

Green Data Center Management

David A. Garbin and Elizabeth W. Chang

Data centers consume two percent of all electricity produced in the United States, which means that applying proven green practices to data center operations makes both environmental and economic good sense.

Recent changes in data consumption, such as video streaming, increased online transactions, and nearly paperless work environments, have created an explosive growth in data centers and data storage needs. Data centers represent a compelling opportunity for change because each data center consumes 10 to 100 times more energy than a regular building, depending on density and the center's cooling method. That consumption level plus the growing number of centers makes the greening of these buildings of paramount importance.

Indeed, between 2000 and 2005, energy consumed by data centers worldwide has doubled and the trend shows no signs of flattening.¹ In 2009, U.S. data centers alone will consume roughly 80 billion kilowatt hours (kWh)—two percent of all electricity produced in the United States.² At a time when greenhouse gas (GHG) emissions are a grave concern, these numbers present a clear call for immediate and informed action.

The data centers themselves are facing increased electricity costs and nearing their maximum power capacity. Some have already begun to implement measures to do more with less electricity using methods such as virtualization (multiple operating

systems on one or more machines), consolidation (eliminating underused or unused servers), and more sophisticated cooling methods. Centers are also beginning to recognize their responsibility to track energy use and reduce electronic waste.

As the sidebar “Why the Lack of Progress?” makes clear, a lack of guiding policy is the main reason that more data centers are not implementing energy-conservation measures. All too often, data center information technology (IT) professionals will implement one or two green practices they've heard about, yet they are never certain how much electricity was saved or how much cost savings such measures yielded. Analytics based on proven metrics combined with an effective monitoring plan will empower an organization to choose conservation methods that result in documented energy savings. Once the organization determines the most effective measures, it can create the necessary policies and guidelines.

First steps

Which method or combination of methods works best depends on the center. Organizations must first create an energy-consumption baseline using metrics to estimate energy at either the center or equipment level and then set up a monitoring system to track their energy use. The most important first step is to set up some kind of energy-monitoring system, whether that's electric meters or specialized software.

Using metrics

To measure energy use, organizations typically use one of two metrics. *Power use effectiveness* (PUE) is the ratio of the total power the data center consumes to the amount of power that only the center's IT equipment consumes:

$$\text{PUE} = \frac{\text{Total power going into data center building}}{\text{Power used for IT equipment}}$$

Inside Track

- Data centers are running out of capacity and need more efficient ways to power and cool equipment.
- All too often, a data center will implement one or two popular green practices, yet its decision makers will never be certain how much savings such measures yield.
- Energy reduction requires first creating a consumption baseline and then setting up a monitoring system to track exact energy savings.
- A center must manage every energy-reduction method, such as virtualization, consolidation, or improved cooling, with an eye toward its pros and cons.

The other metric, *data center infrastructure efficiency* (DCIE), is a reciprocal of PUE, usually expressed as a percentage. Google, whose data center practices are state of the art, claims to operate their data centers with an average PUE of 1.19.³ But typically, a data center has a PUE of around 2, which translates to a DCIE of 50 percent. In other words, a data center spends just as much energy cooling the IT equipment and building (50 percent) as it does to run the IT equipment (the remaining 50 percent).

The PUE (and consequently the DCIE) fluctuates with the outside temperature and changing work load. Obviously, a data center in a cooler climate will be easier to cool.

More standardization on these metrics will be necessary to compare data centers and to assess the significance of a data center's PUE. Accounting for weather fluctuations, in particular, is important when comparing annual energy use and computing savings.

Estimating consumption

The PUE and DCIE metrics are useful in determining an organization's baseline energy use as a prerequisite to implementing green measures. Ideally, the center would already have electric meters at multiple locations, which would make estimating a simple matter of reading the meters. However, sometimes a metering option is not available or is not practical to implement, such as when the data center is in a leased building, and the tenant is not permitted to alter the building infrastructure.

The data center described in the sidebar "Estimating Energy Use without Meters" on p. 30 is an example of when metering is impractical. A government agency wanted to get an idea of its organization-wide carbon footprint. Because the data center is just a portion of that, it made no sense to install electric meters just to get an accurate reading of only part of the carbon footprint.

To meet the challenge of estimating electricity without metering, Noblis developed a method for estimating a data center's energy use. With just an inventory of the data center's equipment, decision makers can derive an estimate on the basis of industry averages. Although metering would have to be installed eventually to get the specific consumption and be able to document the specific energy savings of any green effort, the method provides initial consumption estimates that are useful in getting a ballpark savings figure. This can often be sufficient motivation for management to take the next step in a greening effort.

