)	(2140 x 813 x 1393 mm)
	Distance from front mounting	29.19 in.	29.19 in.
	rail to rear mounting rail	(741.5 mm)	(741.5 mm)
	Weight	339.51 l	264.55 lb
		(154 kg)	(120 kg)
	Weight with packaging	339.51 l	264.55 lb
		(154 kg)	(120 kg)
	Side panels included	Yes	No
	Equipment mounting capacity	42RU	42RU
	Static load capacity	2100 lb (954 kg)	2100 lb (954 kg)
A A	Static load capacity	2700 lb	2700 lb
		(1224.7 kg)	(1224.7 kg)

Table 2. Cisco R42612 Rack Specification

Multirack Deployment Solutions

An optional rack joining kit is available that allows you to attach two or more racks together. The kit includes all required hardware. The doors do not need to be removed to attach the racks together. When racks are joined together, there is no need for the side stabilizers.

Before joining racks together, make sure that the floor can support the weight of all the equipment and the racks themselves.

Please see the Cisco R42612 Rack and RP-Series Metered Input PDU Installation Guide for more details.

Data Center Power Considerations

Power delivery at the facility level requires careful planning. Thorough power planning requires an understanding of server power consumption and redundancy requirements. Poor redundancy planning can result in cascading failures and downtime.

Overview

The proper deployment of power solutions in the data center involves matching the requirements of the IT equipment to the capabilities of the PDU and the facility. Good power planning follows these steps:

Step 1. Gather IT equipment power requirements.

- Power redundancy
- Number of required power receptacles (number of power connections between server PSU and PDU)
- Power capacity
- Step 2. Gather the facility power and cooling parameters. This step involves gathering some basic information about the data center:
 - Input voltage
 - · Single or three-phase power
 - Available plugs
 - · Available cooling
- Step 3. Design the PDU solution. The PDU is what connects the IT equipment and the power source in the data center. After the facility parameters and the IT equipment parameters are defined, the PDU solution can be designed. This step involves making sure the demands of the IT equipment match the capabilities of the facility. It also involves checking that the solution delivers the proper redundancy.

Power Planning

The process described here assumes that you first decide how many servers to install per rack and then design the facility around the rack. Often, however, facility requirements guide the IT solution design. In this case, the process may need to proceed iteratively. For example, if the IT equipment demands cannot be met with the capabilities of the facility, the amount of equipment deployed per rack may need to be reduced.

Gather the IT Equipment Power Requirements

Step 1. Power Redundancy

The first step in sizing a PDU solution is deciding what type of redundancy the end user requires. Clearly, true AC grid redundancy forces requirements on the facility itself because it requires separate and independent AC power sources to be supplied to the racks in the data center. However, redundancy levels also dictate the number of power supplies installed in the equipment itself and hence the number of power plugs required.

In determining the number of PDUs and power sockets required, some background terminology will be helpful:

- Nonredundant power (N): The device being powered is fed power from a single AC power source and is
 equipped with the minimum number of power supplies. Downtime may result from the loss of a power
 supply or the loss of a utility feed.
- N+1 redundancy (N+1): The device being powered receives power from a single AC power source but is
 equipped with at least one redundant power supply. Loss of a utility feed results in downtime, but the device
 can tolerate the loss of at least one power supply.
- Grid redundancy (2N): The device being powered is equipped with twice the required number of power supplies. To maintain true grid redundancy, half the power supplies should derive power from one AC power source, and the other half should be connected to a separate, independent AC power source. 2N redundancy allows the system to tolerate the loss of any power supply or a single AC power source.

In general, full 2N grid redundancy requires double the number of PDUs needed for single-grid power deployment without any power supply redundancy. To have true grid redundancy, each independent AC power source must be provisioned with enough capacity to carry the full load. While both AC power sources are online, the load may be shared between the input sources. Load sharing of the installed devices will cause each feed to appear to have a light load. However, if one of the input power sources fails, the result is an instantaneous spike in power on the remaining source. If each power source is not sized to carry the full expected load, the loss of a redundant feed may cause the remaining feed to become overloaded and result in downtime through cascading failures.

