Sponsored by APC by Schneider Electric

DATA CENTER EFFICIENCY

Saving money on data center power and cooling remains a top concern for companies of all sizes. Smaller organizations in particular are challenged to find the most energy-efficient solutions, due to restricted budgets and limited staff. What's more, they don't always have the proper space to dedicate to their IT infrastructure, and are often left retrofitting locations designed for other uses to become server rooms or small data centers. Without the resources of large organizations that break new ground in power and cooling techniques, smaller organizations need to find solutions that fit their size and budget.

In this eGuide, *Network World* and sister publications *InfoWorld, PC World* and *Computerworld* explore some recent trends in power and cooling efficiency, offer advice to small organizations and examine how modern-day data centers attempt to tackle these problems.

IN THIS eGUIDE

DATA CENTERS SHOW SIGNS OF ‘GREEN FATIGUE’

By James Niccolai, IDG News Service

The decline in interest is more pronounced at smaller data centers, which tend to have fewer engineers and less money to devote to energy efficiency projects.

A new survey from the Uptime Institute suggests fatigue is setting in when it comes to making data centers greener, and it may be partly due to overachievers like Google and Microsoft.

In the Institute’s latest survey of data center operators, only 50 percent of respondents in North America said they considered energy efficiency to be very important to their companies. That was down from 52 percent last year and 58 percent in 2011, and is despite a constant drumbeat of encouragement to reduce energy costs and cut carbon emissions. The decline in interest was more pronounced at smaller data centers, which tend to have fewer engineers and less money to devote to energy efficiency projects, said Matt Stansberry, Uptime Institute’s director of content and publications.

“A lot of these green initiatives, like raising server inlet temperatures and installing variable-speed fans, are seen as somewhat risky, and they’re not something you do unless you have a bunch of engineers on staff,” he said.

But there may be other factors at work. Stansberry suspects that managers at smaller data centers are simply fed up hearing about success stories from bleeding-edge technology companies such as Google, and their survey responses may reflect frustration at their inability to keep up.

“I don’t really think that half the data centers in the U.S. aren’t focused on energy efficiency, I think they’re just sick of hearing about it,” he said. “You’ve got all these big companies with brilliant engineers and scads of money, and then there’s some guy with a bunch of old hardware sitting around thinking ‘What the hell am I supposed to do?’”

The gap in enthusiasm reflects a divide between large and small data centers that is apparent in other areas too. Data centers with more than 5,000 serv-

ENOUGH ALREADY?

"I don't really think that half the data centers in the U.S. aren't focused on energy efficiency, I think they're just sick of hearing about it."
— Matt Stansberry, director of content and publications, Uptime Institute

ers are far more likely to have invested in new infrastructure and expansion projects, he said. Smaller data centers, meanwhile, are maintaining existing facilities and moving more work to online service providers and collocation facilities.

Those service providers and colos are also the ones investing the most in going green, he said. IT energy costs make up a big part of their overall operating costs, so "every penny they save is profit," he said. Other companies, including big retailers and manufacturers, see less incentive to improve efficiency. For some, reliability and security are a bigger priority.

There was also a division along geographic lines. Interest in green IT was higher in Asia and higher still in Latin America—particularly Brazil, where energy costs are high. The interest may be lower in the U.S. partly because energy here is relatively cheap, Stansberry said.

The survey was completed by about 1,000 respondents at data centers worldwide, though predominantly in the U.S. Nearly half the respondents manage three data centers or more, and they were a mix of facilities staff, IT staff and senior executives responsible for both areas.

Other results show that building data centers in a modular fashion has been slow to catch on. The survey defines modular equipment as that which is manufactured off-site and delivered ready for installation, which allows mechanical and cooling equipment to be added in stages to match the IT load. Less than 1 in 10 respondents said they were using prefabricated, modular components, and more than half said they had no interest in doing so.

HOW TO RUN YOUR BUSINESS ON FOUR RACK UNITS

By Paul Venezia, InfoWorld

With a little hardware and a lot of virtualization, you can keep 50 to 100 users happy—and you don't even need much horsepower.

