White Paper

Redefining the economics of running the modern data center

Is your IT organization capitalizing on the latest technology options and best practices?

Dr. Kenneth Uhlman, PE, Director, Data Center Business Development, Eaton

Abstract

IT teams grapple with some inescapable realities. The cost of operating the data center is outpacing the cost of IT resources themselves. High-density equipment pressures power and cooling systems to the max. Moves, adds and changes can turn the power infrastructure into a deck of cards, at risk of overloads, tripped circuits and unplanned shutdowns.

It doesn't have to be that way. There are practical and affordable ways to ease these concerns without making major changes in the power delivery system.

High-efficiency power quality systems, greater visibility into power conditions at all levels, flexible options to power all those dual- and triple-corded servers, modularity for flexible growth… new approaches and technologies such as these are redefining the economics of running the modern data center.

Is your organization taking advantage?

Contents

Redefining the economics of running the modern data center

Is your IT organization capitalizing on the latest technology options and best practices?

Dr. Kenneth Uhlman, PE, Director, Data Center Business Development, Eaton

Although data centers can now pack more processing power into less real estate, high-density computing environments can be a huge drain on operating budgets.

For one, data centers and servers consume a considerable amount of the nation's total supply of electricity—some 61 billion kilowatt-hours (kWh) as of 2006, about 1.5 percent of total U.S. electricity (*Environmental Protection Agency (EPA) Report to Congress on Server and Data Center Energy Efficiency*, 2007). That's about the same as the energy consumption of the entire U.S. transportation industry (including manufacture of cars, planes, trucks and ships). Data centers consume more power than all the color televisions in the country put together—double the demand of only six years ago.

Second, the cost for each of those kilowatt-hours keeps going up. For many IT organizations, energy costs now represent the largest component of total cost of ownership—and the most stifling influence on IT expansion.

If you manage a data center—or engineer the architecture for clients who do—you know how critical these issues have become. How can you support all those growing loads without bringing unwanted governmental scrutiny or utility surcharges for being an energy hog? How can you reduce energy costs without compromising the overall reliability of the data center? How can you exploit new technologies without overhauling the power infrastructure?

Fortunately, there are practical and affordable answers to these questions. Let's look at six strategies you can capitalize on right now to redefine the economics of operating your data center.

1. Optimize IT and Facilities as a synergistic entity.

Ask CIOs about their data center strategies, and they can eloquently discuss how their computing, storage and network plans will support the organization's objectives. Ask them how Facilities or the energy bill fits into the picture, and the answer will likely be, "Oh, that's handled by another group."

In this type of silo organizational structure, critical issues are tossed over the wall between groups, rather than addressed collaboratively. Working separately, neither group is thinking about optimization across both groups, such as how to maximize energy efficiency or strike the optimum balance between cost and reliability.

In fact, one group's solution could actually be creating a problem for the other. Potential efficiencies and advantages are missed. For example, the server that offered the lowest initial price might be so power-hungry that it's the costly choice over the long term. The deployment strategy that met an immediate need might be cumbersome to change and upgrade later. The power distribution scheme that worked yesterday might be an inflexible hindrance to data center adaptation today.

Collaboration between IT and Facilities could reduce costs, increase agility and flexibility, reduce power consumption and risk, increase system reliability, and give both groups access to a broader range of knowledge experts and best practices.

The technology is available today to bridge that gap, to leverage existing data center meters and monitoring systems to reveal costs and efficiencies, not just availability and utilization. (For more, refer to "See IT and Facilities through a single pane of glass," later in this document.)

True collaboration across these groups also requires executive buy-in, cultural change and probably new incentives as well. The business might actually be structured to encourage the wrong behaviors. For instance, transaction and sales volume may have doubled, so sales employees get bonuses, and IT people get kudos for keeping up, while Facilities gets chastised for a 20 percent spike in the utility bill. Such disconnects are surprisingly common, so organizational/process change is key.

In the quest for efficiency, is the data center heating up?
The debate over higher operating temperatures

British Telecom is leading the trend with its announcement that its data centers are now 15 to 20 degrees warmer than the typical 60°F to 65°F. The American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) has loosened its standards for “allowable” data center temperatures by about 5°F and relative humidity about 5 to 10 percent.