For more information about how to connect power supplies to PDUs for each redundancy option, see the Power Supply Redundancy section in the <u>Cisco UCS 5108 Server Chassis Installation Guide</u>.

Step 2. Number of Required Power Receptacles

In general, the minimum number of required PDU power sockets for a rack-level PDU solution depends on the type of IT equipment and desired level of power redundancy. Please refer to the hardware installation guide for your equipment to obtain the definitive requirements. For convenience, Table 3 lists the quantity and type of power supply connections required for selected Cisco UCS and Cisco Nexus hardware.

		Number of Required PDU Power Receptacles per AC Power Feed		
Model	Power Receptacle	Nonredundant	N+1 Redundant	Grid Redundant ¹
Cisco UCS 5108 Server Chassis	C19	1 to 3 ²	2 to 3 ²	2 to 4 ²
Cisco UCS C220 M3 and M4 Rack Servers	C13	1	2	2
Cisco UCS C240 M3 and M4 Rack Servers	C13	1	2	2
Cisco UCS C460 M4 Rack Server	C13	2	3	4
Cisco UCS S3260 Storage Server	C13	2	3	4
Cisco UCS 6200 Series Fabric Interconnects	C13	1	2	2
Cisco UCS 6300 Series Fabric Interconnects	C13	1	2	2
Cisco Nexus 2000 Series Fabric Extenders	C13	1	2	2
Cisco Nexus 5000 Series Switches	C13	1	2	2
Cisco Nexus 9000 Series Switches	C13 or Saf-D-Grid ³	1	2	2

Table 3. Required AC Input Power Connections

Note: 1: Grid redundancy requires dual utility feeds and in-rack PDUs (see step 1 in Power Planning, above).

Note: 2: In general, a Cisco UCS server chassis requires at least two PSUs for nonredundant operation. Only smaller configurations (requiring less than 2500W) can be powered by a single PSU. Cisco UCS Manager turns on the minimum number of power supplies (PSUs) needed to support the requested power budget and comply with the selected redundancy policy.

Note: Cisco Nexus 9000 Series Switches offer a universal PSU that supports 277 VAC input. The higher input voltage requires a different input connector type, Saf-D-Grid.

Safety is a consideration when selecting a cable. Organizations often ask, "Can we use a cable that connects a C13 receptacle on the PDU to a C20 receptacle on the equipment?" It is tempting to use these cables because PDUs typically have many more C13 than C19 receptacles. However, even though these cables are commercially available, their use is generally discouraged. C19 to C20 power connections are rated for a higher current than C13 to C14 connections. Using a cable that goes from C13 receptacles on the PDU side to C19 receptacles on the equipment side risks either tripping the breaker if the PDU receptacles are fused or, worse, creating a potential fire and safety hazard.

Step 3. Power Capacity

The total load that can be carried by a PDU depends on the rating of the facility input plug. If the actual load exceeds the rating on the input plug for a sufficient period of time, the input breaker will trip, and power will be interrupted to everything that receives power from that plug. To design a PDU solution that helps ensure that power is not interrupted, the load for the equipment on the PDU must be estimated by some means.

There are various ways to estimate the power of an IT equipment deployment in your data center. The approach chosen depends on the goal of the end user. A highly conservative approach, such as allocating power based on faceplate numbers, lowers risk but increases capital costs. The <u>Uptime Institute</u> estimates that 1 kilowatt of server capacity costs US\$25,000 for a Tier IV data center, or US\$11,000 for a Tier 1 data center.

The actual power consumption for a server depends on many factors. First, and most obviously, server power depends heavily on the configuration. Even for similarly configured hardware, power consumption can vary from system to system. Platform configuration settings—for example, BIOS and OSPM options—have an impact on power consumption. Further, the application that is run makes a very big difference in power use. Workloads that do not use the CPU floating-point unit (FPU) use much less power. Enterprise applications rarely, if ever, use the FPU, whereas scientific (high-performance computing [HPC]) workloads are FPU intensive. Given all the variability, any general power number that is used for capacity budgeting must be conservative.

The consequence of overprovisioning power capacity is increased cost. The consequence of underprovisioning power is increased downtime risk. Individual businesses must find an approach that allows them to achieve the appropriate balance of cost and risk.

