How to Choose IT Rack Power Distribution

White Paper 202

Revision 0

by Brian Mitchell

Executive summary

One of the challenges associated with rack power distribution units (PDU) has been determining which to choose among the wide array of offerings. In most cases, there are so many to choose from (100-700 models) that vendors must provide product selectors in order to narrow down the choices. Others challenges include maintaining system availability and supporting higher density equipment. Once a rack PDU is selected, IT administrators wonder if it will support the next generation of IT equipment, in terms of power capacity, electrical plug type(s), and plug quantity. Trends such as virtualization, converged infrastructure, and high efficiency add to the need for a comprehensive strategy for selecting rack PDUs. This paper discusses the criteria for selecting IT rack power distribution and the practical decisions required to reduce downtime.

APC by Schneider Electric White Papers are now part of the Schneider Electric white paper library produced by Schneider Electric's **Data Center Science Center**
DCSC@Schneider-Electric.com

Introduction

Rack power distribution units (PDU) are available with a variety of different features, power ratings, and input and output cord combinations. Selecting the right rack PDU for the application can sometimes be difficult, especially without exact information about the equipment that will go into a rack. Trends such as virtualization, converged infrastructure (compute, storage, and networking), and high efficiency add to the need for a comprehensive strategy for managing equipment in IT racks. This paper discusses the criteria for selecting IT rack power distribution and the practical decisions required to reduce downtime.

Determining which rack PDU to use in a particular rack requires information about the equipment in the rack, the site power distribution, and preference for additional functionality. Most commonly, rack PDUs are selected following an “inside out” order:

1. Determine output plug type and quantity
2. Estimate power capacity
3. Determine input plug type
4. Select visibility and control options for branch circuits
5. Select form factor and mounting

Step 1: Determine output plug type and quantity

With the “inside out” approach, the first decisions are made based on the application inside the rack. IT equipment in racks can have several different plug types. The most common plug types in data centers are C-13 and C-19 connectors (see **Figure 1**) as defined by IEC 60320. C-13 connectors are usually found on servers and small switches, while blades and larger networking equipment use the C-19 plug because of its higher current carrying capacity. Nearly any new server, blade chassis, or enterprise-level switch can use either C-13 or C-19 inputs depending on power consumption. Most equipment that cannot use either a C-13 or C-19 cord falls into three categories: Legacy equipment, equipment drawing 30A or more from a single cord, or equipment targeted for other applications including air conditioners, fans, and laptop power packs that have a regional power connector (e.g. NEMA 5-15).

Figure 1

C-13 and C-19 outlets are the most commonly used for IT equipment

(C-13) (C-19)

Selecting the proper combination of outlets for a rack PDU starts with a look at the IT equipment that will operate in the rack. Rack PDUs should have at least as many plugs of each type as the equipment inside the rack so that every piece of equipment can be connected.

Many data centers operators choose rack PDUs with more outlets of each type than are needed for the initial load. This allows additional outlets for future equipment. **Table 1** provides the number of outlets generally required for different mixes of IT equipment. A popular outlet combination from several manufacturers has been (36) C-13 and (6) C-19 outlets because it allows for a mix of either high or low density equipment. From a data center perspective, this means that a single, common rack PDU can serve practically any rack on the floor. Selecting a popular outlet configuration helps to ensure that the rack PDU will be in stock and immediately available whenever it is needed.

Table 1

Typical power density and outlet ranges per rack

IT equipment	Density range	Power cord qty per feed	Outlet type
1U servers	2 to 9 kW	21 – 42	IEC C13, NEMA 5-15R
2U/4U servers, mixed environment	3 to 10 kW	10 – 21	IEC C13, NEMA 5-15R
Blade servers	6 to 20kW	6 – 12	IEC C19
High density networking	6 to 10 kW	2 – 6	IEC C19
Deep, high density networking	6 to 25+ kW	3 – 6	IEC C19
Network storage	4 to 8 kW	10 – 21	IEC C13, NEMA 5-15R
Converged IT infrastructure	6 to 8 kW	10 – 21	IEC C13, IEC C19

Recommendation: A combination of (36) C-13 and (6) C-19 outlets is common and provides sufficient outlet capacity for nearly all rack PDU applications.

Step 2: Estimate power capacity

Several methods can be used to estimate the maximum power required per IT rack. Once the power requirements are established, a rack PDU can be selected that provides at least enough power to support the rack's load.

