How to Choose IT Rack Power Distribution

White Paper 202
Version 2

Energy Management Research Center

by Patrick Donovan

Executive summary

One of the challenges associated with data center rack power distribution units (rPDUs) is 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 elaborate product selectors to narrow down the choices. Other 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. From supporting small IT loads to extreme high density loads like AI training models, this paper recommends and describes five steps for selecting appropriate IT rack power distribution.

RATE THIS PAPER ★★★★★

Introduction

Rack power distribution units (rPDUs) are available with a wide variety of different features, power ratings, and input and output receptacle combinations. Selecting the optimal rack PDU for the application can sometimes be difficult, especially without having exact information about the specific IT equipment and system configurations that will be installed in a rack. Trends such as increasing rack power densities, rapid growth in demand for high density artificial intelligence (AI) compute, and an emerging requirement to track energy consumption for sustainability metrics further increase uncertainty. To address this, IT professionals need a comprehensive selection strategy for powering equipment in IT racks. This paper describes five steps for selecting appropriate IT rack power distribution units and the practical decisions required to reduce downtime.

Determining which rack PDU to use in a particular rack requires general information about the equipment in the rack, the site power distribution, and need/preference for additional functionality. We recommend selecting rack PDUs using the following five steps:

1. Determine rPDU output receptacle type and quantity
2. Estimate power capacity of the equipment installed in the rack
3. Select rPDU(s) with appropriate power capacity
4. Select visibility and control options for branch circuits
5. Select mounting method

Note: this white paper focuses on alternating current (AC) powered IT equipment and not direct current (DC) powered equipment as typified by Open Compute Project (OCP) standards. See White Paper 232, *Efficiency Analysis of Consolidated vs. Conventional Server Power Architectures*, for more information.

Step 1: Determine rPDU output receptacle type and quantity

The first step starts with the IT equipment that will be installed inside the rack. Nearly all IT equipment come with detachable power cords which can have several different plug types. By far, the most common IT plug type in data centers connects to either a C13 or C19 receptacle[1] (see **Figure 1**) as defined by IEC 60320. Note, C13 refers to the female end of the cord and C14 refers to the male end. Likewise, C19 refers to the female end and C20 refers to the male end. C13 to C14 power cords are usually found on servers and small switches, while higher-powered blades, AI training servers, and larger networking equipment use C19 to C20 power cords because of their higher current carrying capacity.

Most equipment that cannot use either a C13 or C19 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 large storage arrays with large pin and sleeve connectors. In cases where the IT equipment cannot accept C13/14 or C19/20 power cords, make sure that the rPDU you select accounts for regional plug needs. You might need to get a separate rack PDU just to account for those plug types.

[1] Note, “receptacles” is also referred to as “sockets”, “outlets”, and “output connections”.

Figure 1

C13 and C19 receptacles are the most used for IT equipment; odd numbers refer to receptacles, while positive refers to the male connector end of the cord. The third number for each type refers to its temperature rating in degrees Celsius.

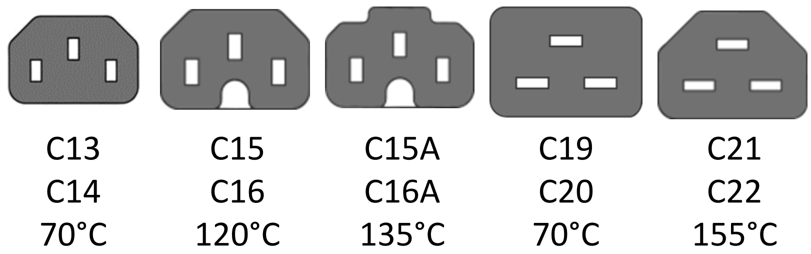

Rack PDUs should have at least as many plugs of each type as the equipment inside the rack so that every piece of equipment can be connected. Many data centers operators choose rack PDUs with more outlets of each type than are needed for the initial load. This allows additional outlets for future equipment. A popular outlet combination from several manufacturers has been (36) C13 and (6) C19 outlets because it allows for a mix of either high- or low-density equipment. From a typical data center perspective, this potentially means that a single, common rack PDU could serve practically any rack on the floor. Selecting a commonly used outlet configuration makes it more likely that the rack PDU will be in stock and immediately available whenever it is needed.