Several options exist for predicting the power for Cisco UCS servers, as listed here. Each approach to allocating power has benefits and trade-offs. If you are unclear about which of the approaches to employ, Cisco suggests that you begin by using the power calculator approach.

Option 1. Faceplate power allocation: The server faceplate power rating is the maximum rated power for a given system. For example, the Cisco UCS 5108 chassis has up to four 2500W output power supplies, and AC input power is 2790 volt-amperes (VA) at 200 VAC. Full power specifications can be found in the <u>Cisco UCS Site</u> <u>Preparation Guide</u> and <u>Install and Upgrade Guides</u> for Cisco UCS C-Series.

- All four supplies are required to achieve 2N redundancy. In other words, the system is provisioned with double the necessary capacity to be able to tolerate the loss of a utility feed.
- The faceplate power consumption for the Cisco UCS 5108 chassis is therefore 5580 VA.
- Based on faceplate power, the Cisco RP208-30M1P-4-8 or RP208-30M1P-6-36 PDU has enough capacity to support only a single Cisco UCS server chassis. However, it is highly unlikely that the actual power for a Cisco UCS 5108 chassis will ever reach the 5580 VA faceplate number.

Option 2. <u>Cisco UCS Power Calculator</u>: The IT industry has recognized that data center provisioning based on the faceplate is too conservative. Users rarely configure servers so that they are equipped with the maximum CPU quantity and performance, memory, disks, I/O, etc. Power calculators can help by allowing users to enter their configuration to provide a closer estimate of power consumption.

• Although power calculators provide an improved means for estimating power, they can be somewhat conservative. Even for a given system configuration, the power can vary widely, depending on the application.

- For the Cisco UCS Power Calculator, the 100 percent power value represents the highest power consumption that can be achieved on the system. The application used to measure that value uses intense floating-point math and is extremely CPU and memory intensive. Therefore, the 100 percent power value likely is the power value that might be achieved with an HPC application.
- Enterprise Cisco UCS deployments realistically will not approach the 100 percent maximum power value, even if CPU use reaches 100 percent. For example, a database that increases CPU utilization to 100 percent will likely approach only the 50 percent utilization power value indicated by the power calculator.
- Therefore, to estimate the power for a typical enterprise Cisco UCS deployment using the power calculator, Cisco suggests using 50 percent as the system workload factor.

Option 3. Direct power measurement: In estimating power, you cannot get more accurate than direct measurement. Direct measurement will almost certainly yield a lower power value than that determined through the power calculator (even at 50 percent) or faceplate.

- However, provisioning circuit capacity on the basis of this value results in a high risk of nuisance breaker trips. Workloads in the data center evolve over time.
- What represents the maximum realistic power today may be well below the maximum power achievable when new software or workloads are introduced.
- Sophisticated end users who accept the risk of this approach carefully control their workloads and usually factor in an appropriate safety margin.
- Cisco UCS Manager power groups are an excellent way to achieve the economic benefits of lean power
 provisioning while maintaining an operational safety net that recovers stranded capacity (Figure 8). A power
 group is a set of chassis, servers, fabric interconnects, and/or fabric extenders that all draw power from the
 same PDU. Users can assign an explicit power budget for the group. This allows users to safely
 oversubscribe a PDU.
- For data center provisioning, the logical way to use capping is to organize all the servers that draw power from one circuit into a single group. A power cap can then be set for the group that protects the data center from tripping circuit breakers.





Power capping can offer significant benefits, but it can also be misused. If the power cap is set below the
actual maximum power, server performance will be affected. In general, performance reduction is the
mechanism used to cap power. For more information on Cisco group power capping, see the Power
Management section in the white paper <u>Cisco Unified Computing System: An Integrated Approach.</u>

Gather the Facility Power and Cooling Parameters

Before the rack and PDU solution can be designed, you must understand the capabilities of the data center facility as well as the requirements of the equipment. The facility must provide space, power, and cooling:

- Space: Physical space in the rack is quantified in rack units (RU), which are units of vertical space. 1RU of space is 1.75 inches high in the rack; 42RU is the amount of vertical space available in the Cisco R42612 Rack. Space is typically the least limiting factor in IT equipment deployment density. Often, the amount of equipment that can be deployed in the rack is limited by the capability to deliver power or remove heat.
- Power: To quantify the power delivery capability of an AC power circuit in the data center, users must know
 its voltage and maximum usable current. Often the usable current is less than the maximum rated current.
 In the United States, all power circuits must be de-rated by 20 percent to meet the <u>National Fire Protection</u>
 (NFPA) 70: National Electrical Code. For example, in the United States, a 20-ampere (A) circuit is de-rated
 to16A continuous current capacity.

