



Deploying High Power to IT Equipment Racks

Table of Contents

Overview	Page 3
Trends in Data Center Power Deployment	Page 3
Drivers for High Power Racks	Page 3
Data Center Power Distribution Around the World	Page 4
What Is High Power?	Page 5
High Power, High Outlet Density	Page 5
High Power, Low Outlet Density	Page 7
Branch Circuit Protection	Page 8
Other Intelligent Rack PDU Feature Considerations	Page 9
The Advantages of Higher Voltage for High Power Racks	Page 10
Is It High Time for High Power?	Page 12
About Raritan	Page 12

Overview

Many data center managers are doing a good job conserving energy – decreasing PUE, raising data center temperatures, using air-side economizers to reduce energy consumption for cooling – but average power consumption at the rack is still going up. In fact, the increased efficiency means more power is available for servers to support data center growth. Data centers are finding that they must deploy more and more power to their racks. This white paper addresses considerations surrounding the deployment of high power.

Trends in Data Center Power Deployment

Data center managers are deploying more and more power to their IT equipment racks to keep up with power-hungry devices. From the chart below, nearly half (49%) of the data center managers polled had a maximum rack power density of 12kW or less. Their expectations were that two years later, only one-third (33%) would have a maximum rack power density of 12kW or less. Some data centers today have racks wired to provide as much as 30kVA.

	2kW or less	>2-4 kW	>4-8 kW	>8-12 kW	>12-16 kW	>16-20 kW	>20-24 kW	>24 kW	Unsure	Approx. Average
Now	1%	5%	20%	23%	16%	9%	8%	7%	11%	12.32
In Two Years	0%	3%	10%	20%	16%	14%	7%	11%	19%	14.57


Maximum power density (in kW) per rack in the data center.

Source: © 2010 Liebert Corporation, "Data Center Users' Group Special Report."

Drivers for High Power Racks

High power requirements at data center racks are driven by several factors, such as high-density racks filled with 1U "pizza box" servers. There are companies now deploying 1U servers in 54U racks. Another example is networking equipment such as Cisco® Nexus 7000 series systems. There are also blade server installations, such as multiple HP® c7000 chassis, in one rack. And network storage devices such as the Dell™ Compellent™ Storage Center FC enclosure which draws 450W for each 2U device.

Below is an ASHRAE chart showing the projected heat load, which is also the power consumption since each watt of energy consumed by IT equipment is converted to one watt of heat. Note that the vertical scale of the chart is logarithmic, so power demands are not leveling off but increasing dramatically.



Source: © 2005 ASHRAE TC 9.9 Datacom Equipment Power Trends & Cooling Applications.

Data Center Power Distribution Around the World

Typical voltages in North America are 120V and 208V. Some typical voltages internationally are 100V (Japan), 230V (Europe) and 240V (Australia). Since IT equipment vendors want to be able to sell their products globally, virtually all IT equipment is designed with power supplies which automatically adjust to voltages up to 240V.

Both single-phase or three-phase circuits can be distributed to racks. In North America, three-phase circuits are typically 208V, though 400V is becoming more common. For the rest of the world, three-phase power distribution is 400V (Europe and most of Asia) and 415V (Australia). Since the maximum voltage conventional IT equipment accepts is 240V, it will be the job of the rack or cabinet PDU to take a 400V input and convert it to 230V or 240V at the PDU outlets.

In many parts of the world, electrical circuits are specified as being rated at 16A or 32A. This is the actual current that these electrical devices are allowed to carry safely. In North America, electrical equipment is typically specified as 15A, 20A, 30A, etc. However, the National Electrical Code (NEC) requires that these values be "derated" by 20% to provide some headroom. So in North America an electrical device specified as 20A is actually rated at 16A (20A x 80%).

Watts (W) is used to specify the actual power consumed (active power). Volt Amps (VA) is used to specify the power that is available (apparent power). Think of apparent power as the design specification. For example, you could have a rack wired for 5.0kVA that is actually drawing only 4.2kW. This white paper will follow this convention, but the terms kW and kVA are often used interchangeably.

