

Optimizing Power Distribution for High-Density Computing

Choosing the right power distribution units for today and preparing for the future

*By Michael Camesano
Product Manager
Eaton Corporation*

Executive Summary

Fueled by the rapid rise of technologies such as virtualization and blade servers, computing densities in today's data centers are climbing dramatically. As a result, server enclosures are requiring more power than ever before.

The result of this is a new and demanding set of power distribution challenges. To meet the power requirements of their increasingly dense server racks, organizations are looking for power circuits and power distribution units (PDUs) that have adequate power capacity and make optimal use of panelboards, but don't drive such high densities that they are prohibitively expensive to cool.

This white paper shows why 30 amp (30A) three-phase circuits and 30A three-phase PDUs meet those demands more efficiently than comparable 20A and 50A alternatives. It also explains why utilizing 30A three-phase circuits and PDUs today puts you in the best position to handle future power capacity needs.

Table of Contents

Why Power Requirements are Escalating	2
Virtualization	2
Blade Servers	2
Issues in Selecting New Power Distribution Equipment	3
Optimizing Panelboard Efficiency	3
Optimizing Server Capacity Requirements	4
Preparing for the Future	6
Conclusion	7
About Eaton	7
About the Author	7

Why Power Requirements are Escalating


Though there are numerous factors behind today's rising computing densities and greater power needs, two increasingly popular technologies bear particular responsibility:

Virtualization

Server virtualization enables a single, physical server to support multiple "virtual machines," each running its own operating system and applications. More and more businesses are utilizing virtualization to reduce management costs and conserve data center floor space by consolidating their server hardware on fewer, more powerful host devices. According to Gartner Inc., nearly 50 percent of server workloads will be running on virtual machines by the end of 2012. While in theory this may seem more energy efficient, it is important to keep in mind that many organizations are hosting virtual machines on larger servers, that typically consume greater amounts of power than the legacy devices they replace.

Blade Servers

These are plug-and-play processing units with shared power feeds, power supplies, fans, cabling and storage. By compressing large amounts of computing capacity into small amounts of space, blade servers help reduce floor space requirements. They also enhance IT agility, since companies can simply plug in additional blades any time their processing needs grow.



Given such advantages, it's no surprise that the blade server market has experienced steady growth in recent years. In fact, blade server shipments will expand at a compound annual growth rate of 20.8 percent annually through 2013, according to Gartner, reaching nearly 22 percent of worldwide server shipments.

However, the high computing density that blade servers make possible also leads to extreme operating temperatures that can raise power requirements per enclosure nearly 1000 percent when compared with power and cooling requirements for conventional servers. Typically when adding a 1U or 2U server to a rack, it would draw a few hundred more watts from the branch circuit. Now, a new blade server can often consume 20 times as much current.

Issues in Selecting New Power Distribution Equipment


Given the high-intensity power requirements of technologies such as virtualization and blade servers, equipping today's server racks with three-phase power circuits is all but mandatory. Selecting the right amperage for those circuits is an important decision and offers a wider range of options. Currently, the three standard options are 20A, 30A, and 50A. When evaluating which amperage provides the best fit for a high-density data center, managers must consider both "upstream" issues related to panelboard capacity and "downstream" concerns related to server rack power requirements.

Optimizing Panelboard Efficiency

A standard 400A panelboard has 42 poles that can support up to 14 three-pole breakers. An optimal power distribution scheme seeks to occupy as many of those breaker positions as possible while also utilizing as much of the panelboard's 400A power capacity as possible. Neither 20A three-phase nor 50A three-phase circuits accomplish those goals as effectively as do 30A three-phase circuits.

20A three-phase circuits

A 400A panelboard can safely support 14 20A breakers, meaning no poles are left unused. However, those 14 breakers can consume a maximum of just 280A, or 70 percent, of the panelboard's load capacity. The remaining 120A goes to waste. Moreover, that waste creates additional inefficiencies upstream along the power chain, where other power distribution components are sized for 120A of energy that will never be used.



50A three-phase circuits

A 400A panelboard can handle no more than eight 50A breakers. Hence, a company that relies on 50A circuits will consume 100 percent of a panelboard's load capacity but leave 43 percent of the board's breaker positions unoccupied. That means you must buy additional panelboards to handle the same number of circuits as a single 400A panelboard can accommodate when you use 20A or 30A circuits. Those extra panelboards would be expensive, take up valuable space in your data center and require additional costly installation assistance from an electrician.

30A three-phase circuits

These circuits provide the most favorable balance of breaker position and power capacity utilization among standard amperage options. A 400A panelboard can safely accommodate 13 30A breaker positions consuming 390A of power. That leaves just three poles and 10A unused.

Optimizing Server Capacity Requirements

While panelboard capacity is currently the chief upstream power distribution challenge in data centers, the biggest downstream issue is growing server rack energy demands. In this case, 30A three-phase PDUs are also a more efficient solution when compared to 20A three-phase and 50A three-phase units.