For a single-phase power circuit:

Maximum power = Average Voltage multiplied by Average Current

For a three-phase power circuit:

- Maximum power = Average Phase-to-Phase Voltage multiplied by Average Current and 1.732, where the constant 1.732 results from the fact that not all three phases are producing the same amount of power at the same time. Each phase's voltage and current move through zero at different times.
- The higher power capacity of three-phase circuits enables fewer power circuits for a fixed load. Many PDU vendors offer PDUs that accept three-phase power input and provide single-phase power to the equipment power sockets. Note that, in many data centers, the capability to remove heat tends to be a bigger constraint than the capability to deliver power. For data centers without adequate cooling, increased power delivery capability does nothing to improve rack density.
- Cooling: All power that is delivered to IT equipment is converted to heat. In many data centers, the
 capability to cool the servers is the limiting factor in rack deployment density. Factors such as the use of
 hot-aisle and cold-aisle designs, the use of blanking panels, and the layout of floor tiles affect the actual
 amount of equipment that can be cooled per rack. Most facilities are built with a cooling limit in mind.
 However, since operating practices affect achievable capacity, true cooling capacity may drift over time in
 any given facility. The only way to make any assumptions about the cooling capacity of a given data center
 is to ask someone familiar with the limitations of that facility.

Design the PDU Solution

The final step in the power solution design process is matching the facility capabilities to the IT equipment power requirements. These two parameters were defined in the previous sections.

The design step is best explained through specific examples, provided in the appendix at the end of this document.

Cisco RP Series Power Distribution Unit (PDU)

Cisco RP Series PDUs are an ideal choice for flexible, reliable, easy-to-deploy power distribution with branch circuit protection. In addition, Cisco offers a complete line of computing, rack, and power products that are optimized and certified for Cisco UCS blade and rack servers. The RP Series includes options for metered inputs for remote monitoring to optimize energy use and circuit protection.

Within these types, the PDU offerings are divided between units designed for U.S.-type facility plugs and units designed for international-type plugs that support single and three-phase input. After identifying the desired input, facility plug requirement, monitoring capabilities, desired outputs, and mounting option, the specific Cisco part number can be determined.

Cisco RP Series Basic PDUs

The Cisco RP Series basic PDU models cost-effectively and efficiently distribute single-input power up to 24 outlets. The basic PDU is available in two mounting options: horizontal and vertical. Horizontal PDUs are installed in the rack and take up one or two rack spaces (1RU or 2RU). Horizontal PDUs are ideal for smaller deployments or to minimize power cable congestion created by multiple vertical PDUs. Because of their smaller size, horizontal PDUs have fewer outlets. Vertical (0RU) PDUs can be mounted on the back or side of the rack enclosure. The Cisco R42612 Rack can support up to six vertically mounted PDUs. Vertical PDUs are ideal for large deployments or when rack space is limited. The Cisco RP Series basic PDUs are designed, tested, and approved for use with Cisco solutions. Budget-friendly Cisco PDUs can extend the value of your existing power systems while improving system reliability and availability. Table 4 summarizes the specifications for the Cisco RP Series basic PDUs.