Monitoring energy use

With accurate and systematic monitoring, a center can't hope to know the impact of any green initiatives. A monitoring system can be as simple as reading an electric meter outside the data center building or as complex as using a commercial software and hardware monitoring package to monitor electricity at the equipment level.

If metering is an option, meter placement is critical, particularly if the data center takes up only part of the building. It is

important to know exactly how much energy is expended in what part of the data center because getting the right numbers on the amount of energy saved is crucial in gaining direction and momentum in the green movement.

Products to monitor energy use down to the equipment level are available, such as IBM's Active Energy Manager and Aperture's VISTA. These programs should tie into the building's cooling systems to also measure the kWh of electricity expended in cooling the IT systems.

With monitoring results in hand, the center can determine if the cooling method is working, what systems are the most energy efficient, and what systems are not being used to their capacity. A center can even configure a monitoring system to perform special functions, such as turning off certain groups of systems at night.

Although the setup of either electric meters or monitoring programs takes planning and effort, the considerable return is that a center can measure exact electricity use in kWh and can thus know the true impact of any green initiative.

Energy conservation measures

Table 1 compares the methods available to data centers to begin conserving energy, including virtualization, consolidation, improved cooling methods, and the implementation of green acquisition and green building practices.

Why the Lack of Progress?

A survey by the Business Performance Management Forum on Green Information Technology (IT) identified the four most common reasons for not undertaking energy conservation measures (in order of prevalence):¹

- lack of policies to guide process,
- time and resources required to push the green agenda,
- too busy to formalize a program, and
- cost of infrastructure improvements.

The good news is that most data center managers realize the need to manage and conserve energy use. In another survey of 600 IT practitioners, only 16 percent responded with "not a priority" on reducing power consumption.² More than a third of the 600 respondents were not involved in paying the electric bills. Even so, many still recognized the importance of reducing power use.

References

1. R. Miller, *Google: Raise Your Data Center Temperature*, Oct. 14, 2008; <http://www.datacenterknowledge.com/archives/2008/10/14/google-raise-your-data-center-temperature/>.
2. M. Fontecchio, "Data Center Energy a Concern, But Metrics Lacking," *Data Center News*; Sept. 10, 2008; http://searchdatacenter.techtarget.com/news/article/0,289142,sid80_gci1329440,00.html.

Table 1. Comparing energy-conservation methods.

Method	Description	Pros	Cons
Virtualization	Multiple operating systems are set up on one computer or server	Users do not notice a difference	Setup cost
Consolidation	Minimize servers that sit idle by using fewer servers more actively	Fewer systems to maintain	Doesn't allow for extra machines to accommodate future capacity
Improved cooling	Techniques to cool IT equipment with less energy	No need to downsize capacity	Hard to measure impact without sufficient metering
Green acquisition	Through programs such as the Electronic Product Environmental Assessment Tool (EPEAT) and Energy Star, buyers buy products made of less toxic, more easily recycled materials	Encourages manufacturers to build more green products	Limits what you can buy
Green building	Using green building practices such as Leadership in Energy and Environmental Design (LEED) certification to build more energy efficient buildings	Energy savings; LEED prestige great for public relations	Can be implemented only when building new data center

Virtualization

Virtualization is having more than one operating system, such as Windows or Linux, on one server or personal computer (PC). The process is transparent to users, can drastically reduce the amount of IT equipment, and saves energy. Energy savings accrue not only from having less equipment but also from having less need to cool that equipment. With a PUE of 2.0, this savings is nearly equal.

PCs can be virtualized with the use of thin clients, which do not use as many resources as a PC. With a thin client, a user can log onto a server and see a screen identical to the screen that would appear during work on a laptop. The cost of thin clients has recently become comparable to that of a laptop or desktop, but thin clients have a longer life because they simply connect to the server and display a screen without processing or storing data.

Although, virtualization involves labor and software cost on the server front, there is no drawback once it is set up. Users notice no difference between logging onto their own dedicated machine or onto a virtual server. Cloud computing is an example of server virtualization that is gaining popularity. When someone goes to Google or Amazon.com to execute a search, there is no knowledge of where the servers that execute that query are located or if different servers are being used from query to query. Cloud computing is not an intentional green strategy, but it is arguably green because it allows fewer servers to complete the same amount of work. The sidebar “How Green Is Cloud Computing?” describes this concept in more detail and why it might be considered a green measure.

