

Is your data center ready for virtualization?

Important power considerations for virtualized IT environments

*By Chris Loeffler, Global Applications Manager, Distributed Power Solutions
Eaton Corporation*

Abstract

Virtualization brings the potential to deliver dramatic savings in terms of server count, footprint, power consumption and cooling requirements for data centers. For all its advantages though, virtualization also brings some unique challenges:

- Overall power consumption will be lower, but it will be highly variable.
- There will be fewer servers, but each one will be more critical than ever.
- Applications can be dynamically reallocated at will, but the support infrastructure cannot do the same.
- Data center footprint will be smaller, but overall efficiency might still be suboptimal.

The power and cooling infrastructure—which may have been quite sufficient for pre-virtualization needs—could easily become inadequate when data center performance patterns are radically altered.

The good news is that there are practical and affordable ways to address these challenges and improve data center efficiency in the process. This paper looks at some of the power-related challenges and the readily available technologies to address them.

Contents

The hidden challenges of virtualization.....	3
The challenge: Each server will be more power-hungry than before.	3
The challenge: Each enclosure will be more power-hungry than ever.....	4
The challenge: Higher rack densities may exceed available UPS capacity.....	6
The challenge: Applications and their processing demands can shift at will.....	8
The challenge: Overall data center efficiency metrics may still not be what they should.	8
Closing thoughts.....	9
About Eaton.....	9
About the author	9

Is your data center ready for virtualization?

Important power considerations for virtualized IT environments

A North American Fortune 500 company was facing challenges that look familiar to many data center managers. Server sprawl had led to an inventory of more than 5,000 distributed servers, most of them under 20 percent utilized. If nothing changed, the data center would soon outstrip its power and cooling capacity. A new facility would be required at a projected cost of more than \$10 million.

Instead, the company turned to virtualization. Mission-critical applications were redeployed onto 150 virtual Linux servers, occupying 20 percent of the original footprint. That left enough raised-floor space to triple processing capacity in the next three years, while cutting support costs in half and trimming \$15 million from total cost of ownership.

If these results sound good, they ought to. Virtualization has the promise to maintain or increase computing power and data center performance while controlling costs and extending the value of existing data center facilities.

Virtualization enables partitioning, whereby a single physical server runs multiple virtual machines, each with its own independent and secure application and operating system. Instead of provisioning a single, physical server with enough spare (often idle) capacity to support the peak load of a single application, you can now dynamically match available processor power across the data center to meet the variable demands of application workloads. The same work gets done, but there's far less idle capacity.

Consider a typical small company, running 240 applications, one application per server, each server operating at about 10 to 15 percent CPU utilization. In a typical scenario, about 200 or more of these applications would be candidates for virtualization, at an average rate of four applications per server. Consolidation would trim the configuration down to 52 physical servers in a virtualized environment and 30 conventional servers each running a single application, for a total of 82 hardware boxes. The savings are dramatic, as shown in the table below:

The ROI of virtualization	Before virtualization	After virtualization
Data center size	240 servers	82 servers
Server power draws	240 @ 200W	30 @ 200W, 52 @ 270W
Total power required	48 kW	20.04 kW
Cooling	13.6 tons	5.7 tons
UPS/electrical loss	7.2 kW	3.0 kW
Electrical cost	\$.08 per kWhr	\$.08 per kWhr
Yearly electrical costs (IT only)	\$33,638.40	\$14,016.00
Additional costs (UPS/electrical)	\$5,045.76	\$2,102.40
Additional costs (cooling)	\$27,078.91	\$11,282.88
Total yearly electrical spend	\$65,763.07	\$27,401.28
Total savings per year		\$38,361.79
Power improvement		58 percent
Reduction in cooling		7.9 tons
Reduction in UPS/electrical loss		4.194 kW

The hidden challenges of virtualization

For all its advantages, virtualization brings some unique challenges. Overall data center power consumption will be lower, but each server will draw more power. There will be fewer servers, but each one will be more critical than ever. Applications can be dynamically moved around as needed, but the support infrastructure cannot do the same. Data center footprint will be smaller, but overall data center efficiency might still be suboptimal.