Note, it is possible for the rack power density to be so high (e.g., AI training clusters and high-performance computing) that the ambient temperature in the rack or hot aisle exceeds the 70˚C rating of the rPDU's standard C13 and C19 connectors. In such cases, 125˚C rated C15/C16 and C21/C22 connectors should be used respectively. Note that it's required to have both the outlet and plug at the higher temperature rating for the system to operate at the higher rating.

For added flexibility, there are some rPDU offers that provide combination outlets where each outlet accepts either C14 or C20 plugs. Figure 2 shows an example of that which offers "4-in-1" combination outlets. This type can be used as either C13, C19, C15, or C21 outlets providing significant deployment flexibility.

Figure 2

The APC Netshelter Rack PDU Advanced highlighting its "4-in-1" combination outlets.

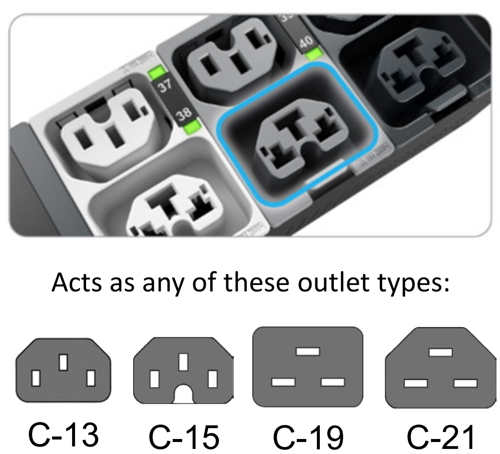

Step 2: Estimate power capacity of the equipment installed in the rack

There are three basic ways to estimate the maximum power required per IT rack. These are listed below from most accurate to least accurate.

1. Use IT equipment manufacturers' online sizing calculators. Many vendors including Cisco, Dell, HP, and Intel offer online calculators that more closely estimate the actual power draw as it accounts for the specific configurations

(redundancy, voltage, number and type of power supplies, drives, etc.) selected.

2. Estimate the power usage of the equipment inside the IT rack using the IT equipment nameplate power ratings. This method is commonly used when very high-power loads are used such as enterprise class servers, blade servers, or high-speed networking equipment. The nameplate rating on this equipment is higher than typical 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 quite conservative, which can lead to significant oversizing. Note, however, you may have historical power draw data available for similarly configured servers and networking gear in data center infrastructure management (DCIM) software. That data (and/or experience) could help estimate how much to de-rate the nameplate power ratings for new equipment to minimize oversizing.

3. Assume a maximum power level of the IT rack based on an estimate of total data center or data hall 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 most of the racks. This rack power density spec 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. Note, being the "least accurate" of the three options listed in this paper, if you plan to use a rack power density specification to estimate the power draw for racks, we recommend reading White Paper 120, *Guidelines for Specification of Data Center Power Density*. This paper describes the science and practical application of an improved method for the specification of both power and cooling infrastructure for data centers.

Once the power requirements are established, a rack PDU should be selected that provides at least enough power to support the rack's maximum load, otherwise you risk tripping a breaker resulting in downtime. Note that North American rPDUs will have 20A circuit breakers and are derated to 80% (16A) while Europe will have 16A circuit breakers, not derated (i.e., 100% rated). If the rPDU has a 20A power cord in North America, the rPDU's outlets will be protected by the building's branch circuit breaker, and hence, no internal circuit breakers are provided. Likewise in Europe, if the rPDU has a 16A power cord, there will be no internal circuit breakers. These rPDU-internal circuit breakers protect their local "branch" outlets.