Given that there are millions of dollars to be saved by turning up the thermostat, should other data center managers be following suit?

Not necessarily, or at least not yet. For one, IT vendors still hold fast to the most stringent ASHRAE recommendations, not the more flexible “allowable” temperature and humidity ranges. You could be at risk of invalidating your equipment warranties if you don’t comply.

Second, a colder data center gives you a greater margin of safety in the event of a component failure. If an AC unit goes down or air vents are plugged, you’d have more time to take action before you hit the thermal limit.

2. Monitor the big-picture perspective and the detail view.

The big picture: real-time view of power conditions

In the past, basic power monitors were used to identify conditions on an electrical distribution system or to evaluate past problems. If you wanted to detect fast voltage transients, you had to bring in portable power quality monitors that cost up to $20,000 and usually required an outside consultant.

Thanks to the latest technologies, we can now match the sampling rate of those high-end systems in compact, affordable devices. These Web-based, permanently installed instruments (with optional local display) monitor, record and analyze critical aspects of an electrical distribution system—so you can optimize energy utilization, process performance and cost.

In designing these power quality instruments, Eaton set out to demystify power quality— to take highly specialized data and convert it into useful information that doesn’t require a power guru or consultant to understand.

The result is an uncommonly easy-to-use interface and new graphical representations of complex power quality data—delivered via email, over the Web and to third-party applications without additional software or licenses. With these capabilities, the IT team can predict and prevent power quality problems before they lead to equipment malfunction, overheated circuits or unplanned downtime.

With the Power Xpert® Server, IT-friendly Web screens show color-coded displays of energy consumption, system health and power disturbances.

Power meters with built-in servers and thin clients demystify power conditions for IT and Facilities managers—and provide the hard for sustainability initiatives.

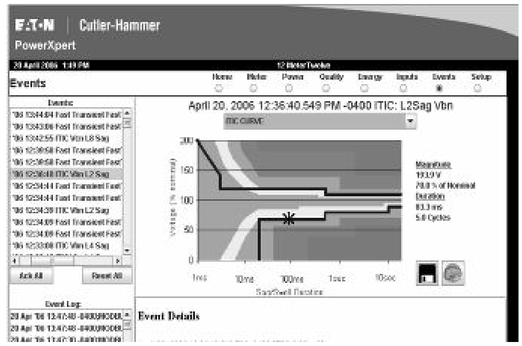

The detail view: monitoring at the branch circuit level

Millions of dollars a year are invested in power protection systems, such as UPSs and generators, yet there still might be a proverbial weakest link at the local level—the branch circuit that serves critical equipment. At this level, you might not be able to see trouble coming until a circuit breaker trips, and that's too late. Systems go down. Valuable data is lost, and business comes to a standstill. It can take hours to recover.

Branch circuit monitoring continuously measures the current on all breaker levels and warns you of impending trouble, so you can take proactive steps. Armed with these insights, data center and facilities managers can more effectively manage energy consumption to prevent overload conditions, optimize power distribution, and, when applicable, allocate energy cost among internal departments.

But what if your data center uses legacy or third-party panelboards, power distribution units and remote power panels that do not have branch circuit monitoring built in? No problem. Eaton offers an upgrade kit that retrofits those devices for branch circuit monitoring. With split-core current transducers (CTs), you don't even have to disconnect the wires to add branch circuit monitoring.

With this capability, you bring the entire power distribution system under the aegis of a unified energy management system. Even if you have a mix of older equipment and elements from other vendors, you get the insights to effectively manage the edge of the power distribution system.

See IT and Facilities through a single pane of glass.

Historically, the IT group monitored IT and networking equipment, reported to the CIO. Facilities monitored the power, cooling, security and real estate functions that underpin the data center—and reported into Corporate Real Estate or some other group. Data center designers focused on availability, not efficiency, and they deployed meters and monitoring systems toward that goal. Facilities functions had a huge impact on the data center—and vice versa—but these functions were monitored quite independently.