For the past eight months or so, I've been running a de facto virtualization test in the lab. I took a variety of lab boxes and internal servers, virtualized them on VMware vSphere 4.1, and ran them all on two physical servers. It may not seem impressive, but it indicates what's possible for small to medium-size corporate computing. And for those looking to consolidate, it's highly relevant.

These are not high-powered boxes. In fact, both servers fit into a single rack unit, using Tyan's YR190-B8228 side-by-side dual-socket server. Within that single rack unit, there are a total of four physical processors: two AMD 4162 EE CPUs with six cores apiece running at 1.7GHz, and two AMD 4170 HE CPUs in the other half of the box, with six cores each running at 2.1GHz. With 32GB of RAM in each side, that's enough horsepower for a surprising array of VMs and services.

The storage is a Synology RS3411RPxs: a 10-spindle SATA array with redundant power and multiple 1G and 10G links, though only the 1G interfaces are in use for this particular experiment. There's a QNAP TS-EC1279U-RP in play to address other tasks, but it also contains backups of the whole shebang, replicated every few hours.

Each physical server has two gigabit links to storage and a single gigabit link to the front-end network – small potatoes indeed. However, the services running on this cluster handle an equivalent load to a moderately sized business (50 to 100 people). Naturally, this is an estimated workload, but these boxes get pushed hard during the day and continue on with various processes in the evening.

There are domain controllers, the vSphere vCenter server, internal mail servers, external mail relays, file servers, backup servers, application servers, database servers, and even a PBX server running Asterisk that handles all the phones, both analog and SIP. It's not tasked with highly transactional databases hit by thousands of queries a second, but it's a relatively busy little data center in a box. Well, three boxes – including the switch it takes up all of four rack units.

In looking at the statistics over time, it's clear there are no bottlenecks anywhere for the load past and present. During normal operation, the CPUs on each server hover around 10 to 15 percent utilization with spikes to 40 percent or so when one VM gets a little more action. RAM utilization averages out around 25 percent per box; if one unit fails, there's easily room for all those VMs to restart on the other side. VMware's DRS handles the load balancing, and HA handles any hardware failure event. In fact, the only thing that could be added might be more RAM if the VM count increases to any substantial degree.

Looking at this data center microcosm, we see I'm easily running a varied workload on low-power, low-cost chips, servers, and storage, and I have plenty of room to grow. If anything, this is where smaller infrastructures need to be concentrating.

When planning and specifying servers for new virtualization projects, it's enticing to go with the biggest, baddest, fastest CPUs in an effort to "future-proof" the resulting infrastructure. The theory is that if you get the newest, fastest chips available, they'll have a longer usable lifespan than if you go with cheaper, slower CPUs. However, the returns on that investment may not always be what you think.

The fact is a significant number of virtualized infrastructures use far less CPU resources than anything else. They chew up RAM, but RAM is cheap and easily added. Spending money on RAM and faster disk is almost always going to provide better overall results than dropping several thousand dollars per CPU on the fastest chips on the market.

There's always a question of redundancy as well. Having only two servers running an entire company's server infrastructure may seem precarious to many, but the reality is that with the load-balancing and high-availability features in modern virtualization solutions, both servers would have to fail at the same time to actually take down everything. If it's that much of a concern, you can always add one or two servers later on instead of just increasing the RAM in the existing boxes. The primary storage in this case may not have redundant controllers, but it does have redundant power supplies and is easily backed up to the QNAP box through the use of simple rsync.

This solution isn't tuned or spec'd for anything beyond normal small- to medium-sized-business computing needs. However, the fact that the cost of all of these components is well under $20,000 and can support a significant number of users is certainly food for thought.

WHAT IT SHOULD KNOW ABOUT AC POWER

By Paul Venezia, InfoWorld

After all, it's just electricity, right? How hard could that be? Surprisingly hard—if you don't know what you're doing. Here's the lowdown.

If there's one thing people tend to screw up in medium-size data center buildouts, it's AC power specification. Even some of the brightest network and server administrators seem to think there's a limitless number of power outlets in any given rack or room, and if there isn't, well, you can just plug a $3.99 power strip into the UPS and be done with it, right?