Step 3: Select rPDU(s) with appropriate power capacity

Once IT equipment power draw for a given rack is known, the next step is to select an rPDU with the required power capacity and correct quantity. Determining this requires knowing:

- The site voltage and whether it's single-phase or 3-phase
- rPDU input power cord amperage rating (a.k.a., "whip amps")
- Redundancy requirements and space constraints inside the rack

Site voltage, number of phases, and the input amperage rating provides the information needed to calculate the rPDU's ANSI or IEC continuous power rating.

Redundancy requirements and rack constraints will factor into determining the appropriate quantity of rPDUs needed. See **Table 2** to determine the rPDU input cord amperage for a given total continuous kW capacity.

At 1N redundancy, the formula to calculate the continuous power rating of an rPDU is shown below.

In ANSI regions governed by the National Electric Code (NEC) the formula is:

$$ANSI\ rPDU\ power\ (watts) = voltage \times \#\ of\ phases \times amp\ rating \times 80\%\ derating$$[2]

In IEC regions the formula is:

$$IEC\ rPDU\ power\ (watts) = voltage \times \#\ of\ phases \times amp\ rating$$

In regions governed by IEC standards, there is no requirement to derate the power rating of the rPDU.

Recommendation: Choose an rPDU model that minimize the number of units required per rack. This means choosing one that uses the highest voltage available in the highest current rating available that meets the load requirements with some room for growth. Minimizing the number of units will help reduce cost and helps avoid space constraints within the rack. Higher voltage provides more power for the same current. For example, a 240V circuit provides twice the power of a 120V circuit. This topic is discussed in White Paper 128, *High-Efficiency AC Power Distribution for Data Centers*.

Power redundancy requirements, of course, also drive the number of rPDUs needed for a given rack. Typically, each piece of IT equipment has redundant power supplies intended to provide power in case one fails. For most 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 dropping the entire load if one power path fails.

Table 1 lists continuous power consumption ratings based on the input voltage, number of phases, and the input cord amps for three different scenarios:

Table 1A: 1ph, 120V/208V (i.e., older data centers, small edge computing sites)
Table 1B: 3ph, 240V/415V (i.e., as found in typical data centers in ANSI regions)
Table 1C: 3ph, 230V/400V (i.e., as found in typical data centers in IEC regions)

Note: All three tables show Line-to-Neutral (L-N) output voltages except for 208V which is Line-to-Line (L-L). The "Total kW" values in Tables 2A and 2B are 80% derated.

[2] Note, for North American ANSI-based systems, the rPDU's nameplate amp rating must be de-rated to 80% of its capacity. So, for example, a 20A-rated rPDU should only be operated at 16A.

Table 1A

ANSI continuous power consumption ratings for rPDUs used in single phase 120V/208V systems that are typical for older data centers and small edge computing sites.

Total kW	Phases	Output Voltage	Breaker amps	Qnty of rack PDUs
1.92	Single phase	120	20	1
2.88	Single phase	120	30	1
3.33	Single phase	208	20	1
3.84	Single phase	120	20	2
4.80	Single phase	120	50	1
4.99	Single phase	208	30	1
5.76	Single phase	120	30	2
5.76	Single phase	120	60	1
8.32	Single phase	208	50	1
9.60	Single phase	120	50	2
9.98	Single phase	208	60	1
11.52	Single phase	120	60	2

Table 1B

ANSI continuous power consumption ratings for rPDUs used in three phase 240V/415V systems that are typical for North American data centers.

Total kW	Phases	Output Voltage	Breaker amps	Qnty of rack PDUs
11.52	Three phase	240	20	1
17.28	Three phase	240	30	1
23.04	Three phase	240	20	2
28.80	Three phase	240	50	1
34.56	Three phase	240	30	2
34.56	Three phase	240	60	1
57.60	Three phase	240	50	2
57.60	Three phase	240	100	1
69.12	Three phase	240	60	2
72.00	Three phase	240	125	1
86.40	Three phase	240	150	1
100.80	Three phase	240	175	1

Table 1C

IEC continuous power consumption ratings for rPDUs used in three phase 230V/400V systems that are typical for European and other data centers located in IEC regions.