1. Estimate the power usage of the equipment inside the IT rack. This method is commonly used when very high power loads are used such as enterprise class servers, blade servers, or high speed networking. The nameplate rating on this equipment is typically higher than actual usage because it is based on the power draw of the power supply at full load. Servers are generally not configured to draw the full power of the supply, so this approach is considered to be fairly conservative.
2. Several manufacturers including Cisco, Dell and HP offer online calculators that can more closely estimate the actual power draw accounting for the specific configurations (number and type of cards, drives, etc.). Schneider Electric offers a sizing tool to help estimate a realistic power usage based on specific models and configurations. **Table 1** provides guidance for typical rack power consumption based on the IT equipment used within the rack.
3. Assume a maximum power level of the IT rack based on an estimate of total data center utilization. For example, if the data center delivers 1 MW of power to the IT load, and the IT load consists of 100 racks, then a maximum power level could be estimated to be around 10 kW for the vast majority of the racks. This is easier to estimate and implement than calculating a maximum for each rack individually. This method is commonly used when a heterogeneous computing environment is expected and the specific IT equipment will be difficult to predict. The user manages the environment by restricting installation of additional servers into IT racks approaching the maximum power level. Schneider Electric's TradeOff Tool™ Power Sizing Calculator, provides guidance on sizing data center power requirements.

Fault current rating

Available fault current, measured in kiloamps (kA), is the maximum amount of current available to "feed" a fault and is dependent on the electrical system design.

The fault current at the input of rack PDUs should be limited to 10kA because this is the typical rating for most cord caps (i.e. input connector plug). For more information on this topic, see White Paper 194, *Arc Flash Considerations for Data Center IT Space*.

Step 3: Determine input plug type

With considerations inside the rack complete, decisions can now be made about the overall row power distribution. The site voltage (i.e. 208V, 240V, etc) should be determined. Additionally, a decision must be made about delivering either single phase or 3-phase power to the rack. To understand the benefits of 3-phase power, see Geist White Paper EP901 Three-Phase Electric Power Distribution for Computer Data Centers.

The site power distribution scheme (voltage and phase) and the estimated power requirements will dictate the input power cord or whip amperage needed for the IT rack. See **Tables 2** and **3** to determine the whip size. The tables show the continuous power capacity in kW at each amperage. Grey shaded values represent uncommon power levels that are generally not recommended by design engineers.

Table 2

ANSI continuous power consumption based on input voltage and amperage

Whip Amps	Single/Split Phase			3-Phase	
	120V	208V	240V	208V	415V
20A	1.9kw	3.3kw	3.8kw	5.8kw	11.5kw
30A	2.9kw	5.0kw	5.8kw	8.6kw	17.3kw
40A	3.8kw	6.7kw	7.7kw	11.5kw	23.0kw
50A	4.8kw	8.3kw	9.6kw	14.4kw	28.8kw
60A	5.8kw	10.0kw	11.5kw	17.3kw	34.5kw

Table 3

IEC continuous power consumption based on input voltage and amperage

Whip Amps	Single Phase	3-Phase
	230V	400V
16A	3.7kw	11.1kw
32A	7.4kw	22.2kw

To mitigate risk of overloads that might impact uptime, common practice is to limit PDUs to ≤40A for single phase or ≤32A for 3-phase applications. Rack PDUs larger than this are seen to expose too much equipment at risk of a single circuit breaker overload. If more power is needed within a single rack, an additional set of PDUs on separate circuits should be installed. Best practice is to use ≤20A rack PDUs when possible because these PDUs do not contain circuit breakers. This eliminates a component from the power path and has reliability benefits. The circuit is still protected by an upstream 20A circuit breaker as always. A good practice for 120V countries is to switch to 240V, similar to 230V countries. This higher voltage provides double the amount of power in a 20A circuit compared to 120V. This topic is discussed in White Paper 128, Switching to 240V AC Distribution in North American Data Centers.

Typically, rack PDUs have locking NEMA (North America) or IEC 309 (EMEA) plugs to connect to the power distribution. In certain cases it is desirable to have waterproof connectors because connections are made under a raised floor which may become flooded with water in the event of a fire, although this is uncommon.

Rack PDUs are commonly available in several different whip lengths. An appropriate length should be selected that allows at least enough slack that the cord can be comfortably plugged into the power outlet and managed if maintenance is ever needed. A whip length of 2m (6') is typically sufficient to reach input power.

Recommendation: Use 16A and 20A rack PDUs when possible because they are easy to work with, enhance reliability, and simplify layouts. Provisioning a rack with 3-phase rack PDUs offers additional power capacity for future loads.

Step 4: Select visibility and control options for branch circuits

Rack PDUs can also monitor power and provide visibility into the instantaneous power consumption in addition to overall power consumption trends over an extended period of time.