Technically, they're right. You can plug that $3.99 power strip into a UPS and run a few servers or small switches off it, even if it pushes the UPS near overload. That doesn't mean it's a good idea or even a safe one. But it will probably work for a time. Oddly, I've found that a few of these same folks bristle mightily when they discover that a user has brought in a $19.99 five-port 10/100 switch to plug into the network jack in their cube – pot, kettle, and so forth.

The power vacuum

The sad fact is that far too many good, solid network and system folks simply don't know much about power. They see an outlet and assume there's juice. Occasionally they get upset when they can't find an outlet that will accept a NEMA 5-20 plug (with one horizontal blade), because the outlet's only 15 amp and has a 5-15 receptacle. I've seen many situations where someone has gone to Home Depot and built a cable with a 5-20 receptacle on one end and a 5-15 plug on the other to "adapt" the power cable. This is known as a bad idea.

Sure, there are plenty of power-savvy IT folks out there, but the trend is toward letting the ephemeral "someone else" deal with the problem. The "someone else" usually turns out to be a UPS and/or generator vendor or a local electrician. That's where the fun begins. It can be like an IT version of "Who's on First?" We could call it "How Many Watts?" or "What's an Amp Again?"

When power guys who know a little about IT and IT guys who know a little about power get together, there's a reasonably good chance that things will turn out OK after some tweaking. When power guys who know nothing about IT get together with IT guys who know less than a little about power, that's generally a recipe for disaster. This is how you wind up with wildly over- or under-spec'd in-rack power situations.

An AC power primer

Let's take a brief look at power, centered around IT needs. (Note: The following is greatly simplified in places. Do not take this as canonical; your service and situation may differ greatly depending on location and a variety of other factors.)

We'll start with the mains. You probably have three-phase AC service to your location. You don't really care about what's coming in unless you own the building outright, but you do care about delivering that power to your IT gear. If you have a generator, then there's a relay panel in the middle of that service, ready to cut over to the generator should main power fail. Thus, everything you place on the load side of that panel must reside below the operational limit of the generator. If you have a 50kW generator, don't expect to be able to run 65kW of equipment, and so on.

If you have a substantial amount of gear and want to use a substantial UPS to protect it, you'll probably wind up bringing three-phase power directly to a PDU and UPS in the room. These two pieces will handle the switching to/from the UPS during a power outage and will be used to distribute power to the rest of the room.

Using three-phase gear can also greatly simplify rack power distribution because three phase basically means that there are three 120v circuits delivered to each rack, generally within a single vertical PDU. By combining two of those circuits, you can run gear that requires 208v, such as beefy switch and blade chassis power supplies. However, within the same PDU, you can also spread out those three circuits into individual 120v outlets for 120v gear, such as standard 1U servers and the like. Without three-phase delivery, your PDUs are generally only 120v or 208v, and unless they have step-down transformers or you decide to make the switch and run all your gear at 208v, you'll need separate PDUs for each set of gear.

Regardless of whether you can use three-phase service, you'll need to be familiar with the range of NEMA plug types and what they mean. For instance, you're certainly familiar with the standard NEMA 5-15 plug that the vast majority of all electronic gear uses. This means it's suited to be plugged into a 120v 15-amp circuit. If that plug is a 5-20, it needs at least a 20-amp circuit. If the designation includes an "L" before the first number, it's a round, locking plug type, such as a NEMA L5-20. That type would again denote that the device requires 120v 20-amp service. If the first number is a 6, as in a NEMA L6-20, then we're talking a maximum of 250v and 20 amps. If it's L6-30, it's 250v and 30 amps. An L6-50? You guessed it, a 250v 50-amp circuit.

There are also IEC 60320 plugs and receptables. You're already intimately familiar with the C13/C14 coupler – you probably have a half-dozen within reach right now with a NEMA 5-15 on the other end. Some larger gear such as a blade chassis will use C19/C20 couplers on the back of the gear and offer a variety of cable options to connect to a number of service types. A common example is a C19-to-L6-20 cable that will deliver 208v 20-amp service to a particular power supply. These are usually found on blade chassis, big storage, and large switching chassis. There are many additional NEMA plug types, but these are some of the more common types used in IT.

There's an entire discussion to be had on how to gauge your current and future power consumption, wattage calculations and so forth, but that's a discussion all unto its own.