Total kW	Phases	Output voltage	Breaker amps	Qnty of rack PDUs
11.04	Three phase	230	16	1
22.08	Three phase	230	16	2
22.08	Three phase	230	32	1
34.50	Three phase	230	50	1
43.47	Three phase	230	63	1
44.16	Three phase	230	32	2
69.00	Three phase	230	50	2
69.00	Three phase	230	100	1
86.25	Three phase	230	125	1
86.94	Three phase	230	63	2
103.50	Three phase	230	150	1
120.75	Three phase	230	175	1

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. For example, connections made under a raised floor may become flooded with water, although this is uncommon.

Rack PDUs are commonly available in several different input 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.

Fault current

Now that you have selected an rPDU with the appropriate power capacity, you must be mindful of fault current ratings. Available fault current, measured in kiloamps (kA), is the maximum amount of current available to "feed" a fault (e.g., downstream of the rPDU) and is dependent on the electrical system design. Average rack densities have slowly increased over the years to well beyond 5 kW/rack. And in the case of AI training clusters, rack densities can exceed 100 kW. This means that the wire size and circuit current ratings have increased, both of which cause the available fault current at the rack to be greater than in lower density racks. The available fault current at a high-density rack today may be as much as 20X the fault current that was found in early data centers. Therefore, the risk of an arc flash increases as the rPDU amp rating increases. These larger diameter wires allow more fault current to pass through.

There are various solutions you can apply to the upstream electrical distribution network to reduce the available fault current at the input of the rPDU. Specifying higher impedance transformers and using inductors are two such solutions. The fault current at the input of rack PDUs should be limited to 10 kA 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*. See White Paper, *Arc Flash Mitigation*, and White Paper 253, *Benefits of Limiting MV Short-Circuit Current in Large Data Centers*, for more details on addressing arc flash hazards.

Step 4: Select visibility and control options for branch circuits

There are three levels of visibility into rack-level power status: basic, local device monitoring, and remote monitoring (often referred to as "metered"). A **basic** rPDU 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.

Rack PDUs with **local device monitoring** provide a display screen on the rPDU that shows instantaneous, "that moment in time" information only. This can be useful in determining phase balancing and general circuit load status. However, this only provides insight for the moment that it is being looked at. Decisions driven by this information will not consider peak usage or cyclical trends. These devices cannot signal outside of the local display and cannot alert remote IT staff about high-risk situations as they happen.

For data center applications, rPDUs that support **remote monitoring** (sometimes called, "metered rPDUs") allow remote visibility when circuits approach the maximum capacity and when input and output breakers are at risk for tripping. Using a built-in network management card, metered rPDUs 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 reduce risk of overloads. This type of proactive planning is a good approach to reducing the risk of downtime from tripped breakers.

Some rack PDUs that support remote monitoring also offer power and energy metering of individual outlets in addition to the branch circuits (i.e., banks of outlets) and input power feed. It empowers IT professionals with the necessary tools for advanced energy management such as tracking energy use by server or by department, for example.

Many metered 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 through a web interface or to DCIM software. IT equipment overheating due to a lack of sufficient cooling is a realistic cause of downtime. In-rack temperature sensors connected to rack PDUs are an easy way to ensure a proper operating environment for servers and other equipment. rPDUs that support remote monitoring can also be discovered and managed by DCIM software.

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 outlet switching controls as well as visibility to power draw and environmental conditions. Switching functionality allows outlets to be cycled on/off remotely via a network connection. This is a common need for remote sites with little to no staffing. 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 supervisor. The supervisor walks into the unfamiliar network closet and looks for cues to identify the troubled device to restart it. Even with good intentions, the wrong equipment is sometimes rebooted. Switched rack PDUs allow knowledgeable and authorized IT staff to manage equipment cycling even if they are off site.