Now the data from IT metering and monitoring systems can be recast to reveal efficiencies and costs, not just availability. The next step is to link these new metrics into the organization's business service management dashboards.

The result would be a unified view of IT and Facilities across the enterprise. An integrated, enterprise-wide monitoring system provides valuable information for the organization to understand the impact on the business and continually improve operations to the benefit of both.

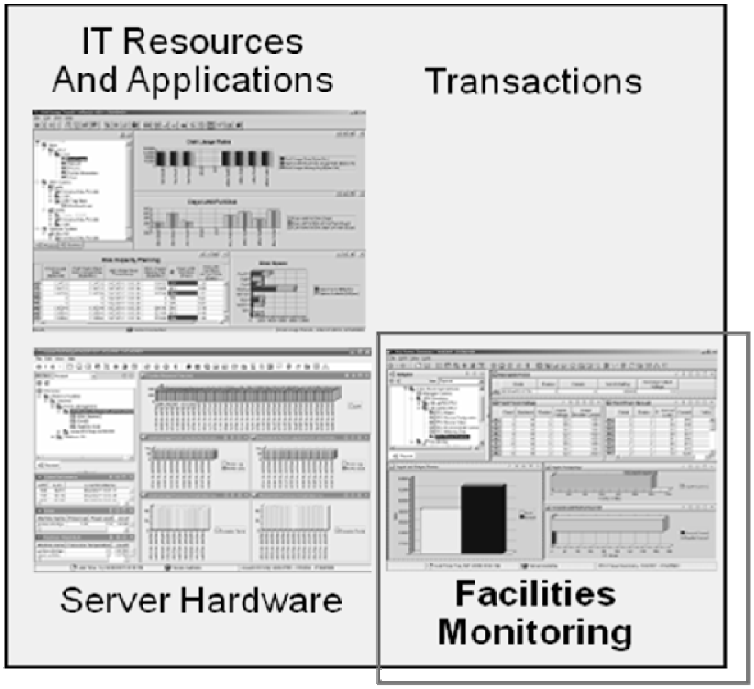

With a converged view of IT and Facilities, managers can make decisions that intelligently optimize the costs, constraints and realities of both functions.

This convergence is coming both from the IT side and the Facilities side:

- *IT monitoring and control systems*, such as IBM Tivoli and Active Energy Manager®, now enable IT operators to provision power management on servers and monitor items such as power, temperature, humidity, air flow and access control. Network capability and performance is also monitored by the network operation center, which is typically hosted by IT.
- *Enterprise-wide power monitoring systems*, such as Eaton's Power Xpert Architecture, enable IT, financial and facilities executives to see all their critical energy-related assets—primary switchgear, generators, building management systems, cooling, access control, UPS, branch circuit monitoring, etc.—in a single dashboard rather than in ad hoc monitoring systems. This alone saves valuable time when diagnosing power incidents or evaluating trends or efficiencies.

The ideal is a high-speed *hypervisor*—an umbrella view that correlates data from many monitoring systems—that enables both teams to collaborate to make the best decisions across the organization.