Before embarking on a new data center design or even before modifying or updating an existing data center, spending a little time researching the power situation at your facility and power delivery in general. It's neat stuff, and adding that knowledge to your arsenal can only help you now and in the long run.

DATA CENTERS ARE BIG POLLUTERS, REPORT SAYS

By Christina Desmarais, PC World

Data centers across the world use about 30 billion watts of electricity, about the same as the output of 30 nuclear power plants, with digital warehouses in the U.S. accounting for one-quarter to one-third of that load.

Data centers across the world use about 30 billion watts of electricity, about the same as the output of 30 nuclear power plants, with digital warehouses in the U.S. accounting for one-quarter to one-third of that load, The New York Times reports.

Those are huge numbers for an industry that often puts forth an image of environmental friendliness. Consider Apple's soon-to-be-built research and development complex near its Cupertino, California, headquarters. The building, which will look something like a spaceship, will be entirely surrounded by a thick layer of trees and will be powered with its own energy center that will run mostly off the grid.

Yet in a year-long investigation, The Times found that "most data centers, by design, consume vast amounts of energy in an incongruously wasteful manner." It adds, "Online companies typically run their facilities at maximum capacity around the clock, whatever the demand. As a result, data centers can waste 90 percent or more of the electricity they pull off the grid."

In addition, to prevent power failure, they rely on vast numbers of generators that spew diesel exhaust in amounts that often violate clean air regulations. In fact, many Silicon Valley data centers appear on California's Toxic Air Contaminant Inventory, a list of the biggest stationary polluters.

"It's staggering for most people, even people in the industry, to understand the numbers, the sheer size of these systems," Peter Gross, a man who has helped design hundreds of data centers, told The Times. "A single data center can take more power than a medium-size town."

Not all are environmental abusers. For example, eBay's 245,000 square-foot data center in Delta, Utah, is LEED Gold certified and is 50 percent cheaper to operate and 30 percent more efficient than previous eBay facilities, partly due to the 400,000-gallon water cooling cistern that collects rainwater and can keep the building cool for 7000 hours without drawing any electrical current.

GREEN IT: IN SEARCH OF AN ENERGY YARDSTICK

By Mary Brandel, Computerworld

New green-IT metrics help measure data center efficiency, but each has its strengths and weaknesses.

Last fall, Google made an announcement that in many ways foretells the future of data center efficiency metrics. The search giant not only disclosed its total power consumption and carbon emissions (mostly attributable to its data centers), but also released estimates of its per-user and per-search energy consumption. With this information—and given that a billion Google searches occur per day—it was possible to calculate that searches account for 12.5 million watts of the company's 260 million watt total.

The idea of quantifying the kilowatts of energy required to perform a useful unit of work is now considered the holy grail of data center metrics. The most widely used metric today—The Green Grid's Power Usage Effectiveness (PUE) measure—compares the amount of data center energy consumed by IT equipment to the facility's total usage; but no metric yet reveals energy consumed per unit of work, mostly because it is difficult to define a standard "unit of work."

"For cars, it's miles per gallon, but for hardware, is it a MIP or bits transmitted per second? The industry hasn't zeroed in on what 'work' is," says Steve Press, executive director of data center facilities at Kaiser Permanente. "PUE doesn't tell you what work the data center is doing—it could be blocking virus traffic or moving data for the National Weather [Service]. Once we settle on [that], it will have a big effect on overall competitiveness."

"If you're a bank or credit card company, your unit of useful IT work is a transaction, whereas if you're a Web company, it's Web pages served or ad click-throughs," says David Ohara, a green data center consultant who runs the Green Data Center Blog. "Ideally, there should be one fairly homogeneous definition."

The Green Grid and other industry groups are working on metrics that address the measurement of productive energy consumption, but issues remain. John Tuccillo, chairman of the board and president at The Green Grid, says, "Once that kilowatt gets to the rack, blade, processor or processor core, how can you then quantify how that unit of energy is being used for the intended purpose? That's the holy grail."

Meanwhile, companies are using combinations of available metrics, as well as coming up with their own metrics, to paint a picture of data center efficiency and productivity. "Overall, the industry struggles to create practical measures that provide the comprehensive insight that steers appropriate investment," says Simon Mingay, an analyst at Gartner.