There are three levels of visibility into rack-level power status: basic, monitoring, and metering. A **basic** view provides no information about power consumption at all. In this view, the rack may be very close to tripping a breaker, but nobody is able to identify this as a problem.

A **monitoring** view provides a local screen that displays instantaneous information only. This can be useful in determining phase balancing and general circuit load status, but provides insight only into one instant. Decisions driven by this information will not take into account peak usage or cyclical trends. These devices cannot signal outside of the local display, and cannot alert IT staff about high-risk situations as they happen.

For datacenter applications, a **metering** view allows visibility when circuits approach the maximum capacity and when breakers are at risk for tripping. Through the use of a built-in network management card, metered PDUs can alert users when loads approach predetermined thresholds through email, text message, visual alerts on displays, and through other means *before* problems arise. IT staff can use stored power consumption history to analyze trends and make more informed decisions about where to add new devices so circuits can be balanced and never at risk of overload. This type of proactive planning is a solid approach to eliminating downtime from tripped breakers.

Some rack PDUs offer metering of individual outlets instead of just branch circuits. While most data centers are not prepared to take advantage of this deeper look into power consumption in the rack, it empowers IT professionals with the necessary tools for advanced energy management.

Many rack PDUs are also able to measure other dimensions of the operating environment. Commonly, temperature and humidity sensors can connect to metered rack PDUs. PDUs can then display that information on the local display and store or transmit the data across the network. After power availability, server temperature is the most likely cause of equipment failure and downtime. In-rack temperature sensors connected to rack PDUs are an easy way to ensure a proper operating environment for servers and other equipment.

Recommendation: Use metered rack PDUs if there is no dedicated up-stream metering of the rack PDU's branch circuit.

Rack PDUs can also provide control as well as visibility. **Switching** functionality allows outlets to be cycled on/off remotely via a network connection. This is common for remote sites. Servers or switches occasionally need physical reboots, even in an office or facility that does not have on-site IT staff. A common way to do this involves calling the local administrator. The administrator walks into the unfamiliar network closet and looks for cues to identify the troubled device in order to restart it. Even with good intentions, the wrong equipment is sometimes rebooted. Switched rack PDUs allow knowledgeable IT staff to manage equipment cycling even if they are off site.

With colocated equipment, one of the most common calls is to request manual restarts of hung servers. In order to do this, the colo admin must walk to the rack, identify the correct server, restart it, and then verify with the customer that the task is completed successfully. This takes time, and there is also typically a fee (generally $100-$200/hour) associated with any work done by the colo to manage customers' equipment. This lag time, risk, and expense can be eliminated with rack PDUs that can switch outlets remotely. **Figure 2** shows an example of control interface for controlling rack PDU outlets.

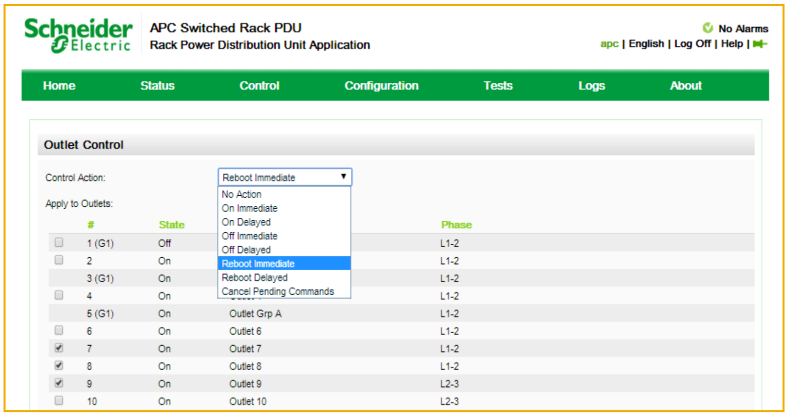

Figure 2

In a colocation environment, rebooting servers can be managed remotely by cycling outlets with a switched rack PDU

Accuracy of metering equipment

Like all measurement devices, metered rack PDUs have a tolerance band of meter accuracy. Tolerances for the Hall Effect current transducers (CTs) typically used in rack PDUs are affected by operating temperature and current load. As the ambient temperature increases, so does the CT's tolerance as illustrated in **Figure 3**. High rack densities and rack air containment strategies tend to drive hot aisle temperatures up to 32ºC (90ºF) or higher. The data center operating environment provides some practical bounds for operating temperature ranges that rack are likely to encounter. Cold aisle temperatures are unlikely to be lower than 15ºC (60ºF), and server exhaust temperatures are unlikely to rise above 45ºC (113ºF). Above this range, server reliability begins to become a concern as electronic components are not intended to operate continuously in elevated temperatures. Note that rack PDUs are typically mounted at the rear of an IT rack which can experience temperatures as high as 45ºC (113 ºF) in contained hot aisles. Ensure the rack PDU is rated for continuous use in high temperature applications.