The power and cooling infrastructure—which may have been quite sufficient for pre-virtualization needs—could easily become inadequate when data center performance patterns are radically altered. The good news is that there are practical and affordable ways to address these challenges, and improve data center efficiency in the process. Let's take a look at some of the power-related challenges.

The challenge: Each server will be more power-hungry than before.

The scenario: On an un-virtualized platform, the average server CPU (central processing unit) runs at only 10 to 15 percent of capacity. With virtualization, that figure jumps to about 70 to 80 percent. The higher the CPU utilization, the higher the power consumption per server. For example, an HP ProLiant DL360 Generation 5 server consumes an average of 190W at 10 percent utilization, 258W at 70 percent utilization, and 270W at 80 percent.

Other rack equipment is even more power hungry. For example, blade servers can consume from 3 kW to 6 kW per blade chassis, as much as 8 kW in high-performance compute configurations. Cisco Catalyst VoIP switches can consume up to 10 kW each.

Is your power distribution system up to these realities, or was it designed for the old 1U server world of 60-100W per U? Can it distribute enough in-rack power for high-density servers? Does it provide the needed visibility into power at the server level? Traditional power strips are no longer enough.

The solution: Increase the density of enclosure-level power protection and distribution. Enclosure-based power modules are available that can distribute up to 36 kW in only a few U of rack space. The newest generation of enclosure-based power distribution units distributes power in an organized manner to four to 45 receptacles for a wide range of power densities:

- **Standard power**—for applications less than 4 kW, 15A to 30A single-phase—to support up to six 1U/2U servers or peripherals
- **Mid-range power**—4-10 kW, 20A three-phase to 60A single-phase—to support up to 16 1U/2U servers or two blade server chassis
- **High-range power**—10-15 kW, 30A three-phase to 50A three-phase—for high-density applications, such as up to 25 1U/2U servers or three blade server chassis
- **Ultra power**—15 kW and up, 60A three-phase to 100A three-phase—for an enclosure that will be loaded with up to four blade server chassis and other power-hungry pieces of IT equipment

An advantage of modern enclosure-based PDUs, compared to traditional power strips, is the ability to securely view the status of each circuit from anywhere on the company intranet or the Internet, and receive automated alerts of potential trouble. As servers are pressed to carry multiple applications, availability at the server level becomes all the more important.

Beyond the power strip*Networking and intelligence for enclosure-level power distribution*

Enclosure-based PDUs now come with communication options that enable each device to be directly tied into the TCP/IP network with its own IP address for monitoring and management. Combine a smart PDU with power management software, and you can aggregate information from thousands of PDUs in one location—making it easier to understand their relationship to the compute environment.

Some models offer individual receptacle monitoring and control, enabling you to remotely turn an outlet on or off, or have it cycle. The ability to disable an outlet provides remote control of power delivery—preventing equipment installers from accidentally overloading power distribution devices and ensuring that power is available for new equipment according to plan. The ability to cycle an outlet enables you to remotely reboot a locked-up device without dispatching a technician to the site.

The challenge: Each enclosure will be more power-hungry than ever.

The scenario: Think back to our hypothetical company with the 240 servers, condensed to 82 servers after virtualization. The traditional configuration had spread out power consumption across eight full racks (assuming 12U of storage and networking equipment in each rack). The “after” scenario reduces power consumption overall, but increases power consumption per U and per enclosure.

Traditionally, data center managers could plan for about 60 to 100 watts of power consumption per U of rack space. A full rack of equipment averaged 3 to 4 kW of power. With today’s blade servers, that figure has escalated to 600 to 1000 watts per U, and growing. Power consumption may soon reach up to 40 kW per rack.