Similarly, with colocation equipment, one of the most common calls is to request manual restarts of hung servers. To do this, the colocation 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 colocation provider to manage customers' equipment. This lag time, risk, and expense can be eliminated with rack PDUs that can switch outlets remotely. Figure 3 shows an example of control interface for controlling rack PDU outlets.

Figure 3

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

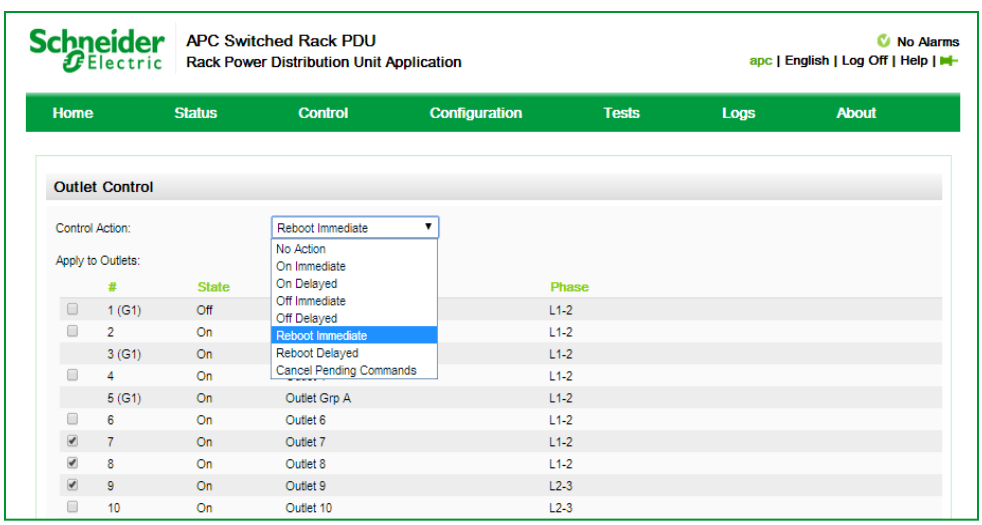

Accuracy of metering equipment

Like any measuring instrument, rPDU meters have inherent limitations in accuracy. And, by accuracy, we mean a measurement's degree of correctness to actual reality. The accuracy of a given rPDU's output current meter is driven by several factors including the vendor's selection of components, electronics design, the extent of calibration performed at the factory, ambient temperature during operation, and the amount of load connected to the rPDU. Meter accuracy, degree of factory calibration, and the number of measuring points (i.e., individual outlet metering, vs. whole device or outlet bank-level metering) are cost drivers of a given rPDU.

Because meter accuracy can vary from one vendor to another, from one SKU to another, and from one set of environmental conditions to another, it is important to understand each vendor's specifications when evaluating a specific solution. Some less scrupulous vendors might "cherry pick" and publish only one particularly good number (e.g., "+/-1%") without informing you that the actual accuracy varies as the load and/or temperature changes.

Typical rack PDU output current meter accuracy is in the +/-3% range for most operating conditions. This should be more than adequate for most users, particularly if the main reason for the metering is to ensure you do not overload and trip breakers. For this use case, its important that the accuracy performance is accounted for when setting high current alarm thresholds for the device or in your data center infrastructure management (DCIM) software.

If power is to be metered for reselling or true "billing grade" is desired, a rack PDU certified to ANSI C12.1-2008 or IEC 62052-11 or 62053-21 is required[3]. These standards define a more restricted accuracy range and a lengthier, special calibration procedure that provide the finer accuracy resolution for utility meters. This resolution and calibration are not needed for typical data center metering applications including most colocation power metering. **Note, rack PDUs simply labeled as "utility grade", "revenue grade", or "billing grade" but are not certified to one of these standards does not meet the criteria for metered energy reselling.**

Step 5: Select mounting method

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 6, are:

- Horizontal 482.6mm (19in) rack mountable PDUs (a) mainly used with open frame racks and with audio/video equipment.
- Vertical "Zero-U" PDUs 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 generally require less cable management. This orientation provides a clearer and more visible power path for every cord.