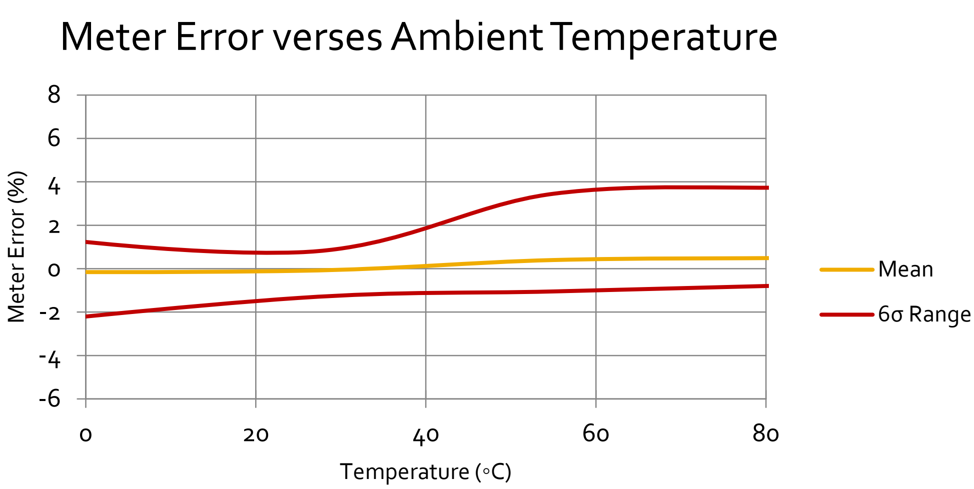

Figure 3

Current metering tolerance varies with temperature and may drift outside a practical range if used in extreme environments

Similarly, CTs are less accurate at the upper limits of their current sensing capacity as shown in **Figure 4**. This divergence is of little concern since the meter's primary purpose is to protect against over-current breaker tripping, which naturally coincides with the meter's higher accuracy range.

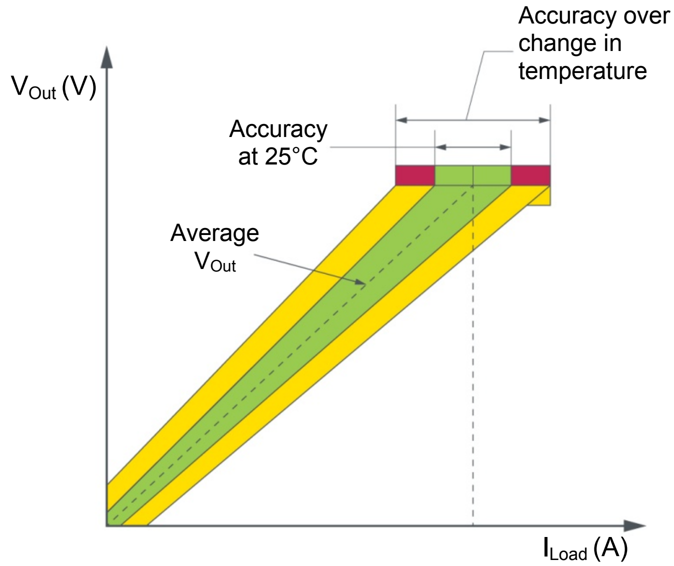

Figure 4

Current metering tolerance increases as the current load increases

An accuracy of ±5% or less is generally considered to be acceptable for data center use. For example, for a 20A breaker in the U.S., the rack PDU has an allowable continuous load rating of 16A. This provides sufficient buffer to prevent inrush currents at device startup from overloading the circuit. At 16A, a meter with ±5% accuracy will give readings from 15.2A-16.8A. 16.8A is still comfortably below the breaker threshold even accounting for inrush current spikes. This accuracy range is sufficient for power metering for departmental chargeback as well.

Meters with tighter accuracy offer little incremental benefit for data center applications, and are often specified for only a particular set of operating conditions. When PDUs are used outside this narrow band of operating parameters the tolerance grows wider.

Recommendation: Ensure the rack PDU is rated for continuous use in applications up to 45ºC (113ºF). Select a metered rack PDU with a ±5% accuracy tolerance or less since this accuracy provides sufficient visibility into circuit loading to prevent circuit breaker overloads.