20A three-phase PDUs

A typical 20A three-phase PDU delivers 5.7 kW of power. A typical set of blade servers, however, draws 4 kW of power. That means companies relying on 20A three-phase PDUs can place no more than one blade chassis, along with perhaps three 1U servers rated at 500W each, on any given server rack. This means 75 percent of a standard server rack is left empty, which is inefficient with regards to both floor space utilization and power and cooling.

In addition, relying on 20A three-phase PDUs to support blade servers can cause over capacity on upstream breakers. In theory, a PDU with 5.7 kW of capacity should have no trouble handling a set of blades drawing just 4 kW of power. However, some blade server chassis utilize two primary power supplies (along with a second set of redundant backup power supplies). Each primary power supply utilizes a separate line drawing 2 kW of power. 2 kW translates to 9.6A, which is well within the capacity of a 20A PDU.

Upstream of the server rack, however, those two separate lines combine to become a single three-phase circuit. To calculate the total load on that circuit, you multiply 9.6A by 1.73, producing 16.6A. Electrical code requirements limit the load on a 20A circuit to no more than 80 percent of total capacity, or 16A. So a 4 kW set of blade servers with two primary power supplies drawing 9.6A each will overload upstream 20A three-phase breakers.

Some blade products feature three primary power supplies (with a second set of three in reserve). In this case too, however, using 20A three-phase PDUs can result in overloaded upstream breakers. Though each power supply draws only 9.3A, that still adds up to 16.1A upstream, which is above the 16A maximum for a 20A breaker based on current electrical code.

50A three-phase PDUs

These products deliver a whopping 14.4 kW of power, more than enough to support high-density server racks. Unfortunately, 14.4 kW of power generates more heat than most conventional cooling solutions can handle. Currently, most organizations dissipate data center heat by placing computer room air conditioning (CRAC) units around the periphery of their server floor. This strategy is generally only effective on racks equipped with PDUs rated up to 10 kW of capacity. PDUs with more capacity than that usually require more "exotic" cooling solutions, such as in-row or overhead liquid cooling. And while these technologies are effective, they also impose multiple burdens, including the following:

- Safety concerns: Running water in close proximity to valuable electronic gear is inherently risky
- Costly maintenance fees: To guard against potentially hazardous leaks, businesses that use liquid cooling technologies must spend heavily on annual maintenance contracts
- Reduced floor space: Exotic cooling solutions take up floor space that could be dedicated more productively to server hardware instead

Above all, liquid cooling solutions are far more expensive than conventional CRAC-based cooling systems, as they require you to install and maintain costly chilled water and leak detection equipment. So-called passive cooling solutions, which pump chilled air from ordinary CRAC units directly through servers that have been sealed with foam or another material to keep cool air locked inside, are significantly less expensive. But they too cost more to install and maintain than ordinary CRAC-based room cooling systems.

Of course, organizations can always limit a thermal impact of a 50A PDU, and thereby eliminate the need for exotic cooling, by running it below capacity. But why invest extra in a 50A product if you intend to use only part of its capacity?



Figure 3 : “Exotic” data center cooling technologies like this in-row liquid cooling system are costly to install and maintain. They also expose sensitive server hardware to the dangers of leaky pipes.

30A three-phase PDUs

These systems enable organizations to meet server rack power requirements without enduring the disadvantages that come with 20A three-phase and 50A three-phase products.

A typical 30A three-phase PDU delivers 8.7 kW of power. This is enough to support two sets of blade servers in a given enclosure, or one set of blades plus another six conventional 1U servers. In addition, under electrical code requirements, a 30A three-phase circuit can support up to 24A. So whether your blade chassis uses two primary power supplies drawing 16.6A total or three primary power supplies drawing 16.1A total, upstream overloading won't be an issue.

Additionally, 8.7 kW power output of the 30A PDU fits comfortably below the roughly 10kW limit for traditional CRAC-based data center cooling technologies. Thus, using 30A three-phase PDUs frees you from the expense and risk associated with installing a more exotic cooling system and also spares you from further disadvantages:

- Excess capacity: 50A PDUs distribute twice as much power as the typical 6 kW to 8 kW server enclosure requires. Paying for that excess capacity is a costly endeavor. 30A three-phase PDUs provide more than adequate power to meet the needs of most server racks yet generally cost hundreds of dollars less than 50A products.
- Higher cabling costs: Due to their higher power capacity, 50A PDUs use thicker cables than 30A PDUs, which typically cost about ten percent more.
- Inferior load segmentation: Traditional 50A and 30A PDUs split loads across three breakers, limiting the productive capacity of that unit. However, thanks to its 14.4 kW capacity, each breaker on a

50A PDU is responsible for protecting 4.8 kW worth of IT infrastructure. That's the equivalent of a full set of blade servers or nine 1U servers. Any time a breaker trips on a 50A PDU, then, a substantial data center disruption results. A 30A PDU, by contrast, has 8.7 kW of total capacity, so each of its breakers is responsible for just 2.9 kW worth of load. That means that when a breaker on a 30A PDU trips, it impacts far fewer IT resources.