With all this in mind, the chart below depicts a compilation of metrics that data centers are using to calculate their energy efficiency.

DEFINE WORK

"For cars, it's miles per gallon, but for hardware, is it a MIP or bits transmitted per second? The industry hasn't zeroed in on what 'work' is."
— Steve Press, executive director of data center facilities, Kaiser Permanente

Power Usage Effectiveness (PUE)

PUE is by far the most widely used data center efficiency metric. Developed by The Green Grid, it illustrates how much energy in the data center is allocated to IT equipment – servers, networks and storage – and how much is used for cooling, lighting and power equipment. Its popularity is attributable in part to its simplicity, both in concept and in application.

PUE is calculated by taking total facility energy (cooling, lighting, power and IT) and dividing it by IT equipment energy. A "perfect" PUE score, then, would be 1. According to The Green Grid benchmarks, a PUE of 2.2 indicates "needs improvement," while 1 to 1.4 is best in class. Some organizations prefer to use the Data Center Infrastructure Efficiency (DCiE), metric, which is the inverse of the PUE ratio.

"While not perfect, PUE does a good job of achieving a snapshot of how much electricity is powering what the data center is there to do," says Doug Washburn, an analyst at Forrester

CADE Corporate Average Data Center Efficiency

Developed jointly by the Uptime Institute and McKinsey, CADE is a set of four metrics intended to rate the business performance of a single data center or the weighted average performance of several data centers.

1. **Facility asset utilization:** Measures current IT energy use, divided by maximum IT energy use.
2. **Energy efficiency:** Measures current IT energy use, divided by total facility energy use.
3. **IT asset utilization:** Measures average server CPU utilization.
4. **IT energy efficiency:** Measures useful IT work, divided by watts. Since industrywide definitions of useful IT work are still under development, CADE uses a baseline value of 5%.

According to Uptime, CADE helps quantify how much of a data center's expenditures are devoted to information processing by separating energy efficiency from capital asset utilization and IT from facilities.

While the actual definitions of the four metrics could use some fleshing out, John Stanley, research analyst at The 451 Group, says CADE is a useful framework for understanding the issues of assessing facility efficiency and utilization.

Research. "It shows you where the biggest energy consumers are, and you can take steps to make changes from there."

In March, The Green Grid released more explicit guidelines to help standardize PUE data collection. However, limitations remain. For instance, while it can be helpful to compare your PUE with an industry average, it can be misleading as a comparison benchmark, given that no two data centers are alike, says John Hickox, U.S. practice leader of climate change and sustainability services at KPMG. "Drawing a reasonable comparison can be difficult and set unrealistic expectations as to what a target state performance level might or should be," he says.

PUE also does not help monitor effective use of the power supplied, Mingay says. For example, if a data center manager began using virtualization to increase average server utilization, the resulting efficiency improvements would not be reflected in the PUE, he says.

Katherine Winkler, a board member and secretary at The Green Grid, agrees that PUE is not an end-all, be-all metric. Rather, like any metric, it should be used with a specific goal in mind, such as benchmarking, comparing alternate data center designs or tracking your own efficiency. "You might say it disincents consolidation and virtualization, but it's about the efficiency of your cooling and power infrastructure," she says. "If your PUE goes up [after server consolidation or virtualization], that's an indicator you don't have a resilient infrastructure that can shrink when you're shrinking demand," she says.

Most agree that PUE has a lot of value. "Simply getting the industry to start doing measurements is a great step forward," says Jack Glass, director of data center planning at Citigroup. It also helps bridge the gap between the IT and facilities departments. "It's a new collaboration that didn't exist before in many data centers," Winkler agrees.

"Despite its very well-publicized limitations and frequent misuse, PUE is a good measure when it is clearly understood by those using and interpreting it, and when it is measured in a consistent way," Mingay concludes.

Carbon Usage Effectiveness (CUE) and Water Usage Effectiveness (WUE)

CUE, recently released by The Green Grid, addresses data-center-specific carbon emissions. It is calculated by measuring total CO2 data center emissions and dividing by IT equipment energy. WUE, also a relatively new metric, measures the amount of water used for IT and is similarly calculated by measuring total water usage in the facility and dividing by IT equipment energy.