What Is High Power?

High power consumption at a rack may take the form of a few devices, each of which consumes a lot of power, such as blade servers and blade chassis requiring 5kW or more per chassis, or many moderate power consumption devices, such as a 42U rack filled with 42 1U “pizza box” servers, each server requiring 200-300 watts. There are several ways of deploying power in these scenarios and an approach which works for a high-outlet-density situation may also work for a situation where a lot of power needs to be deployed to a few power hogs.

Some data center managers add power by running more circuits. But, in general, it does not make sense to run several whips (power cables) to devices with multiple power supplies such as blade servers. It is easier and more economical to run two high-power feeds, either by under-floor whips or an overhead system, to a pair of high-power rack PDUs. From the high-power rack PDUs, short cables can be run to the power supplies, making for a much cleaner, e.g., less under-floor air obstruction, and more manageable deployment. Economics also improve with savings in copper and component costs.

When considering power demand, it is important to determine and design for peak actual demand. Designing to IT equipment nameplate ratings is excessively high. Designing for average power consumption may not be sufficient for periods of peak demand.

High Power, High Outlet Density

In the case of a large number of devices, each demanding a moderate amount of power, many power outlets will be necessary on the rack PDU.

A typical dense “pizza box” deployment would include two rack PDUs for redundant power where each PDU is loaded to 40%, so that if one power feed fails, the other feed will not exceed the NEC requirement of 80% (for North America). Typical outlets for “pizza box” servers are IEC C-13 (up to 250V, 10A international, 15A UL) and NEMA 5-20R (up to 125V, 20A, 16A rated). In this application, it is not uncommon to see a three-phase 208V 50A rack PDU with up to 54 outlets providing up to 14.4kW power per rack.

208V single phase vs. 208V three phase

If each server consumes an average of 200W, then the total power consumption is $42 \times 200W = 8.4kW$. The fully populated rack in this case requires 8.4kW. Therefore, as you size a rack PDU to support this load, you'll need to look for something that supports greater than 8.4kW. While rack PDUs are advertised in the market at a certain voltage, phase and amperage, kW ratings on rack PDUs typically already account for the NEC requirement of 80% load.

Because for three-phase power the sine waves are 120 degrees out of phase, calculating VA is slightly more complex than for single phase because we need to include the square root of 3, which is 1.732. The apparent power formula for three phase is $V \times \text{Derated } A \times 1.732 = VA$. A three-phase Delta deployment provides three separate circuits and more than 70% more total power than a comparable single phase, single circuit.

For example, a 30A 208V three-phase rack PDU will support 8.6kW. The math works as follows:

$$24A \text{ (80\% rating of 30A)} \times 208V \times \text{sq. rt. } 3 \text{ (or 1.73)} = 8.6kW$$

NOT $30A \times 208V \times 1.73 = 10.8kW$

Alternatively, if single-phase circuits are being run to the rack, then to support an 8.4kW load at the rack, you would need a rack PDU that provides a minimum of 60A. The math works as follows:

$$48A \text{ (80\% rating of 60A)} * 208V = 10kW$$

Finally, if you believe you need additional headroom for growth for potential increased server utilization that drives power consumption greater than an average of 200W, then an appropriate rack PDU could be 50A 208V three-phase which will support 14.4kW. The math works as follows:

$$40A \text{ (80\% rating of 50A)} * 208V * \text{sq. rt. } 3 \text{ (or } 1.73) = 14.4kW$$

Three-phase power is a way for one whip or rack PDU to deliver three circuits instead of just one. The whip or input power cord on the rack PDU will be larger for three-phase power because, instead of three wires (hot, neutral and ground), a three-phase cable will have four (Delta) or five (Wye) wires. Three-phase cables may be slightly larger than single-phase cables, but it is important to remember that one slightly thicker three-phase cable will be significantly smaller and weigh less than three single-phase cables for the same voltage and amperage. In addition, consider also that a single-phase cable at higher amperages can sometimes be larger than a three-phase cable at lower amperages.