If power is to be metered for reselling, a rack PDU certified to ANSI C12.1-2008 or IEC 62052-11 or 62053-21 is required. These standards define a more restricted accuracy range and special calibration procedure that provide the finer accuracy resolution for utility meters. This resolution and calibration is not needed for typical data center metering applications including most colocation power metering. Rack PDUs labeled as “Utility Grade Metering”, “Revenue Grade Metering”, or “Billing Grade Metering” but are not certified to one of these standards does not meet the criteria for metered energy reselling.

Step 5: Select form factor and mounting

Commonly, each piece of IT equipment has redundant power supplies intended to provide power in case the other fails. For data center applications, these power supplies are generally connected to separate redundant rack PDUs which are, in turn, fed from separate sources or circuits. This prevents the entire load from dropping in the event of a fault along one power path.

Rack PDUs install into the back of a server cabinet and provide convenient outlets accessible to both IT equipment and users that must configure the equipment. Two primary mounting orientations, illustrated in **Figure 5**, are:

- Horizontal 482.6mm (19in) rack mountable PDUs (a) mainly used with open frame racks and with audio/video equipment.
- Vertical 0-U PDUs that distribute outlets closer to the equipment they power. This style is the preferred orientation in data centers because they consume no rack U space and allow shorter power cords and require less cable management. This orientation provides a clearer and more visible power path for every cord.

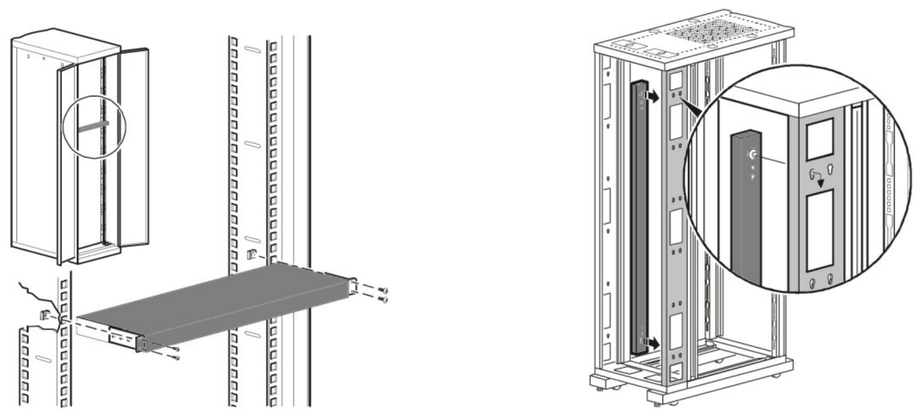

(a) (b)

Figure 5

Horizontal (a) and vertical 0-U (b) mounting styles

Recommendation: For datacenter applications, select a vertical 0-U form factor whenever possible for maximum available space and reduced cable clutter.

Conclusion

Selecting a rack PDU should start with an understanding of the type of equipment that will be installed. IT gear will dictate outlet type and quantity, as well as power consumption requirements. Using the power requirement estimates and information about the site voltage, whips and input cords can be appropriately sized.

Additional consideration should be given to preventing overloaded circuits and high-temperature applications. Metered rack PDUs can alert administrators before circuits are overloaded and reduce the risk of downtime. They provide historical power usage data which can be used to make better decisions than instantaneous readings alone.

For most cases, it is possible to standardize on one or two rack PDUs that are sufficient for practically any cabinet in the data center. The most common rack PDUs include:

- (36) C-13 and (6) C-19 outlets
- 20-30A (ANSI) / 16-32A (IEC) input power circuit
- Plug with at least 2m (6') of cord
- Branch metering with ±5% accuracy or better
- 0-U form factor

About the author

Brian Mitchell is Product Manager for Racks & Enclosures for Schneider Electric's IT Business Unit. With 8 years of professional experience, Brian has worked for Schneider Electric, Eaton and Emerson managing various projects related to power distribution and telecommunications. Brian holds Bachelor's degrees in economics and mechanical engineering and a Master's degree in mechanical engineering from the Missouri University of Science and Technology, as well as an MBA in business from Baldwin Wallace College.

Resources

Rack Powering Options for High Density
White Paper 29

How Monitoring Systems Reduce Human Error in Distributed Server Rooms and Remote Wiring Closets
White Paper 103

Switching to 240V AC Distribution in North American Data Centers
White Paper 128

How to Choose and IT Rack
White Paper 201

Planning Effective Power and Data Cable Management in IT Racks
White Paper 203

Browse all white papers
whitepapers.apc.com

Browse all TradeOff Tools™
tools.apc.com

Contact us

For feedback and comments about the content of this white paper:

Data Center Science Center
dcsc@schneider-electric.com

If you are a customer and have questions specific to your data center project:

Contact your Schneider Electric representative at
www.apc.com/support/contact/index.cfm