	20A three-phase PDUs	50A three-phase PDUs	30A three-phase PDUs
Overload upstream breakers	Load balancing on power supplies can cause overload conditions	Manageable, however 50A breakers are much less common.	Transparent balancing allows ease of installation
Safety concerns	Lack of local breakers increase user safety risk	Exotic cooling requires running liquids in close proximity to valuable electronic equipment	Local breakers provide additional protection
Exotic cooling requirements	Can not support enough load to drive exotic cooling	Expensive, large cooling systems would be needed to dissipate excess heat	Not required for loads of 9K or under
Reduced floor space	Not needed	Exotic cooling solutions would take up valuable floor space	Not applicable
Higher cabling costs	Minimized	More cabling would be needed to run cooling units	Minimized
Empty rack space	75% of the rack would be left empty	43% left empty	Only 7% left empty
Inferior load segmentation	Splits loads across three breakers, limiting productive capacity	Each breaker is responsible for 4,8kW worth of load - equivalent to full set of blade servers	Each breaker is only responsible for just 2.9kW worth of load

Table showing Optimizing Server Capacity Requirements- possible disadvantages

Preparing for the Future

Based on historic and current power requirements, it is safe to predict that data center power capacity requirements will continue to rise in the years ahead. In fact, power consumption in high-performance computing applications may soon reach up to 40 kW per rack. 30A three-phase circuits and PDUs are currently the best fit for today's power distribution challenges and position you to be prepared for future power demands.

Many companies will be tempted to address tomorrow's massive power requirements by replacing their existing circuits and PDUs with systems rated 80A or higher. But organizations that use 30A three-phase circuits and PDUs now needn't make such expensive investments. Instead, they can simply "mirror" their existing power distribution infrastructure by deploying a second panelboard as well as a second 30A three-phase circuit and PDU for each of their server racks. This approach enables them to double their power capacity while also maximizing pole usage and minimizing energy waste on the new panelboard. While the combined 17.4 kW of power produced by their dual 30A PDUs will require an exotic cooling solution, using PDUs rated 80A or above would require even more powerful—and hence costly—cooling technologies.

Conclusion

The widespread adoption of blade servers and virtualization has affected server rack power demands and has created new power distribution challenges. Data center managers need power distribution technologies that deliver adequate but not excessive power to server enclosures, minimize both wasted poles and stranded power on panelboards and prevent operating temperatures from rising to levels that only costly cooling solutions can handle.

Of the various amperage options available today, 30A three-phase circuits and PDUs are best equipped to achieve this difficult balance. Eaton's line of power distribution units, known as Eaton ePDUs, offer 30A three-phase circuits and provide a solution to growing power demands and challenges. A 30A three-phase circuit utilizes roughly 93 percent of the breaker positions and 97.5 percent of the power capacity on a 400A panelboard. Similarly, a 30A three-phase PDU provides more than ample wattage to a typical server enclosure, yet doesn't overwhelm conventional, CRAC-based cooling systems. 20A and 50A PDUs can't match these benefits. They're also less effective helping you prepare for future data center power capacity requirements.

While there's no such thing as a "one size fits all" solution to any IT infrastructure need, organizations coping with the effects of high-density server enclosures would be wise to consider making 30A three-phase circuits and PDUs a central part of their power distribution scheme.

About Eaton

Eaton Corporation is a diversified power management company with 2009 sales of \$11.9 billion. Eaton is a global technology leader in electrical components and systems for power quality, distribution and control; hydraulics components, systems and services for industrial and mobile equipment; aerospace fuel, hydraulics and pneumatic systems for commercial and military use; and truck and automotive drivetrain and powertrain systems for performance, fuel economy and safety. Eaton has approximately 70,000 employees and sells products to customers in more than 150 countries. For more information, visit www.eaton.com. In addition to UPSs, Eaton is your source for a comprehensive range of data center solutions, including power distribution, power protection, rack enclosures and accessories for network closets, computer rooms and data centers.

To find out more, visit www.eaton.com/powerquality or call us at 800-356-5794.

About the Author

Michael Camesano is a product manager in Eaton's Power Quality Solutions organization in Raleigh, NC. A 10-year veteran in the IT industry, Camesano has experience in IT hardware and software marketing, development and sales. You can reach him at michaelcomesano@eaton.com.

Tutorials on demand

Download Eaton white papers to learn more about technology topics or explain them to customers and contacts. Maintenance bypass, paralleling, UPS topologies, energy management and more are demystified in free white papers from our online library: www.eaton.com/pg/whitepapers.