Delta vs. Wye Three-Phase Power

The two three-phase alternatives are Delta and Wye. A three-phase Delta system will have four wires: Line 1 (hot), Line 2 (hot), Line 3 (hot) and a safety ground. Individual circuits are formed by combining lines. Three circuits are available L1+L2, L2+L3 and L1+L3.


A three-phase Wye system will have five wires: Line 1 (hot), Line 2 (hot), Line 3 (hot), a neutral and a ground. Individual circuits are formed by combining lines and by combining a line with the neutral. As an example, a three-phase 208V Wye rack PDU supports three 208V circuits (L1+L2, L2+L3, L1+L3) and three 120V circuits (L1+N, L2+N, L3+N). Three-phase Delta and three-phase Wye have the same apparent power, but the three-phase Wye can provide two different voltages, while the three-phase Delta can provide only one voltage.

In North America, there may be a requirement for 120V convenience outlets such as NEMA 5-15R (120V, 15A, 12A rated) or 5-20R (120V, 20A, 16A rated). These can be supported by 208V three phase Wye PDUs where wiring between lines (L1, L2, L3) and between lines and the neutral can provide power to both 208V and 120V outlets.

Whether the three-phase wiring is Delta or Wye, the rack PDU voltage is always referenced to the line-to-line voltage, not the line-to-neutral voltage. This is even true in the 400V three-phase example on page seven where all the outlets are wired line-to-neutral.




Delta
(Δ)



Wye
(Y)

Since the difference between Delta and Wye is whether or not there is a neutral, many data centers are wired for Wye and use whips terminated with Wye receptacles such as NEMA L21-30R. This means the data center can use Wye PDUs which support 120V/208V or use Delta PDUs which support only 208V without needing to change the data center wiring. A Delta PDU would use a NEMA L21-30P (the mating Wye plug), but would not use a neutral inside the PDU. This is a perfectly acceptable practice. For example, a data center could deploy Delta PDUs to racks where there is only a need for 208V and Wye PDUs to racks where there is a need for both 120V and 208V.



See Raritan's posting, ["How to Calculate Current on a 3-phase, 208V Rack PDU \(Power Strip\)"](#) where you will be able to download a three-phase calculator.

High Power, Low Outlet Density

In the case of high power consumption at a rack for a few devices, each of which consumes a lot of power, such as blade servers requiring 5+kW per blade chassis or data center network or storage devices, the total amount of power required might be comparable to, or even more than, the high outlet-density example above, but the number and type of outlets may be different.

Density for devices such as blade servers depends on their number of power supplies (often between two and six for redundancy), how the power supplies are configured (often a few power supplies are run close to their maximum while the others are idle because power supplies are most efficient when they operate at their maximum) and how many devices will be deployed in a rack.

In the case of a few devices demanding a lot of power, a large number of outlets may not be needed, but outlets capable of delivering substantial power may be required. Typical outlets for high demand devices such as blade servers at 208V or 230V are IEC C-13 (up to 250V, 10A international, 15A UL) or C-19 (up to 250V, 16A international, 20A UL) or, less commonly, NEMA L6-20R (up to 250V, 20A, 16A rated) or L6-30R (up to 250V, 30A, 24A rated) locking outlets. In some cases, blade server manufacturers have even supplied blade chassis with power connectivity for 30A three-phase plugs/receptacles.

