



## Power Distribution Considerations for Data Center Racks

### Overview

As data centers strive to become highly available as well as efficient, an important aspect of the entire data center infrastructure that needs special attention is their power distribution strategy. Traditional ways of distributing power to the racks are being constantly challenged with power densities as well as power consumption costs among data centers constantly rising. A recent study conducted by IDC found that the number one challenge facing data centers today is power/cooling. In this study, more than half of the respondents is taking a serious look at how to reduce data center costs by controlling power consumption. In addition, an EPA study mentioned that the power consumption costs of a data center can account for as much as 30 percent of a company's IT budget.

The following paper discusses the considerations that need to be made while choosing the appropriate power distribution products at the rack level. Interestingly enough, several of these considerations also help data center operators manage and control the largest portion of data center power consumption, which are the actual IT loads.

### Five major power distribution considerations for data center racks

- At what voltage do you run all the IT equipment?
- What electrical circuits do you bring to each rack?
- Number and type of outlets on the rackmount power distribution units
- What features do you look for within your rackmount power distribution units?
- Remote management of power distribution

## At what voltage do you run all the IT equipment?

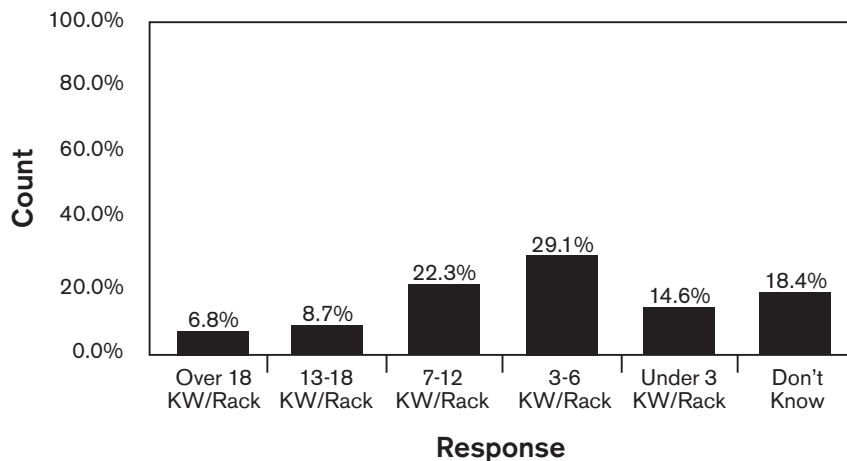
Most modern day IT equipment typically requires 1-ph power to operate. However, most IT equipment can accept voltages within the range of 100V – 250V. The reason is that all manufacturers like to build universal power supplies that could be used around the world. In North America, the electrical infrastructure within data centers is capable of providing 120V (phase to neutral) as well as 208V (phase to phase), requiring data center operators to make a choice between high and low voltage. In such situations, it is recommended that customers choose the higher voltage (208V). Running IT equipment at the higher voltage means a lower current draw for the equipment which translates into lower losses and better efficiency. Industry studies have revealed that running servers at 208V can bring as much as two percent efficiency advantage to the servers alone. Counting the added benefit of lower cooling required, the real gains might be as high as three percent to four percent. In light of this benefit, several enterprise level Unix servers, blade servers as well as large networking switches are being built with power supplies that can only accept 200 – 250V.

Outside of North America, choosing the right voltage is not an issue as the electrical infrastructure provides only one voltage option (220, 230 or 240V) that falls within the acceptable range.

## What electrical circuits do you bring to each rack?

Most IT equipment today comes with dual redundant power supplies. To be able to leverage the redundancy built within the IT equipment, data center operators should plan on bringing a minimum of two sets of completely independent circuits (N + N) into each rack. Ideally, both sets of circuits should come from different power sources and should be capable of handling the entire anticipated load capacity of the rack.

Given the rising power densities within data centers, operators need to pay very close attention to determining the anticipated loads within the racks. As shown in the figure below, a recent survey of 100 large data center operators revealed that 81 percent had load densities greater than the typical traditional power density of 3kW /rack.



Tables below mention the maximum load capacity that can be handled by various common circuits used globally.