	RP208-30-1P-U-2	RP208-30-1P-U-1	RP230-32-1P-U-2	RP230-32-1P-U-1
RU size	0	1	0	1
Input voltage	200 to 240 VAC ± 10%			
Frequency	50 to 60 Hz			
Amperage	30A	30A	32A	32A
UL rating	24A	24A	32A	32A
Input plug	NEMA L6-30P	NEMA L6-30P	IEC 309 32A 2P+E	IEC 309 32A 2P+E
Cord length	100 ft (3m)	100 ft (3m)	100 ft (3m)	100 ft (3m)
Output voltage	200 to 240 VAC			
Recentacles	20 IEC 320 C13	2 IEC 320 C13	20 IEC 320 C13	2 IEC 320 C13
Receptuoles	4 IEC 320 C19			
Circuit breaker	2 two-pole 20A (UL 489)			

Cisco RP Series Metered Input PDUs

The Cisco RP Series metered input PDUs are an excellent choice for flexible, reliable, easy-to-deploy power distribution with remote monitoring for energy optimization and circuit protection. A metered PDU has the same features as a basic PDU, plus a digital load meter for remote current monitoring to enable load balancing and prevent overloads. These PDUs include flexible mounting options and power levels ranging from 5.7 kW to 22.1 kW. Tables 5 and 6 summarize the specifications for the Cisco RP Series metered input PDUs.

	RP208-30M1P-4-8	RP208-30M1P-6-36	RP230-32M1P-6-36
RU size	1	0	0
Input voltage	200 to 240 VAC ± 10%	200 to 240 VAC ± 10%	200 to 240 VAC ± 10%
Frequency	50 to 60 Hz	50 to 60 Hz	50 to 60 Hz
Amperage	30A	30A	32A
UL rating	24A	24A	32A
Input plug	NEMA L6-30P	NEMA L6-30P	IEC 309 332P6
Cord length	100 ft (3m)	100 ft (3m)	100 ft (3m)
Output voltage	200 to 240 VAC	200 to 240 VAC	200 to 240 VAC
Receptacles	8 IEC 320 C13 4 IEC 320 C19	36 IEC 320 C13 6 IEC 320 C19	36 IEC 320 C13 6 IEC 320 C19
Circuit breaker	2 two-pole 20A (UL 489)	2 two-pole 20A (UL 489)	2 one-pole 16A

 Table 5.
 Specifications for Single-Phase RP Series Metered Input PDUs

Table 6. Specifications for Three-Phase RP Series Metered Input	PDUs
---	------

	RP208-60M3P-12	RP208-30M3P-6-30	RP230-60M3P-12-9	RP230-32M3P-12-12
RU size	2	0	0	0
Input voltage	200 to 240 VAC ± 10%	200 to 240 VAC ± 10%	200 to 240 VAC ± 10%	346 to 450 VAC ± 10%
Frequency	50 to 60 Hz	50 to 60 Hz	50 to 60 Hz	50 to 60 Hz
Amperage	60A	30A	60A	32A
UL rating	48A	24A	48A	32A
Input plug	IEC 309 460P9	NEMA L15-30P	IEC 309 460P9	IEC 309 532P6
Cord Length	100 ft/3m	100 ft/3m	100 ft/3m	100 ft/3m
Output Voltage	200 to 240 VAC	200 to 240 VAC	200 to 240 VAC	200 to 240 VAC
Receptacles	12 IEC 320 C19	30 IEC 320 C13 6 IEC 320 C19	9 IEC 320 C13 12 IEC 320 C19	12 IEC 320 C13 12 IEC 320 C19
Circuit Breaker	6 two-pole 20A (UL 489)	3 two-pole 20A (UL 489)	6 two-pole 20A (UL 489)	6 one-pole 16A

Cisco RP Series PDU Input Plug Types

Cisco RP Series PDUs support a wide range of input voltage types. The input voltage, rated current capacity, and region (North America or international) determine which input plug is required. The maximum PDU power for international versions is higher because (1) the input plug is rated for higher current and (2) de-rating requirements defined by the <u>National Fire Protection (NFPA) 70: National Electrical Code</u> do not apply outside of North America. Table 7 summarizes the capabilities of the various input plugs.