One example of such a rack PDU would be 60A 208V three-phase rack PDU with 12 IEC C-19 outlets. In this scenario, one could support up to three blade chassis, each with six power supplies and running up to 5.7kW, or four blade chassis each with six power supplies and running up to 4.3kW. The math works as follows:

400V Three-Phase Example

$$48A \text{ (80\% rating of 60A)} * 208V * \text{sq. rt. } 3 \text{ (or } 1.73) = 17.3kW$$

$$17.3kW / 3 \text{ chassis} = 5.7kW$$

$$17.3kW / 4 \text{ chassis} = 4.3kW$$

400V three phase

As shown in the 208V / 120V example earlier, three-phase Wye wiring is a convenient way to step down voltage. This is particularly true for 400V power. A good and accepted practice to deliver a lot of power to densely packed racks is via 400V three-phase Wye rack PDUs. A data center designer could specify 400V Wye whips to 400V Wye rack PDUs. Since a lot of data center equipment operate on voltages up to 240V, the 400V Wye PDU can provide three circuits – L1+N, L2+N, L3+N – each supplying 230V ($400V / 1.732$). Note that 400V Wye rack PDUs do not lend themselves to supporting 120V outlets as do 208V Wye rack PDUs.

480V / 277V, a new high voltage deployment scenario

Facebook and OpenCompute have reexamined data centers and the equipment that goes into them. In order to create the most efficient data centers, they have looked into 480V three-phase Wye power where each line is wired to the neutral so the outlets deliver 277V. This Wye configuration, with lines wired to the neutral, is the same wiring configuration as the 400V / 230V wiring described above. The Facebook/OpenCompute approach increases efficiency, but it is highly customized. Today, most IT equipment doesn't have power supplies that can support 277V. Very common data center receptacles are IEC C-13 and C-19. These outlets don't support 277V.

There are savings and efficiencies (1% - 2% over 400V/230V three-phase systems) using the Facebook/OpenCompute approach, but it requires custom triplet racks, custom servers with custom power supplies, custom battery/UPS and 480V / 277V rack PDUs. It is an excellent concept, but because it is not currently an industry standard, applicability to more than a few data centers could be challenging.

Branch Circuit Protection

Since April 2003, Underwriters Laboratories (UL) has required branch circuit protection, either circuit breakers or fuses, for PDUs where the inlet current is greater than the outlet current, e.g., 30A (24A rated) input plug, 20A (16A rated) outlets. 15A and 20A (12A and 16A rated) rack PDUs can be supplied without branch circuit breakers because circuit breakers in upstream panel boards are deemed to provide the necessary protection. Rack PDUs with breakers or fuses are like mini-subpanels. For example, a 208V 30A (24A rated) three-phase PDU has three circuits and each circuit/set of outlets has a 20A circuit breaker.

Circuit breakers: single pole vs. double and triple pole

An important consideration is the reliability and flexibility of the branch circuit breaker configuration. Typically, circuit breakers are available as single, double or triple-pole devices. It is less expensive to use double-pole (or triple-pole) circuit breakers for two (or three) circuits, but there are some drawbacks. Double-pole breakers will trip if either of the two circuits they protect is overloaded. This means double-pole breakers are less reliable. Double-pole breakers are also limiting because if you choose to shut off a circuit, for maintenance for example, you have no choice but to shut off both circuits. Alternatively, some rack PDUs protect each circuit with a single-pole circuit breaker. This is more expensive, but single-pole breakers are more reliable and less limiting. Look for rack PDUs that allow only one circuit to be de-energized for improved reliability and flexibility.

Circuit breaker and line metering

Circuit breaker metering is a useful feature on any rack PDU with branch circuit breakers, but it is particularly important when dealing with high power because the consequences of tripping a breaker can be disastrous if it means losing several blade servers. With circuit breaker metering the end user sets a threshold. When that threshold is crossed, an alert is delivered so the end user knows power demand needs to be reduced or there is the risk of tripping a circuit breaker.

Line metering, intended for three-phase rack PDUs, is very useful for balancing the power drawn over each line. Overdrawing power from one line relative to another wastes available power, and unbalanced lines can place excessive demands on the neutral in Wye-configured PDUs.

Fuses vs. circuit breakers

There are several disadvantages to fuses relative to circuit breakers. Spare fuses must be stocked in inventory, in many instances fuses must be installed by a licensed electrician and the correct fuse must be used to ensure reliability and protection. However, if individual outlets need to be protected fuses may be the only practical alternative.