North American Electrical circuits		International Electrical circuits	
Circuit	Max. Power (kVA)	Circuit	Max. Power (kVA)
120V, 15A, 1-ph	1.4	230V, 10A, 1-ph	2.3
120V, 20A, 1-ph	1.9	230V, 10A, 1-ph	3.7
208V, 20A, 1-ph	3.3	230V, 32A, 1-ph	7.4
208V, 30A, 1-ph	4.9	400V, 16A, 3-ph	11
208V, 20A, 3-ph	5.7	400V, 32A, 3-ph	22
208V, 30A, 3-ph	8.6		
208V, 60A, 3-ph	17.2		

While choosing the appropriate circuits to bring to the rack, it is important to note that all data centers typically have 3-phase power coming in. Traditionally, because of the lower power densities that could be easily handled by single phase circuits, operators chose to separate out the phases at the output of floor mount power distribution units and feed individual phases to separate racks. However, for load densities >5kW / rack, customers should consider bringing 3-phase power all the way to the rack level. Bringing 3-phase power to the rack level provides several benefits:

- Lower cost of cabling: Loads > 5kW within a rack can be handled by a single 3-ph circuit as against multiple 1-ph circuits. Lower number of circuits coming into each rack directly translates into lower cost of cabling.
- Higher reliability of electrical infrastructure: Bringing 3-phase power into the rack and utilizing rackmount PDUs with phase level metering capabilities allows data center operators to better balance the loads across all three phases. Balanced loads minimize harmonics and overheated neutral wires.
- Higher reliability of IT infrastructure: Bringing a lower number of circuits into each rack maximizes the air flow both under the raised floor and within the racks. Better airflow keeps IT equipment from overheating.
- Scalability: Higher capacity of 3-ph circuits provides more room for growth in the future. More equipment can be plugged in without bringing down power to existing equipment.


If a data center operator chooses to use 3-ph power for some or all the racks, the next decision point is whether to bring in a Wye-supply or a delta supply to the rack. Wye supply would also pull in a neutral wire to the rack whereas a delta supply would not. Internationally, since all loads require a neutral wire to be able to run 220 – 240V, a Wye supply is a must. In North America, if the operators have chosen to run all equipment at 208V, a neutral is not required and hence, a delta supply could be used. However, if there is any piece of equipment running at 120V, a Wye-supply should be used.

# Power Distribution White Paper


## Physical characteristics of power distribution units

Three major physical characteristic considerations for rackmount PDUs that need to be considered are:

- Form factor: Rackmount power distribution units come in vertical as well as horizontal form factors. Horizontal PDUs are typically limited to 10 outlets or less within 1U of rack space. If more outlets are required, vertical versions should be considered. However, before choosing vertical versions, it is important to consider whether or not there is adequate rail space to install those within the rack without interfering with the installation and movement of the IT loads.
- Type of outlets: Most modern day IT equipment comes with an IEC input power connector and modular input cords. This is a common practice among all manufacturers so that they are able to ship the same equipment worldwide and just localize the power cord. While the input power cord supplied with each piece of IT equipment can vary significantly, the IEC connectors are fairly standardized. Figure below shows the most common IEC connectors and plug combinations available:




IEC 320C13 Socket 10A  
IEC 320C13 Plug 10A



IEC 320C19 Socket 16A  
IEC 320C19 Plug 16A

In most situations, it makes sense for data center operators to choose PDUs with IEC outlets. Jumper cables that connect the IEC plugs on the equipment to the corresponding IEC outlet on the PDU are available both from equipment vendors and PDU vendors. The only time a PDU with IEC outlets would not work for customers in North America is when they have some equipment that has a hardwired input cord with a NEMA plug. In such cases, a PDU with the appropriate NEMA plugs would have to be chosen.

- **Number of outlets:** The PDUs should have the appropriate number of outlets required to power all IT load within the racks. While choosing the number of outlets required, it should be kept in mind that several enterprise pieces of equipment can have multiple redundant power supplies each requiring one or more power connections. For example, a Cisco Catalyst 6500 switch built with 8700W power supplies will require three power cords / power supply. A rack with several blade servers will typically require lower number of high power outlets, while a rack with several 1U servers will require high number of low powered outlets.



**Cisco Catalyst switch power supply with IEC input power connectors**

#### **Important features within power distribution units**

In order to ensure the highest availability levels for all mission critical equipment, it is important that data center operators consider power distribution units that have metering as well as switching capabilities.

Like most electrical circuits, whether it is residential or commercial, the cumulative current capability of all outlets on a branch circuit is always higher than the current rating of the branch circuit itself. With most data centers having several unused outlets available on the power distribution units, the plugged in loads are always susceptible to tripped breakers or fuses whenever new equipment is plugged in. To minimize chances of overloads, metering capability within PDUs along with the capability to set current thresholds is very important. If the PDU chosen has several branch circuits or it is 3-phase, metering capability should be available for all branch circuits or phases separately. Given the fact that nameplate current ratings for most IT equipment are highly overrated, metering capability among PDUs can also provide datacenter operators a gauge of the actual power draws of IT equipment. This information could allow data center operators to right size their upstream electrical infrastructure when either adding new racks or a new data center is being designed and built. Optimizing the electrical infrastructure improves the efficiency of the entire power chain, thereby cutting down on the power consumption costs.