Winkler points out that both are still evolving. For instance, while a "perfect" CUE or WUE is 0.0—indicating that no carbon or water use is associated with data center operations—there are no current benchmarks defining "average" or "good" CUE or WUE levels.

Mingay cautions that for CUE to be useful, it needs a very tight definition, particularly in terms of the emissions factors that get used. "There is way too much opportunity to misrepresent this number," he says.

PUE and WUE quantify elements that the pure energy metrics miss, says John Stanley, research analyst at The 451 Group. Combining the three gives you a fuller view; for instance, Washburn says, you might have a very low PUE but a high CUE if your main power source is coal. Or you may have a low PUE and CUE but a high WUE if you rely on hydroelectricity or even more efficient cooling systems, which typically use more water, says David Schirmacher, chief strategy officer at Fieldview Solutions. "If you use one metric without the other, you can be fooling yourself," Winkler agrees.

Energy Reuse Factor (ERE)

Since PUE does not account for reusing excess energy from

AN IDEA WHOSE TIME HAS NOT YET COME

"I don't know of anyone using DCeP beyond perhaps a few members of The Green Grid or those closely associated. It's attempting to get to something very important, but it remains impractical for most."

— Simon Mingay, analyst, Gartner

the data center, The Green Grid introduced ERE earlier this year for companies that reuse energy for heating or cooling. ERE is calculated by adding up total facility energy, subtracting reused power and dividing by IT equipment energy. An ERE of 1.0 means you only need to supply the facility with an amount of power equal to 100% of the IT energy. Because of its relative newness, ERE is not widely used at this time.

Server Compute Efficiency (ScE) and Data Center Compute Efficiency (DCcE)

The Green Grid's ScE is intended to look beyond CPU utilization to CPU efficiency. It does this by expressing the fraction of time during which a server is providing "primary services" – the services it is intended to provide – versus "secondary" services, such as virus scans, security updates and hard drive defragmentations. Data Center Compute Efficiency (DCcE) aggregates ScE across all servers in the facility.

For example, if an email server is processing emails 75% of the time, then its ScE would be 75%, Mingay says. "It looks at not just 'Is your system doing something?' but 'Is it doing something useful?' " Stanley adds. The idea is to identify underutilized servers and apply virtualization, consolidation or decomissioning strategies to increase efficiency.

The Green Grid emphasizes that these metrics are not productivity metrics—they don't measure how much work is done, but rather what proportion of work is useful. That distinction is important, Mingay says, since a server with a slow processor might spend an hour processing 100 transactions, while a server with a fast processor might complete 100,000 transactions in that time. However, both would get an identical ScE of 100% if 100% of their time during the hour was devoted to doing that work.

Data Center Energy Productivity (DCeP)

With DCeP, The Green Grid seeks to quantify the ratio of useful work produced by a data center to total energy consumed by it. Because the metric requires a definition of "useful work," which the industry has yet to define, "it's more of a framework," Stanley says.

In lieu of a definition for useful work, The Green Grid has produced a set of eight productivity proxies intended to be useful for estimating DCeP. Examples include weighted CPU utilization, operating system workload efficiency and bits transmitted per kilowatt per hour. None are perfect individually but will likely be used in combination, Winkler says.

For instance, "having traffic on the network is an indicator [that] the data center is doing a lot of stuff, but it could be transmitting garbage data and getting a great score on the metric even though it's not accomplishing anything," Stanley says.

So far, the DCeP has not had a lot of uptake by the industry, Mingay says. “I don’t know of anyone using this beyond perhaps a few members of The Green Grid or those closely associated,” he says. “It’s attempting to get to something very important, but it remains impractical for most.”

SPECpower

The Standard Performance Evaluation Corporation developed the SPECpower_ssj2008 metric to evaluate server power and performance by measuring server operations per second per watt. “They stress the CPU at different levels and measure the amount of work being done against the amount of power being consumed, and then they come up with an average number,” Schirmacher says.

A benefit of SPECpower is the large public data set for a variety of servers, Stanley says. On the downside, measurements are performed by the vendors themselves, which can achieve ratings beyond what’s possible in a traditional deployment, he says. Additionally, the SPEC definition of a Java unit of work could differ drastically from the machine’s intended use in the data center. “SPEC tried to come up with a workload that is broadly representative, and it exercises different parts of the system, but it’s still different from any specific workload,” Stanley says.