Other Intelligent Rack PDU Feature Considerations

Remote power monitoring and metering

Remote power monitoring and metering should include all the relevant power information, not just current draw. Such a selection of data ought to include current (amps), voltage, power (kVA, kW) and energy consumption (kWh). Since kWh is the measure to compare or bill for energy usage, it ought to be accurate, such as ISO/IEC +/- 1% kWh billing accuracy. And since equipment for different departments or customers may be in the same rack, ideally, there should be a dedicated energy-metering chip per outlet.

To manage power at the rack, power information from individual outlets, units, lines and circuit breakers should be available with user-adjustable power data sampling and buffering. Alerts via SNMP, email, SMS messages and syslog should be sent when thresholds are exceeded. Individual outlet and outlet group switching within a rack PDU and across multiple rack PDUs should be available. User-configurable, outlet-level delays allow data center managers to power sequence devices to reduce in-rush currents and to establish logical boot-up sequences.

Compatibility with energy management software is handy, especially when configuring many rack PDUs or aggregating the data from them. Usage reports by department, location or device; plots of usage over time and available capacity relative to actual usage help data center managers run their operations more smoothly and efficiently.

Flexible and versatile connections

To make remote monitoring, metering and management possible requires secure remote access via Ethernet and serial connections. To ensure security an intelligent rack PDU should have strong encryption and passwords and advanced authorization options including permissions, LDAP/S and Active Directory.

USB-A (host) and USB-B (device) connections are useful in taking advantage of many USB devices including memory sticks for downloading firmware and standard configurations, to collect logged data and to attach webcams. Intelligent PDUs may have other connections for special features, such as a port to connect sensors or network connectivity via WiFi.

Environmental sensors and maximum operating temperatures

For IT equipment, every Watt drawn becomes a Watt of heat. Racks drawing high power must have sufficient cooling, and equipment in the rack must be able to handle the high temperatures. It is important to ensure that the cooling system is adequate for the IT load in high-power racks. Some intelligent rack PDUs can support external sensors to monitor the cool air drawn into servers to ensure that the cooling system can maintain, say, 25°C (78°F).

But consider equipment located in other areas of the rack. Specifically, rack PDUs are typically located in a much hotter hot aisle or at the rear of the rack. For higher power racks, the exhaust temperatures from IT equipment will be much higher. This is why some intelligent rack PDUs are certified to operate at 60°C (140°F).

Rack PDU sensor compatibility with energy management software is handy, especially when aggregating data from several different sensors attached to several rack PDUs. The ability to plot sensors on a relative humidity vs. temperature cooling chart allows data center managers to determine if all the locations are within ASHRAE, equipment vendor or corporate environmental envelopes. It is also useful to plot environmental conditions over time to determine if equipment or process changes have affected temperatures, airflow or air pressure at one or more locations.

The Advantages of Higher Voltage for High Power Racks

Running higher voltages at lower currents means smaller cables, which use less copper, weigh less, take up less space and cost less. Running three-phase power instead of single-phase power means fewer cables, which simplifies deployment as well as the benefits of less copper, less weight and less cost.

	30A	30A	30A	60A
Phase	Single-phase	Three-phase Delta	Three-phase Wye	Three-phase Delta
Input Voltage	208V	208V	400V	208V
Output Voltage	208V	208V	230V	208V
Apparent Power	5.0 kVA	8.6 kVA	16.6 kVA	17.3 kVA
1U Servers	24	41	80	83
2U Servers	24	41	80	83
Blade Chassis	1	1-2	3-4	3-4

North America: 208V vs. 400V PDU power density

	32A	16A	32A	63A
Phase	Single-phase	Three-phase Wye	Three-phase Wye	Three-phase Wye
Input Voltage	230V	400V	400V	400V
Output Voltage	230V	230V	230V	230V
Apparent Power	7.4 kVA	11.1 kVA	22.1 kVA	43.6 kVA
1U Servers	35	53	107	83
2U Servers	35	53	106	83
Blade Chassis	1	2	4	8-9

Europe and Asia: 230V vs. 400V PDU power density

Plugs and receptacles are less expensive at higher voltages and lower current ratings. For example, a 30A 400V three-phase Wye (16.6kVA) plug (Hubbell NEMA L22-30P) costs \$32 and the receptacle costs \$41. A 60A 208V three-phase Delta (17.3kVA) plug (Mennekes IEC309 460P9W) costs \$166 and the receptacle costs \$216. The plug/receptacle combination is \$73 vs. \$382, respectively.