It’s typical to find there are not enough circuits to support these high-density racks, or enough power drops to add more loads. A typical electrical panelboard has 42 breaker poles to support either: 42 120V single-pole circuits, or 21 208V double-pole circuits, or only 14 208V three-pole (three-phase) circuits, or some combination of these circuit types.

A typical computer room that was wired five to 10 years ago was designed to feed one 20A 208V circuit to each rack (less than 3.5kW per rack). If you now have to support 20 kW of equipment in each rack, it could take up to six of these 20A circuits. The existing electrical infrastructure will be unable to support this load growth, and could easily run out of circuits or run out of capacity, especially with the growing prevalence of dual- and triple-corded loads.

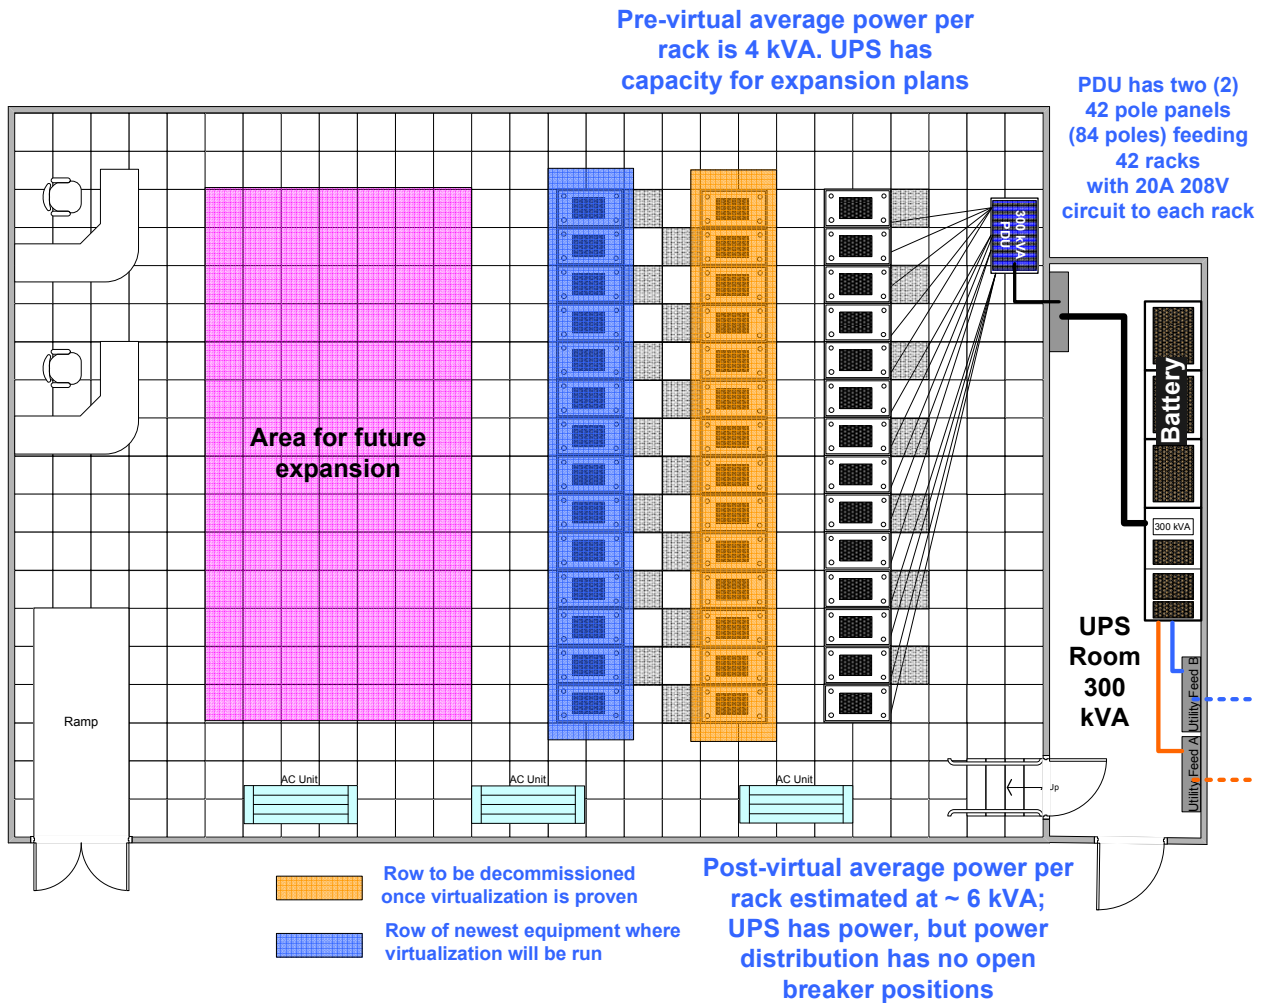


Figure 1. This computer room is constrained by its power distribution system.
How do you add a viable power distribution strategy with minimal impact to the raised floor?

The solution: Build out a power sub-distribution strategy. Instead of running individual cable drops from your large UPS and PDUs to each rack, run higher powered subfeed circuits to an intermediate remote power panel (RPP), power distribution rack (PDR) or rack mounted power distribution device—and from there to enclosures:

- RPPs are now available with 400A-rated panelboards to support the higher-powered PDUs now required in enclosures. These RPPs provide up to 168 factory-installed branch breakers in a freestanding enclosure. If floor space is at a premium, RPP functionality can be provided in a wall-mounted package that supports up to 84 factory-installed branch circuit breakers.
- A PDR provides RPP functionality in an enclosure that matches your data center enclosures.
- Rack mounted power distribution devices come in ratings up to 125A, and can sub-distribute power to 12 breaker poles with matching single- or three-phase receptacles. Any of these sub-distribution power panels can be added to existing distribution systems to expand power distribution to match the density increases that come with virtualization.