Something to keep in mind is that there isn’t a big drop-off in power consumption for servers running at low capacity, Schirmacher says, so even if you purchase servers with good ratings, you also need to ensure you’re achieving high utilization, which is measured by ScE.

Gartner Power to Performance Effectiveness (PPE)

Gartner developed PPE to analyze the effective use of power by existing IT equipment, relative to the performance of that equipment. With PPE, companies measure the average performance, energy consumption and utilization of their equipment and then compare that with optimal targets they have previously defined. PPE is represented, then, as actual power performance, divided by optimal power performance. “Pushing IT resources toward higher effective performance per kilowatt can have a twofold effect of improving energy consumption—putting energy to work—and extending the life of existing assets through increased throughput,” Mingay says.

ASHRAE’s Thermal Guidelines for Data Processing Environments

The American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) continually revises its data center temperature standards, revealing that facilities don’t need to be as cool as previously thought. In some cases, inlet temperatures can safely be boosted to 80 degrees Fahrenheit. While not a metric per se, the standards “provide the foundation for any changes you make that could improve the efficiency of the data center,” says Glass.

Data center efficiency metrics will continue to be an area of big concern, as companies look to cut costs and improve their sustainability ratings. “If you’re serious about energy management, you have to measure it,” Press says. “And you need to pick a measurement that’s right for you in terms of the granularity you want.”

“The metric in itself doesn’t have to be perfect,” agrees William J. Panian, facilities manager at PricewaterhouseCoopers. “A crooked stick is much better than no stick.”

DATA CENTER EFFICIENCY RESOURCES

Practical Options for Deploying IT Equipment in Small Server Rooms and Branch Offices

Small server rooms and branch offices are typically unorganized, unsecure, hot, unmonitored, and space constrained. These conditions can lead to system downtime or, at the very least, lead to “close calls” that get management’s attention. Practical experience with these problems reveals a short list of effective methods to improve the availability of IT operations within small server rooms and branch offices. This paper discusses making realistic improvements to power, cooling, racks, physical security, monitoring, and lighting. The focus of this paper is on small server rooms and branch offices with up to 10kW of IT load.

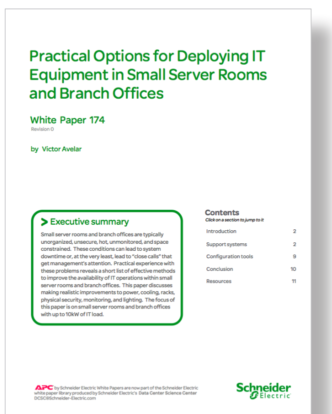

Learn more >>

Creating Order from Chaos in Data Centers and Server Rooms

Data center professionals can rid themselves of messy racks, sub-standard under floor air distribution, and cable sprawl with a minimum of heartache and expense. Whether the data center mess is created over years of mismanagement or whether the cable-choked data center is inherited, solutions for both quick fixes and longer term evolutionary changes exist. This paper outlines several innovative approaches for dealing with the symptoms of chaos and for eliminating the root causes of disorder.

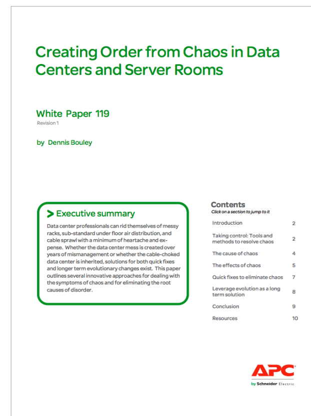

Learn more >>

IT Challenges Solved by Infrastructure for Small IT Environments

With the ease of adopting cloud computing, it is inevitable that some data and applications will eventually move there. Virtual environments are becoming easier to implement and manage, but they still rely on critical power and cooling support from a robust IT infrastructure. And sometimes, IT equipment needs to be deployed in creative ways in small or inadequate spaces. Read this eBook to learn how having the right infrastructure solutions to house, maintain, and secure your network and server equipment is essential to the well-being and continuance of your business.

Learn more >>