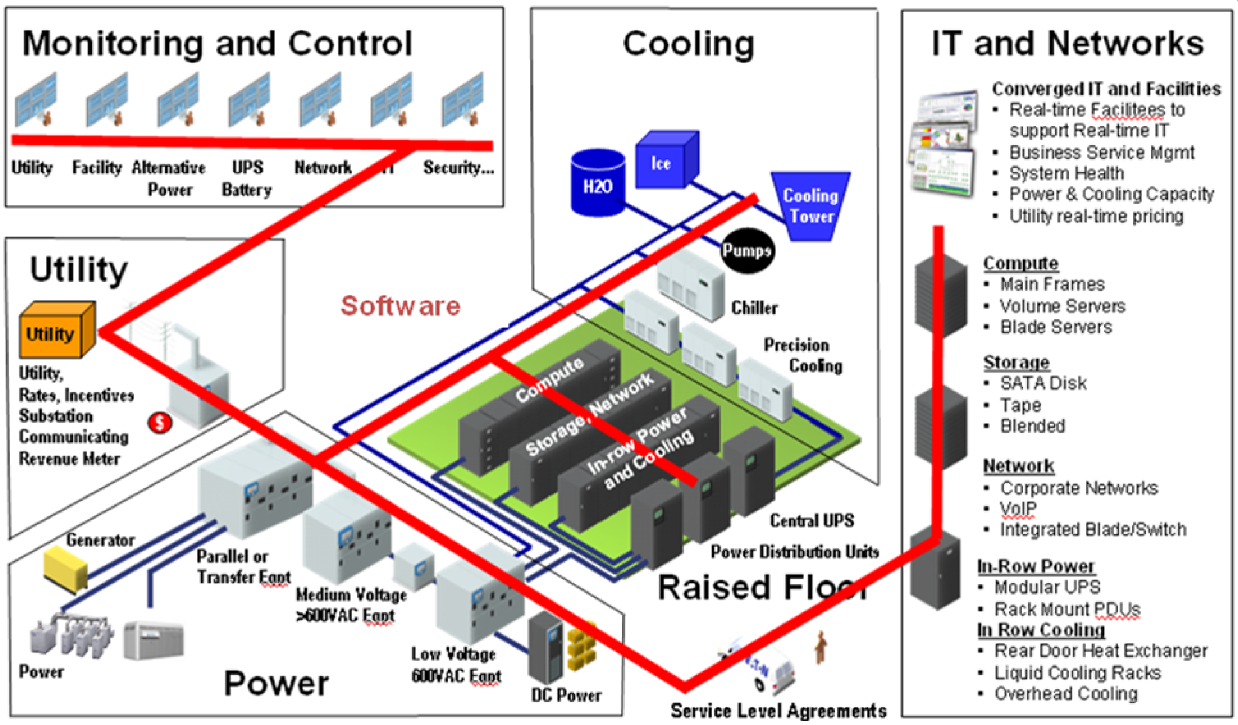

With real-time information from separate monitoring systems, organizations can better understand the current health of the IT and Facilities infrastructure, optimize it as a cohesive unit, and proactively address issues or opportunities on the horizon.

The system provides a high-speed scan of hundreds or thousands of monitoring points, with customized views for each authorized user. With a unified monitoring system in an open protocol environment, IT and Facilities gain a common language to address their mutual issues: availability, efficiency, SLA performance, and proactive diagnostics and analysis.

3. Know where you are and where you're going.

Many data center managers don't know the efficiency of their IT equipment or the site infrastructure—or have a clear path in mind for maintaining and improving that efficiency. There's a lot of low-hanging fruit being overlooked, readily available opportunities to substantially reduce energy costs and become "greener" in the process.

How much of the data center power budget goes to IT systems, and how much to support systems such as climate control, security and power distribution? For every kilowatt-hour of power being fed to IT systems, how much real IT output do you get? The answers to these questions provide a picture of how much power is consumed for every unit of real business productivity, such as Web pages served, transactions processed or network traffic handled.

Although there are no true industry benchmarks for IT efficiency, there are some industry benchmarks for *site* infrastructure efficiency. The non-profit organization, The Green Grid (www.thegreengrid.org) recommends a metric called the Power Usage Effectiveness (PUE) ratio—or its inverse, Data Center infrastructure Efficiency (DCiE)— where:

PUE = Total data center power in / total power used by IT equipment
DCiE = 1 / PUE