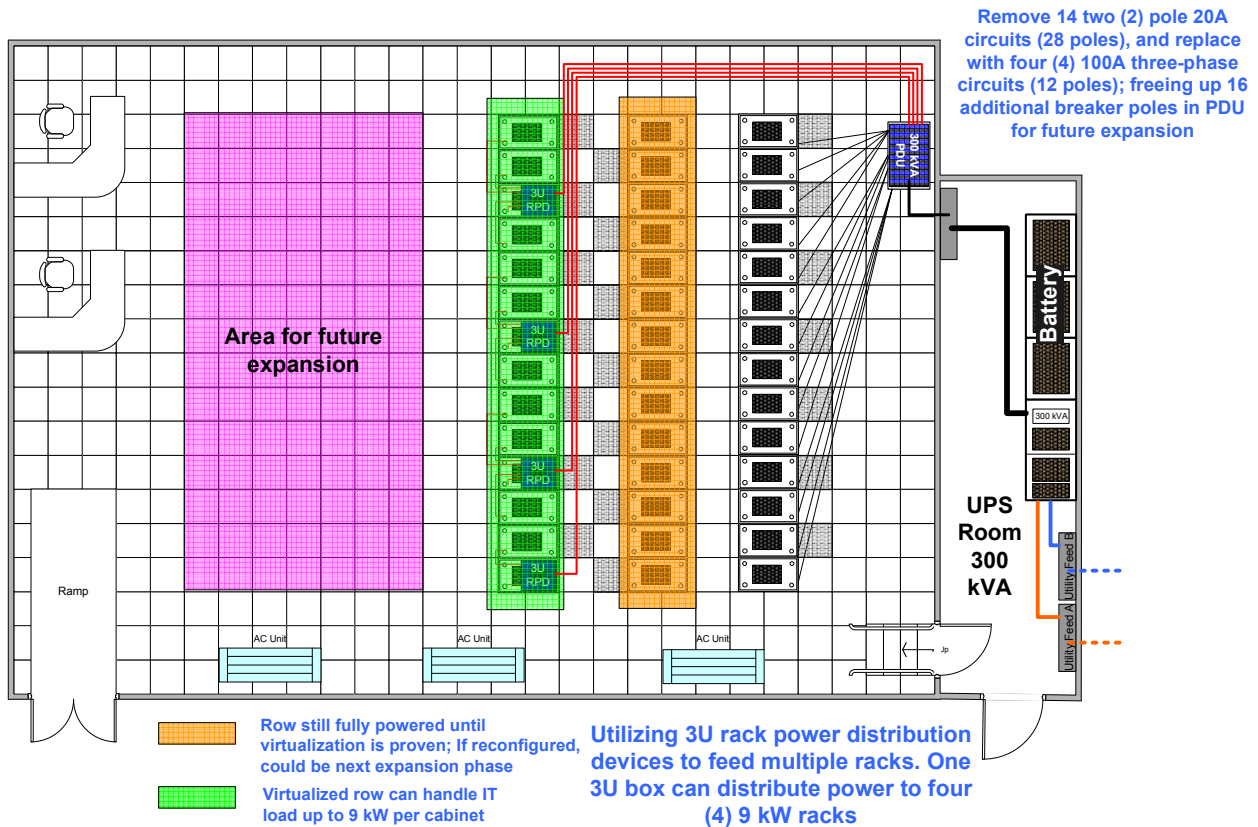


Figure 2. A sub-distribution strategy extends the capacity of the existing distribution system.

The challenge: Higher rack densities may exceed available UPS capacity.

In traditional data center designs, one or two large, three-phase UPSs stood alone in a separate room, providing conditioned power and battery backup for the whole data center, perhaps even the entire building. These UPSs fed large PDUs on the data center floor.