There are other benefits to higher voltages. By eliminating voltage transformations, 400V power reduces energy costs by approximately 2%-3% relative to 208V distribution and approximately 4%-5% relative to 120V distribution.

The case for running higher voltages in data centers

Consolidating data centers may reduce total power consumption overall but concentrate power demand in one data center or one set of high-density racks. Looking at the examples from the North America table above, a data center manager could increase power to a rack of IT equipment, which can operate on voltages up to 240V, by moving from 30A, 208V single-phase power which delivers 5.0kVA. Remaining at 30A, but replacing the 208V single-phase power with 208V three-phase power brings a better than 70% increase in power to 8.6kVA. If there is an even more significant demand for power at the rack, moving to 400V three-phase power while remaining at 30A increases the power to 16.6kVA, a more than 90% increase over the 208V three-phase and over triple the power compared to the original 208V single-phase power.

There will be some increase in the size of the cabling because of the move from single-phase power to three-phase power but, since the amperage was kept at 30A, the physical size of the cables will only increase modestly to accommodate the additional phases while the power available to the racks is dramatically increased. To accomplish a similar increase in power while remaining at 208V single-phase power, the amperage would have to increase to 50A to approximate the power available from the 208V three-phase deployment and to 100A to match the 400V three-phase deployment. These cables would be substantially larger than the 30A cables, there would be three times the number of cables and the data center would not realize the efficiencies that come with higher voltages.

Is It High Time for High Power?

Whether you operate a large, a medium or even a small data center, it may be time for you to consider deploying high power to at least some of your racks. Good candidates are racks that will be packed with 1U servers, racks with blade servers and racks with data center networks or storage devices. And, there are side benefits. Moving to higher voltages, whether single phase or three phase, reduces transmission losses which leads to energy savings. Higher voltages, especially when deployed as three-phase power, are a good way to increase rack power capacity without adding to cable clutter and blocking cooling air in under-floor plenums. High power racks, coupled with in-row or overhead local cooling, also eliminate the energy waste from moving air across the room since cooling is now localized.

There are several high-power alternatives from which to choose. A few representative examples were presented in this white paper. The best alternative for you depends on your current situation and plans for the future. But, high-power deployments, even three-phase 400V, are becoming more common and accepted and should be on your short list of deployment options.

High-density racks can be deployed in small, medium or large data centers. Even in our own small data center, we've increased temperature set points to where our cooling capacity has increased to support higher density rack loads. The total power consumption of a small data center may not be great, yet there may be racks with multiple blade servers or densely packed 1U servers that consume as much power as similar racks in a multi-megawatt data center.

About Raritan

Raritan is a proven innovator of power management, infrastructure management, KVM and serial solutions for data centers of all sizes. In more than 50,000 locations worldwide, Raritan's award-winning hardware and software solutions – including intelligent rack PDUs; energy management software; DCIM software solutions to effectively manage data center capacity, assets and change; KVM-over-IP and Serial-over-IP access products – provide IT and facility directors, managers and administrators with the control they need to increase power management efficiency, improve data center productivity and enhance branch office operations. Based in Somerset, NJ, Raritan has offices worldwide serving customers in 76 countries. Visit us at Raritan.eu or follow us on the Raritan blog. Raritan is an active member of the Green Grid, Climate Savers Computing Initiative and the Leadership in Energy and Environmental Design associations. The company was recognized by the EPA for its contribution to the agency's data center initiative.