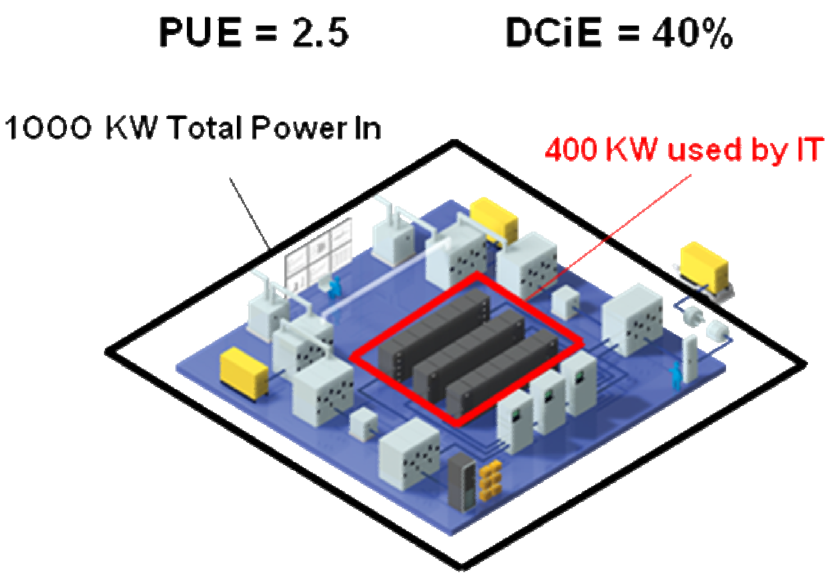

For example, a data center that consumed a total of 1000 kW, where the IT equipment consumed 400 kW, would be said to have a PUE of 2.5 and a DCiE of 40 percent.

Naturally, PUE and DCiE will vary by data center tier rating, application, geography and weather conditions. For instance, a Web hosting center would have a more favorable PUE (a lower number) than the data center supporting a trading floor. A Tier IV data center would have a higher PUE (lower efficiency) than a Tier II design, due to added redundancy.

The ideal would be a PUE of 1.6, but any well-designed and operated data center could realistically achieve a PUE of 2 while meeting business objectives.

Data center	PUE	DCiE
Microsoft goal	1.125	89%
Ideal PUE	1.6	63%
Target PUE	2.0	50%
Typical PUE	2.4 to 3.0+	40-33%

If you know where you start on this metric, you can track efficiency over time and reveal opportunities to maximize IT output while lowering input power—such as using server power management, server virtualization and blade servers—and where you can reduce losses and inefficiency in support systems.

4. Demand more real power from your UPSs.

Advances in UPS technologies have greatly improved the efficiency of these systems. In the 1980s, most UPSs used silicon-controlled rectifier (SCR) technology that operated at a low switching frequency and were 75 to 80 percent efficient at best. With the advent of new isolated gate bipolar transistor (IGBT) switching devices in the 1990s, switching frequency increased, power conversion losses decreased accordingly, and UPSs could run at 85 to 90 percent efficiency.

When even higher-speed switches became available, there was no need for UPS solutions to include transformers, which helped boost efficiency to 90 to 94 percent. The new Eaton BladeUPS™—optimized for today's IT equipment power supplies—operates at 97 percent efficiency.

Don't take vendors' efficiency ratings at face value. When evaluating a UPS, it's not enough to know the peak efficiency rating it can deliver at full load (the efficiency figure usually given). You are unlikely to be operating the UPS under full load. Since so many IT systems use dual power sources for redundancy, the typical data center loads its UPSs at less than 50 percent capacity, as little as 20 to 40 percent in some cases. You would expect efficiency to be lower when the UPS is operated at partial loads, but to what degree?

Previous-generation UPSs (those bought before 1990) are markedly less efficient at low loads. Even most of today's UPSs are noticeably less efficient under the low loads typically expected of them. For the very energy-conscious IT manager, the new Eaton BladeUPS system offers a very high efficiency profile of 95+ percent, all the way down to 20 percent loading.

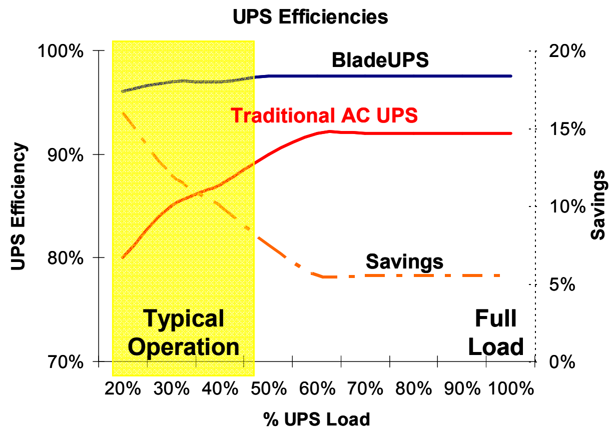

With dual and triple-corded loads fed from redundant power supplies, today's UPSs are typically loaded at 50 percent or less. When comparing efficiency statistics, be sure to compare UPS performance at lighter loads, not full load.