Table 7.Input Plug Details

	North America			International	
Plug type	L6-30P	L15-30P	IEC309 460P9	IEC309 332P6	IEC309 532P6
Plug image			0	())	()))) ())
Nominal input voltage	208V	208V	208V	230V	400V
Maximum usable current	24A	24A	48A	32A	32A
Maximum total power	5.0kW	8.6kW	17.3kW	7.4kW	22.2kW
Supported models	RP208-30-1P-U-1 RP208-30-1P-U-2 RP208-30M1P-4-8 PR208-30M1P-6-36	RP208-30M3P-6-30	RP208-60M3P-12 RP208-60M3P-12-9	RP230-32-1P-U-1 RP230-32-1P-U-2 PR230-32M1P-6-36	PR230-32M1P-12-12

For More Information

Within Cisco

Cisco R42612 Rack and Cisco RP Series Metered Input Power Distribution Units

Cisco R-Series Rack and RP-Series PDU Installation Guide

Cisco RP Series Basic Power Distribution Units Data Sheet

Cisco RP Series Metered Input Power Distribution Units Data Sheet

Power Management in the Cisco Unified Computing System: An Integrated Approach White Paper

External

One excellent resource for information and best practices is <u>the Green Grid</u>. The Green Grid comprises equipment manufacturers, data center design firms, and end users and has a wealth of in-depth information regarding IT energy efficiency standards and practices.

Also, <u>ASHRAE (American Society of Heating, Refrigeration and Air-Conditioning)</u> has provided excellent data center facility guidelines through <u>Technical Committee 9.9 (TC 9.9)</u>.

Appendix: Sample Designs

Example 1: Medium Deployment (Rack and Blade Server)

The first step involves (1) determining the power redundancy requirement based on availability requirements, (2) determining the number of receptacles based on solution hardware requirements and redundancy requirements, and (3) determining the power capacity requirements based on hardware configuration. For this example, the solution includes a pair of fabric interconnects, blade servers, and a rack server. The specific hardware configuration is described in Table 8.

Quantity	Base SKU	Configuration Details
2	UCS-FI-6332	2x 650W Power Supply (UCS-PSU-6332-AC)
1	UCSB-5108-AC2	4x 2500W Power Supply (UCSB-PSU-2500ACDV) 2x 2304 IOM (UCS-IOM-2304)
8	UCSB-B200-M4	2x E5-2683 v4 (UCS-CPU-E52683E) 16x 32GB DDR4-2400MHz RDIMM (UCS-MR-1X322RV-A) 2x 300GB 6Gb SAS15K RPM SFF HDD (UCS-HDD300GI2F105) 1x VIC 1340 mLOM (UCSB-MLOM-40G-03)
1	UCSC-C220-M4L	2x 770W Power Supply (UCSC-PSU1-770W) 2x E5-2698 v3 (UCS-CPU-E52698D) 24x 8 GB DDR4-2133MHz RDIMM (UCS-MR-1X081RU-A) 4x 600 GB 6 Gb SAS 15K RPM LFF HDD (UCS-HD600G15KHY-E) 1x VIC 1225 (UCSC-PCIE-CSC-02

 Table 8.
 Example 1 Solution Configuration

For this example, the availability requirement is grid redundant. This means two redundant AC line feeds are required. During normal operation, the equipment load is shared across both feeds, where a single feed is capable of supporting the entire equipment load during a planned or unplanned loss of a single feed. The number and type of receptacles is documented in Table 9.

Table 9. Example 1 Receptacle Requirem	ents
--	------

Quantity	Base SKU	Total Inputs by Receptacle Type	Receptacle Type per AC Input Feed
2	UCS-FI-6332	0 C19 + 4 C13	0 C19 + 2 C13
1	UCSB-5108-AC2	4 C19 + 0 C13	2 C19 + 0 C13
1	UCSC-C220-M4L	0 C19 + 2 C13	0 C19 + 1 C13

Again, these requirements are based on a redundant power policy. The number of plugs required depends on the desired redundancy and expected power consumption. In general, a Cisco UCS 5108 Blade Chassis with a grid redundancy power policy requires two PSUs wired to each AC input feed. If a different redundancy policy (N+1 or nonredundant) is selected, fewer PSUs and fewer inputs may be required.

The next step is to determine how much power is required to support the solution. The <u>Cisco UCS Power</u> <u>Calculator</u> is used to estimate the expected power consumption. The estimated power consumption ranges from 1500W while idle to 4900W while executing the heaviest possible workload under the worst case operating conditions. See Figure 9 for more details. The estimate for typical power consumption is 3200W, and is based on a 50 percent system workload factor. In general, a system workload factor of 50 percent is recommended for sizing the power required by the solution.