Virtualization projects often result in a cluster of high-density racks in a data center that was designed for lower density racks. In these cases, the capacity of the existing, centralized UPS can become a bottleneck to virtualization efforts, especially as redundancy becomes more important.

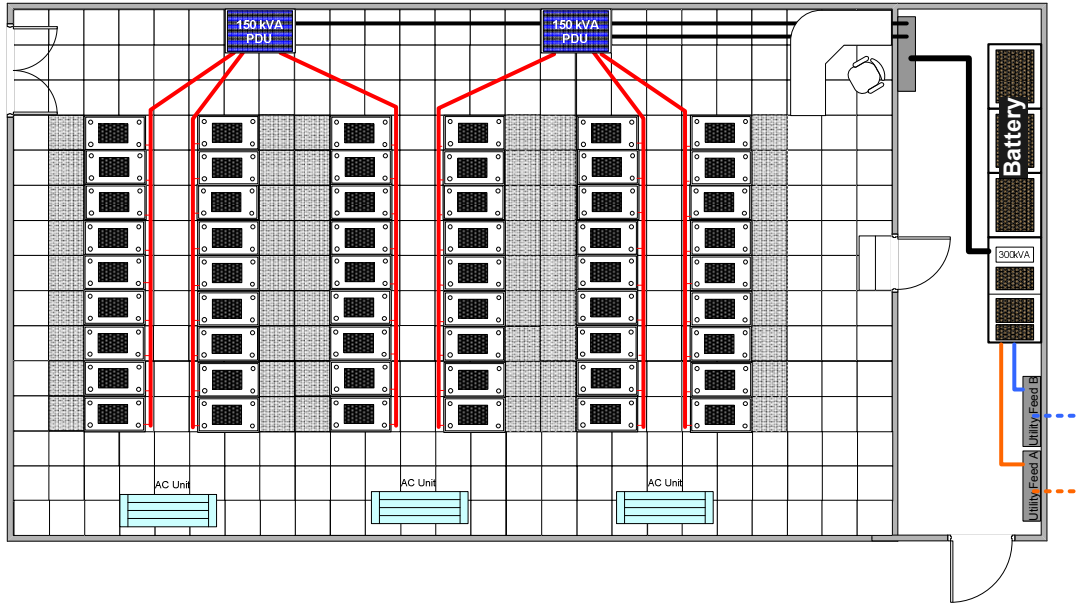


Figure 3. The centralized UPS in this 1500-sq.ft. computer room cannot support additional equipment.