Even small increases in UPS efficiency can quickly translate into thousands of dollars, in more real power and lower cooling costs. In a one megawatt data center, a 10-year-old UPS is probably wasting about 150 kW of power and dissipating a lot of heat. Replacing that vintage equipment with new, high-efficiency UPSs can free up about 120 kW of that power to support new IT equipment and reduce the burden on cooling systems.

"Replacement of aging equipment with alternatives that consume less power may be more viable than you might think, when long-term energy pricing is factored into the total cost of ownership (TCO) evaluation."

Stephen Prentice, in *Why Cool Is Now 'Hot' for IT Planners*
Gartner RAS Core Research Note G00138041, May 2006

---

5. Consider costs in total context.

It's easy to look at the one-time capital cost of a piece of power infrastructure, and overlook more obscure factors that can drive up total cost of ownership. For a UPS, for example:

How efficient is the UPS? Line-interactive or standby UPSs operate very efficiently, but they do not protect against all power anomalies, such as harmonics and frequency variations. Double-conversion UPSs protect against the full gamut of common power problems, but at a cost. They're less efficient, especially at the light loads that are common in dual-bus data centers.

This trade-off is not an either/or proposition any more. With Eaton's unique energy-saving topology, the UPS operates in efficient line-interactive mode unless power conditions warrant a switch to the higher protective level of double-conversion mode.

What kind of power performance does the UPS deliver? If you choose a line-interactive topology for higher efficiency, will you need input filters to address total harmonic distortion? Some vendors advertise an attractive efficiency rating, but then they must use additional input filters to achieve the stated THD performance, which drops their stated efficiency and may take up additional floor space. If the difference in efficiency is even one percentage point, at 500 kVA that translates into more than $10,000 a year.

Manufacturer	kVA	Full load efficiency	Output power factor	THD at full load
Vendor A	500	94.3 percent	0.9	<5 percent
Vendor B	500	94.0	0.9	7 percent

Vendor C	500	93.8	0.9	7 percent
Vendor D	500	93.6	0.9	6 percent
Vendor E	500	93.0	0.9	7 percent

Power performance must be viewed as the combination as system efficiency, output power factor and input THD.

How costly will it be to deliver and install the UPS? For large data centers, large centralized UPS systems offer the highest availability—compared to smaller systems—as long as you can address delivery and installation issues. Will you be able to get the system into your freight elevator, or will you have to dismantle it or rent a crane to get it to the data center?

Manufacturer	kVA	Density (kW per ft^2)	Weight	Footprint (sq. foot)	Dimensions (WxDxH, in.)	Cross-country freight cost
Vendor A	550	29.6	2977	16.7	73.6x32.7x73.7	$893
Vendor B	500	24.1	6900	18.7	69.0x39.0x82.0	$2,070
Vendor C	500	24.0	5226	18.8	80.7x33.5x76.7	$1,568
Vendor D	500	23.9	5512	20.9	94.0x32.0x71.0	$1,654
Vendor E	500	23.1	5795	19.5	72.0x39.0x78.0	$1,739
Vendor F	600	20.6	6373	29.1	99.8x42.0x80.0	$1,912
Vendor G	500	15.1	6930	29.9	114.2x37.7x79.7	$2,079

A high-density system—delivering the power in less footprint—can be far less expensive to ship and occupy much less raised-floor space.

Will it require significant on-site assembly and cabling, or is the unit prefabricated and tested at the factory? In a prefabricated UPS, everything is pre-wired, so there's minimal cabling to be done at installation. Traditional multi-module installations involve three times more connections for inter-unit power and communications. Factor in all the installation costs and complexity.

Is there the opportunity for future expansion? You don't want to face a forklift upgrade in a few years because the system couldn't grow with the data center.