Non-responsive servers or network equipment is a common issue that every IT Manager faces. Switching capability defined as the ability to remotely recycle, turn on or off power, is very important to minimize the downtime associated with hung equipment. It also helps keep the data center's operational costs low by reducing the number of expensive trips to the physical location of equipment. As power consumption costs among data centers rise, switching capability provides a way to control costs by providing the ability to turn off underutilized servers remotely.

Several modern day servers today ship with embedded service processor technologies such as IPMI, Dell's DRAC or HP's iLO. These service processors provide the capability to recycle power to the server. However, the best switching solution is still provided through a switched PDU offering. They allow power control of all IT equipment, irrespective of whether or not they have an embedded service processor which today is limited only to server models built in the second half of 2004 or later. Switched PDU solutions could provide several other enhanced capabilities that are important to improve the availability of the loads. Some of these features are:

- Sequential startup and shutdown: It is a well known fact that IT equipment draws a much higher current called the "inrush current" at the initial startup which is significantly higher than the current drawn during normal operation. For example, based on the specifications of an HP DL380 G5 server, its nominal current draw during normal operation is 8.8A but the inrush current is rated at a maximum of 30A for a maximum duration of three milliseconds. Unless the electrical circuit is capable of handling the inrush current of all equipment (which is seldom the case and is very expensive), starting up all loads within a rack at once can easily lead to tripped upstream breakers and unavailable loads. Switched PDUs provide the capability to stagger the turning on of outlets such that the current limits of the circuit never get exceeded at the time of startup of all equipment within a rack. It is also important to note that applications running on several servers can be dependent on each other and hence, they need to be started up or shutdown in a particular order. Switched PDUs allow that to happen.
- Electronic overcurrent protection: A PDU with just metering capability can certainly monitor the current draw and send alarms when programmed thresholds are exceeded. Switched PDUs can take further advantage of the metering capability built within them. When current thresholds are exceeded on a branch circuit, most switched PDUs have the capability to lock any unused outlets that reside on that specific branch. This ascertains that no additional equipment gets accidentally plugged into that branch, ensuring continuous availability to the powered loads.

## Remote management

In order to make the maximum utilization of monitoring and control features, and to help realize the vision of a “lights-out” data center, it is important that switched and metered PDUs are able to be managed remotely. When evaluating the remote management options available, the following questions should be considered:

- Does the remote management strategy of the PDUs align with the tools used to manage all other IT devices? For simplicity of use, it is important that IT managers have an integrated interface for all the devices they are responsible for, including PDUs.
- Do the PDUs also provide an out-of-band management option? In-band management tools are great to use, but their limitation is the vulnerability to problems from the very devices they manage.
- Does the remote management tool provide a global integrated view of all PDUs within the datacenter?
- Does access to the outlets follow secure authentication schemes?
- Does the remote management tool automatically discover all the attached PDUs?
- Does the tool provide grouping capabilities for outlets? Grouping allows power control of multiple loads together. This feature comes in especially handy when a server with multiple power supplies is hung up. By recycling power to all power supplies together, grouping ensures that the power to the entire server gets recycled.

## Summary

A well planned rack level power distribution strategy is important to ensure continuous power to all IT equipment and to monitor and control power consumption. Proper consideration should be given to the electrical, physical as well as the manageability characteristics of the PDUs. Electrically, IT equipment should be chosen to run within the 200 – 250 range whenever possible. When loads within a rack exceed 5kW, a 3-phase PDU should typically be chosen while a 1-ph PDU is good enough for loads below 5kW. Physically, PDUs need to be easily installable and they need to have the appropriate number of outlets to power all equipment. Remotely manageable PDUs with switching and metering capabilities provide the ideal feature set for highest levels of availability as well as data center efficiency.

## About Avocent Corporation

Avocent (NASDAQ:AVCT) delivers IT operations and infrastructure management solutions for enterprises worldwide, helping customers to reduce costs and simplify complex IT environments via integrated, centralized in-band and out-of-band hardware and software. Through LANDesk®, Avocent also is a leading provider of systems, security, and process management solutions. Find more information at [www.avocent.com](http://www.avocent.com).