UCS Power	Calculator	Ŵ	How to Order 🛛 🛱 Put	rchase 🏼 🄞 Sr	martPlays	😪 Lo	ocal Reseller
					Projects	Edit	Export 👻
Results							
	System Workload					Cooling	Sustam
	Factor (W)	Max Power (W)	Max Current Draw (A)	Idle Power (V	V) ,	Workflow	(BTU/hr)
Fabric	465.2	544.2	2.76	387.6		158	7.2
B-Series Servers	2373	3825.8	18.5	960.5		809	6.9
C-Series Servers	320	541.7	2.62	101.6		109	1.9
Invicta Systems	0	0	0	0		0	(
M-Series Servers	0	0	0	0		0	(
Total	3158.1	4911.7	23.89	1449.8		107	76
Annual Cost	\$2,904.85	\$4,517.75		\$1,333.53			

Figure 9. Example 1 Power Consumption Estimate

The next step is to determine available power and cooling capacity supported by the data center. This example uses a typical single-phase facility receptacle for the United States. Each rack is allocated two L6-30R receptacles, where each receptacle is connected to a different AC power distribution circuit to support grid redundancy. Each feed is capable of providing 4992W of continuous power. The cooling solution is sized to match the available power per rack. Therefore, each rack is sized to provide approximately 5 kW of continuous power and cooling for installed IT equipment.

The amount of rack space required by the solution is described in Table 10. The R42612 rack provides 42RU of usable rack space.

Table 10.	Example 1	Rack S	pace Red	quirements

Quantity	Base SKU	Total RU
2	UCS-FI-6332	2RU
1	UCSB-5108-AC2	6RU
1	UCSC-C220-M4L	1RU

Note: 33RU of physical space remains for PDUs and/or other equipment. Space that is not consumed should be populated with filler panels.

The final step is to select the appropriate rack PDU. The Cisco RP Series includes four models with an L6-30P. The maximum estimated power consumption for the solution is 4400W. Each L6-30 plug is capable of providing 4992W of continuous power. Therefore, a single Cisco RP Series PDU is more than capable of supporting the solution. A second Cisco RP Series PDU is required to support the redundancy requirements. For the example, the customer plans to monitor input power consumption using Cisco UCS Manager. Therefore, a basic model is selected for this example. There is sufficient space in the Cisco R42612 Rack to support a horizontal or vertical mounting option. The horizontal rack PDU (RP208-30-1P-U-1) has six total outlets (two IEC 320 C13 and four IEC 320 C19). Although there are enough outlets, there are not enough IEC 320 C13 outlets.

In this case, an IEC 320 C19 outlet can be used; however, a special adapter cable (IEC 320 C13 to IEC 320 C20) is required. This type of adapter cable is not available from Cisco and must be sourced through a partner.

For this example, the vertical mounting model (RP208-30-1P-U-2) is selected. This model has sufficient outlets as well as additional outlets to support additional equipment deployed in the rack.

Example 2: Large Deployment (Blade Server)

The first step involves (1) determining the power redundancy requirement based on availability requirements, (2) determining the number of receptacles based on solution hardware requirements and redundancy requirements, and (3) determining the power capacity requirements based on the hardware configuration. For this example, the solution includes a pair of fabric interconnects, blade servers, and a rack server. The specific hardware configuration is described in Table 11.

Quantity	Base SKU	Configuration Details
2	UCS-FI-6332	2x 650W Power Supply (UCS-PSU-6332-AC)
1	UCSB-5108-AC2	4x 2500W Power Supply (UCSB-PSU-2500ACDV) 2x 2304 IOM (UCS-IOM-2304)
8	UCSB-B200-M4	2x E5-2683 v4 (UCS-CPU-E52683E) 16x 32 GB DDR4-2400MHz RDIMM (UCS-MR-1X322RV-A) 2x 300 GB 6 Gb SAS15K RPM SFF HDD (UCS-HDD300GI2F105) 1x VIC 1340 mLOM (UCSB-MLOM-40G-03)

 Table 11.
 Example 2 Solution Configuration

For this example, the availability requirement is grid redundant. This means two redundant AC line feeds are required. During normal operation, the equipment load is shared across both feeds, where a single feed is capable of supporting the entire equipment load during a planned or unplanned loss of a single feed. The number and type of receptacles is documented in Table 12.