What hidden energy costs are being overlooked? Look beyond cost per kilowatt-hour. In the commercial and industrial worlds, you must contend with fluctuating rates based on peak demand periods and surcharges for poor power factor performance, usually caused by variable-speed motors.

The rate your organization pays will be based on the worst 15 minutes of the billing period, and that demand charge is typically carried forward for the next 11 months. Add in penalties for power factor, and the organization could be spending far more for energy than it should.

6. Build only what you need, and expand on demand.

"Building a whole data center to the highest energy point is expensive and a needless over-engineering exercise," wrote Rakesh Kumar and Philip Dawson in "*The data center as living organism: Why history is not a good guide to the future*" (Gartner RAS Core Research Note 153516, December 7, 2007). "The design principle in Generations 1 and 2 [mainframe and client/server environments] was around a single static structure. New designs need to be modular, with built-in expansion capabilities."

In the power world, modularity can mean UPSs that scale for added capacity or redundancy, extended battery modules to customize backup runtime, and plug-and-play power distribution components that break down room-level wiring into row- or rack-level modules.

For all these interpretations of modularity, the underlying concept is the same: provide a function or process in building-block fashion, and enable users to add, remove or redeploy those building blocks to create variations of the original function or process.

Modularity enables you to:

- Pay only for the functionality you need in the short-term.
- Plug in new capacity or functionality when the time is right.
- Expand at your own pace, without starting over with a new platform every time you want to grow or upgrade.

UPSs systems can meet this requirement with modular components that use *paralleling*. In paralleling, two or more UPSs are electrically and mechanically connected to form a unified system with one output—either for extra capacity or redundancy. As a conjoined system, each module stands ready to take over the load from another module whenever necessary, without disrupting protected loads.

Paralleling provides an excellent solution for matching growth while extending the value of existing UPSs. Paralleling is available on UPSs of widely varying sizes. For instance:

- **For large applications**, the Eaton 9395 UPS expands in 275 KVA modules up to 1.1 MVA. Distributed systems with prefabricated internal static switches can be paralleled up to approximately 4 MVA. UPS modules of unequal sizes can be paralleled up to 4MVA using a centralized bypass system.
- **For rackmount applications**, the Eaton BladeUPS system delivers 12KW of power protection in a 6U unit—expandable to 60kW N+1 in one 19”, 42U enclosure. Paralleling is accomplished using a BladeUPS Bar—a plug-and-play bus structure that mounts in the back of the equipment rack.

BladeUPS uses the same patented paralleling technology that Eaton developed for much larger systems, with the same philosophy, controls and algorithms. Each UPS module contains its own internal battery trays, built-in static switch, maintenance bypass capability and hot swappable components for ease of replacement or upgrade. Going modular does not have to compromise reliability.