The solution: Enclosure-based UPS systems with power densities of 2 kW or more per U are available to meet these new high density computing demands. Modular, scalable UPS systems can be deployed in a variety of system architectures for centralized, zone or distributed power protection—with redundancy to meet businesses objectives. Partner the UPS with a high powered distribution system to flexibly deliver power to loads of various voltages, power cords and layouts.

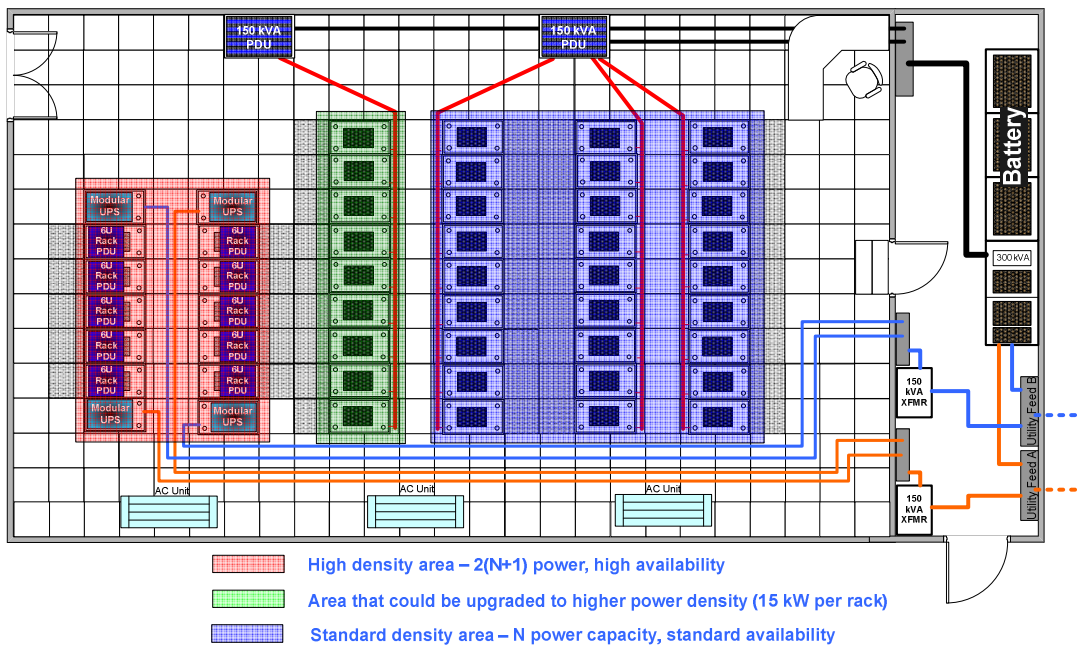


Figure 4. Modular UPS systems on the computer room floor add capacity where needed.

The challenge: Applications and their processing demands can shift at will.

Virtualization makes it possible to deploy, move, or clone an application from one platform to another at will, even while it is running. That means energy demands can shift around the data center at will. Yet power is a fixed asset, tied to a physical infrastructure that cannot be pressed beyond the limits specified by the National Electrical Code.

On-demand migration of applications demands new levels of visibility into how IT applications affect power, and vice versa. Is the IT equipment that runs the most important applications receiving computer-grade power? If you move processing-intensive applications, will you overload a branch circuit? Which rack has enough power to accommodate new virtualized servers and their applications?