Table 12. Example 2 Receptacle Requirements

Quantity	Base SKU	Total Inputs by Receptacle Type	Receptacle Type per AC Input Feed
2	UCS-FI-6332	0 C19 + 4 C13	0 C19 + 2 C13
1	UCSB-5108-AC2	24 C19 + 0 C13	12 C19 + 0 C13

Again, these requirements are based on a redundant power policy. The number of plugs required depends on the desired redundancy and expected power consumption. In general, a Cisco UCS 5108 Blade Chassis with a grid redundancy power policy requires two PSUs wired to each AC input feed. If a different redundancy policy (N+1 or nonredundant) is selected, fewer PSUs and fewer inputs may be required.

The next step is to determine how much power is required to support the solution. The <u>Cisco UCS Power</u> <u>Calculator</u> is used to estimate the expected power consumption. The estimated power consumption ranges from 6150W while idle to 23,500W while executing the heaviest possible workload under the worst case operating conditions. See Figure 10 for more details. The estimate for typical power consumption is 14,700W, and is based on a 50 percent system workload factor. In general, a system workload factor of 50 percent is recommended for sizing the power required by the solution.

UCS Pow	ver Calculator	Ē	How to Order 🛛 🛱 Pure	chase 🍈 S	SmartPlays	<u>6</u> L	ocal Resell
					Projects	Edit	Export
Results							
	System Workload Factor (W)	Max Power (W)	Max Current Draw (A)	Idle Power	(W)	Cooling Workflow	System (BTU/hr)
Fabric	465.2	544.2	2.76	387.6		158	37.2
B-Series Servers	14237.8	22954.6	111	5763.3		485	81.6
C-Series Servers	0	0	0	0		(0
Invicta Systems	0	0	0	0		(0
M-Series Servers	0	0	0	0		(0
Total	14703	23498.8	113.76	6150.9		501	68.7
Annual Cost	\$13,523.83	\$21,614.17		\$5,657.5	9		

Figure 10. Example 2 Power Consumption Estimate

The next step is to determine available power and cooling capacity supported by the data center. This example uses a typical single-phase facility receptacle for the United States. Each rack is allocated two 460R9 receptacles, where each receptacle is connected to a different AC power distribution circuit to support grid redundancy. Each feed is capable of providing 17.3W of continuous power. The cooling solution is sized to match the available power per rack. Therefore, each rack is sized to provide approximately 17.3 kW of continuous power and cooling for installed IT equipment.

The amount of rack space required by the solution is described in Table 13. The R42612 rack provides 42RU of usable rack space.

Quantity	Base SKU	Total RU
2	UCS-FI-6332	2RU
6	UCSB-5108-AC2	36RU

Note: 4RU of physical space remains for PDUs and/or other equipment. Space that is not consumed should be populated with filler panels.

The final step is to select the appropriate rack PDU. The Cisco RP Series includes two models with a 460P9 plug. The maximum estimated power consumption for the solution is 22.5 kW. Each 460P9 plug is capable of providing 17.3 kW of continuous power. Although the maximum power consumption estimate is higher than 17.3 kW, it is highly unlikely that the solution will reach the maximum power estimate. System workloads are dynamic in nature, and it is highly unlikely that each chassis will simultaneously consume maximum power. For this reason, the use of a system workload factor is recommended for sizing. The estimate for a system workload factor of 50 percent is 14.7 kW. This is well within the rated capacity of a single 460P9 plug. To ensure that the solution does not exceed the rated capacity, a power group with an explicit power budget of 17.3 kW is created in Cisco UCS Manager. See the Cisco UCS Manager Configuration Guide for creating Power Groups in Cisco UCS Manager.

A second Cisco RP Series PDU is required to support the redundancy requirements. There is sufficient space in the Cisco R42612 Rack to support a horizontal or vertical mounting option. The horizontal rack PDU (RP208-60M3P-12) does not have a sufficient number of outlets. Therefore, the vertical mounting model (RP230-60M3P-12-9) is selected.



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.

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. 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)

Printed in USA