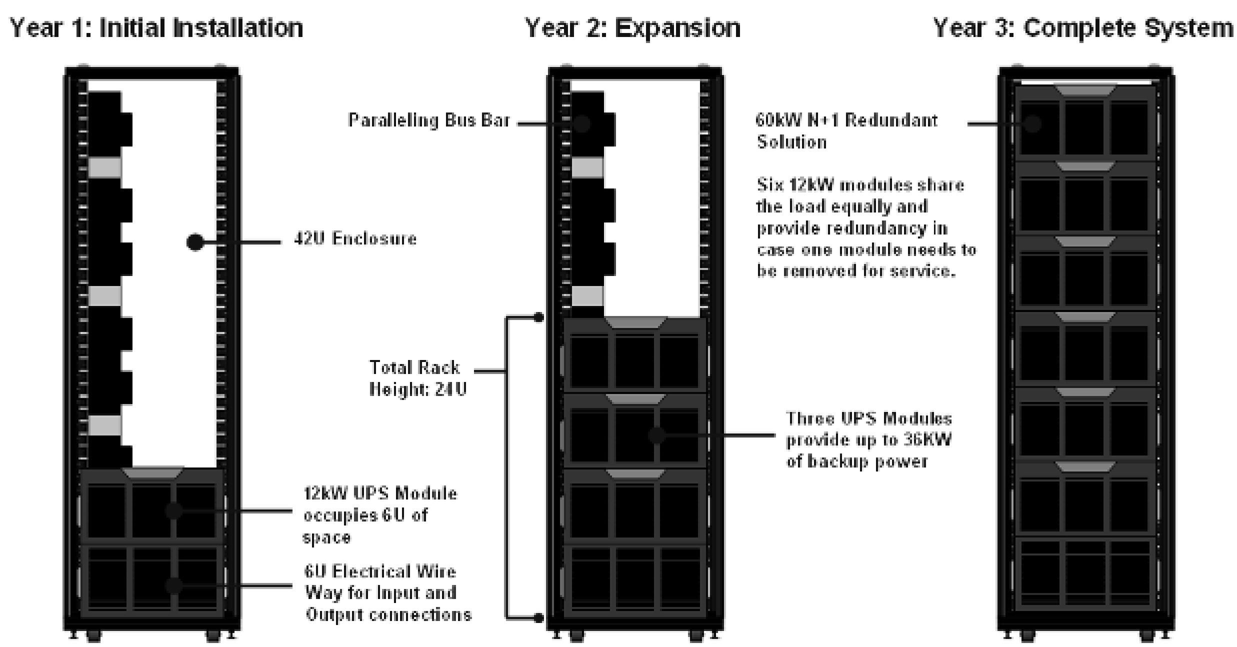

A modular UPS, such as the Eaton BladeUPS (pictured) expands in building-block increments. You can start with one 12 kW module and expand to a 60kW configuration (with an extra module for redundancy) as you add IT equipment and need more power.

Closing thoughts

In a greenfield data center or major expansion/upgrade of an existing data center:

- Expect cultural/process change, so you'll need to have executive-level sponsorship and form a cross-functional team to develop an energy strategy for IT operations.
- Include energy efficiency as a key requirement in design criteria, alongside expectations for capital cost, reliability and uptime.
- Consider energy efficiency in calculations of total cost of ownership when selecting new IT, power and cooling equipment.

With best practices and the right choice of equipment, data center managers can reduce energy consumption by nearly 50 percent. That means that almost half of the power utility bill will fuel actual IT processing, compared to less than 25 percent of the power supplied to a nominal data center today.

With a more efficient allocation of power, you will not only reduce utility bills and total operating cost, but also achieve more with available backup power and cooling systems – delaying the point where those systems would have to be upgraded to match data center expansion.

For more information

Download the following free white papers from www.powerware.com:

- **Next-Generation Power Quality Meters**
- **Who Tripped the Circuit?—***The clear-cut case for branch circuit monitoring in data centers*
- **Do IT and Facilities Speak to Each Other?***—Collaboration between these groups offers advantages to each—and has dramatic impacts on overall business performance.*
- **How to Select the Right UPS for Any Application***—Eight factors to determine the optimum match of UPS model, options and power distribution strategy*
- **Build for Today. Expand on demand.** *Modularity in the data center power infrastructure*
- **Top 10 Ways to Save Energy in Your Data Center***—Even a small data center can save tens of thousands of dollars through smart choices in equipment and best practices.*
- **Increase the efficiency of power distribution in your high-density data center**
- **The Green Imperative***—Power and cooling in the data center*
- **Is Power Your Weakest Link in Data Center Flexibility?** *Key considerations for power systems in adaptive IT environments*

About Eaton

Eaton Corporation is a diversified power management company with 2007 sales of $13 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.

About the author

Dr. Kenneth Uhlman, PE, is the director of data center business development for Eaton Corporation, where he is responsible for Eaton's global data center strategy. He focuses on improving efficiency, availability and business service management for data centers, including the convergence of IT and Facilities.

He holds three U.S. patents, is a licensed Professional Engineer in California and has been honored with many awards, including Eaton's Pinnacle Award and GE's Musketeer Award.

He holds a doctorate in Organizational Development from the University of Phoenix School of Advanced Studies and a BSEEE from North Dakota State University. His dissertation was titled, "Corporate transformations and collaborative partnerships in mission critical facilities."

You can reach him at KennethLUhlman@Eaton.com.