[3] We recommend visiting the official website of the International Electrotechnical Commission (IEC) or contacting the relevant standards organization to ensure these are the latest version of the standards. They will have the latest information on the status and updates of these standards.

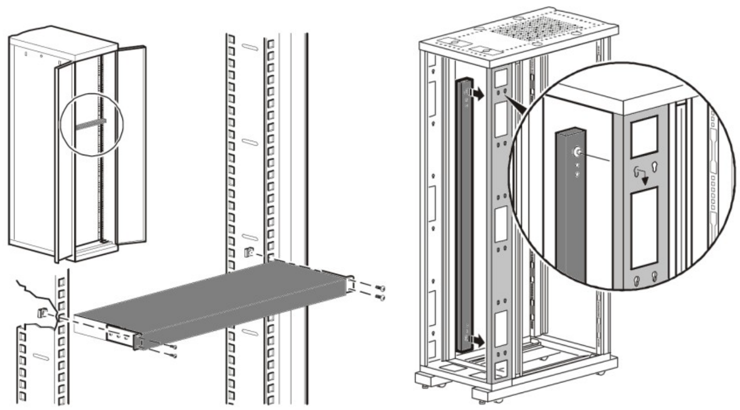

Figure 6

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

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

Conclusion

Selecting a rack PDU should start with understanding the equipment that will be installed in the rack. IT gear will dictate rPDU outlet type and quantity, as well as power capacity requirements. Selecting a specific rPDU that meets those requirements means knowing:

- the site voltage and whether its single or 3-phase
- rPDU input power cord amperage rating (a.k.a., "whip amps")
- redundancy requirements and space constraints inside the rack

An rPDU should be selected that minimizes the number of units required per rack. This means choosing one that uses the highest voltage available in the highest current rating available that meets the load requirements with some room for growth. Minimizing the number of units will help keep cost down and helps avoid space constraints within the rack.

In the case of more extreme rack power densities (i.e., 20 kW or more), either multiple rPDUs with standard, commonly available breaker sizes (e.g., 60A or 63A) should be used or higher-capacity custom rPDUs will need to be ordered (see **Table 2**).

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 and energy 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) C13 and (6) C19 outlets; or combination-style 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
- 0U form factor

About the author

Patrick Donovan is a Senior Research Analyst for the Energy Management Research Center at Schneider Electric. He has over 29 years of experience developing and supporting critical power and cooling systems for Schneider Electric's Secure Power Business unit including several award-winning power protection, efficiency, and availability solutions. An author of numerous white papers, industry articles, and technology assessments, Patrick's research on data center physical infrastructure technologies and markets offers guidance and advice on best practices for planning, designing, and operation of data center facilities.

Acknowledgements

Special thanks to **Brian Mitchell** for authoring the original content of this white paper.

RATE THIS PAPER ★★★★★

Resources

Efficiency Analysis of Consolidated vs. Conventional Server Power Architectures
White Paper 232

Guidelines for Specification of Data Center Power Density
White Paper 120

High-efficiency AC Power Distribution for Data Centers
White Paper 128

Arc Flash Considerations for Data Center IT Space
White Paper 194

Benefits of Limiting MV Short-circuit Current in Large Data Centers
White Paper 253

Browse all white papers
whitepapers.apc.com

Data Center Power Sizing Calculator
TradeOff Tool 1

Browse all TradeOff Tools™
tools.apc.com

Note: Internet links can become obsolete over time. The referenced links were available at the time this paper was written but may no longer be available now.

Contact us

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

Schneider Electric Energy Management Research 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