The solution: These questions are answered with 7x24 power quality metering, monitoring and management at the branch circuit level, which can be conducted at several points in the data center power distribution system. The right strategy will be a trade-off between the number of IP addresses you're willing to allocate, the number of devices to be monitored, and the degree of detail required.

Looking further, with virtualization it becomes important to align power conditions with the shifting applications being supported—and to optimize IT and facilities issues in holistic context. After all, application demands affect power consumption, and power consumption affects application availability and performance—yet rarely do these two interdependent groups work in unison.

This reality calls for a monitoring system that unifies IT and facilities under a single pane of glass, one portal to view the status of IT resources, power and thermal conditions. A converged monitoring system gives both groups real-time visibility into power conditions, with power circuits correlated to IT resources and linked to business services. With this type of holistic perspective, IT and facilities managers can easily see when abnormal conditions threaten the IT infrastructure—and exactly which business services are at risk, so proactive measures can be taken.

The challenge: Overall data center efficiency metrics may still not be what they should.

Virtualization delivers dramatic energy savings, but could it be even better? In a typical one-megawatt data center, a 10-year-old UPS is probably wasting about 150 kW of power and dissipating a lot of heat.

The solution: Capitalize on newer, more efficient UPSs. Recent advances in UPS technologies have greatly improved efficiency. In the 1980s, a state-of-the-art UPS was 75 to 80 percent efficient at best. With the advent of faster switching devices in the 1990s, efficiency jumped to 85 to 90 percent and later approached 95 percent, when transformer-free, large UPS designs became viable.

Even higher efficiency is now possible. In 2007, high-density UPSs became available in very efficient designs. One modular, three-phase UPS operates as high as 97 percent efficiency—and retains high efficiency even when loaded at less than 30 percent utilization, where you would typically expect much lower efficiency.

Even small increases in UPS efficiency can quickly translate into tens of thousands of dollars. For example, a 10-year-old UPS delivering 60 kW and operating at 90 percent efficiency uses an average of 584,000 kWhr (kilowatt-hours) of power per year. In contrast, a UPS at 97 percent efficiency would consume only 541,855 kWhr—a reduction of more than 40,000 kWhr. Since the more efficient UPS also dissipates less heat—and therefore requires less cooling on the data center floor—you'd save approximately \$30,000 in energy costs in less than five years (assuming a utility rate of 10 cents per kWhr).

Closing thoughts

Server virtualization enables a data center to meet its service level agreements with fewer physical servers, which in turn results in dramatic savings in energy and footprint. However, to successfully implement a virtualization project, IT teams need to account for some new realities, such as dynamic changes in power demands, higher server and rack power densities, and the critical need to protect applications with sufficient UPS capacity.

Fortunately, solid solutions are available to address these power considerations—to extract the full potential of electrical savings offered by virtualization and server consolidation, while ensuring that power capacity meets demand at the row, rack and server levels.

About Eaton

Eaton Corporation is a diversified power management company with 2008 sales of \$15.4 billion. Eaton is a global technology leader in electrical 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 75,000 employees and sells products to customers in more than 150 countries. For more information, visit www.eaton.com/powerquality.

About the author

Chris Loeffler is Global Applications Manager for Eaton Corporation, specializing in data center power solutions and services. With more than 18 years of experience in the UPS industry, he has overseen product management of more than 20 UPS products for data center and industrial applications, including the Eaton 9330, Eaton 9395 and most recently, the new Eaton BladeUPS.

Loeffler has held a variety of positions within Eaton, including roles in service engineering, application engineering and more than 10 years in product management. He has authored a number of articles for trade publications and written several white papers on energy efficiency and UPS topologies for industrial and/or data center applications. Loeffler's educational background is in electronics engineering, and he is a member of the Blade System Alliance.

You can reach him at 919-870-3371 or ChrisALoeffler@Eaton.com.