Consolidation

For many reasons, a typical data center carries numerous idle and underused machines. Data center administrators might prefer that each server have a specific function, for example, so they refrain from installing multiple programs on the same machine. A data center customer who is leasing space in a center nearing capacity usually leases more than the immediate need to leave room for expansion. This results in unused or underused machines that consume energy today to support an envisioned need later.

Consolidation is putting more programs on fewer machines and turning off or removing the remaining machines. It is an immediate and relatively inexpensive green initiative because, like virtualization, it cuts back on IT equipment without cutting back on the center's capabilities.

The main drawback is to the data center customers—organizations that lease from a data center business. These days, many data centers are nearing capacity. If an organization does lease additional capacity to accommodate future expansion and the data center doesn't have that capacity when the organization is ready for it, the organization will be forced to go elsewhere and either divide its data storage or move all its data to a larger center.

Improved cooling

Figure 1 shows the share of total energy that cooling occupies at various levels of green practice. On average, about half the energy expended in data centers is for cooling, which is a huge op-

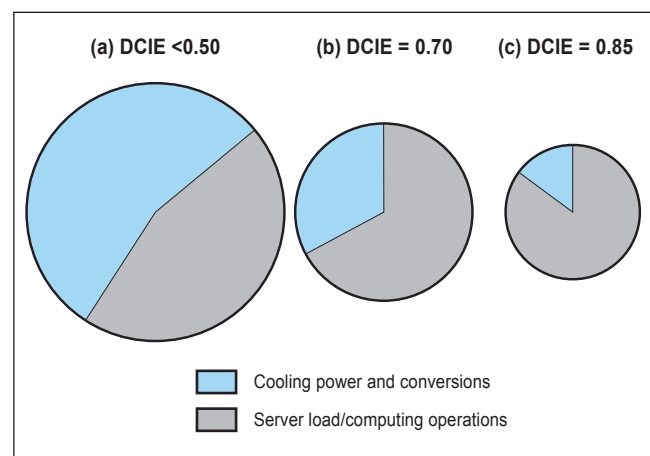


Figure 1. Cooling's contribution to energy use. Relative to the energy used for the server and computing operations, cooling and power conversion can occupy far less than it does at present. (a) In typical practice, the data center infrastructure efficiency (DCIE) is slightly less than 50 percent. (b) In better practice, it jumps to 70 percent, and in best practice, it reaches 85 percent. Figure based on an image by Paul Scheihing, U.S. Department of Energy; http://www1.eere.energy.gov/industry/saveenergynow/pdfs/doe_data_centers_presentation.pdf.

How Green Is Cloud Computing?

Cloud computing comes in several flavors, depending on the service offering or cloud type (public, private, or hybrid), but at its core is the notion that information technology (IT) resources are delivered on demand over the Internet, expanding and contracting according to need.

According to the National Institute of Standards and Technology, cloud computing has two key characteristics:¹

- the architecture surrounding massive clusters of computers is abstracted from the applications using it, and
- a software and server framework provides clients scalable utility computing capabilities.

The latter characteristic, usually achieved through virtualization, means that cloud computing can elastically provide many servers for a single software-as-a-service style application or host many such applications on a few servers. On-demand delivery of services allows organizations to pay according to use, scaling up and down as needed to reduce costs and energy.

Rackspace Hosting turned to cloud computing to even out its bouts of server activity and inactivity. The server would “sit there” underused and because Rackspace was charging by bandwidth, there was a clear opportunity to cut costs by going to cloud computing. The power demand rises only marginally when the server is working harder, so it made sense to have one server running full throttle all the time, getting more work for the same amount of energy. With cloud computing for e-mail hosting, for example, each Rackspace server runs about 2,000 e-mail boxes. In stark contrast, the average business runs 10 e-mail boxes per server.²

Cisco Systems avoided the cost of an entire new data center through cloud computing and virtualization. In the new setup, it was using 68 percent of its server arrays instead of 20 percent. “Going virtual” made it possible to delay building out a new data center, which translated to \$40 million annual savings.³

Some argue that cloud computing does not produce significant cost savings and that energy efficiency can be difficult to quantify. Certainly, not all services, information, and processes are candidates for cloud computing. Security requirements, performance needs, integration points, and existing dependencies must be considered before implementation. Cisco cites insufficient planning, weaker security and lack of industry standards as “minefields” to virtualization and eventually to cloud computing.³ Third-party providers, such as Google and Amazon need to consider how to measure and report on computing resources in the cloud to accurately measure efficiency, replacement costs, and cost savings to the consumer.

References

1. *Perspectives on Cloud Computing and Standards*, National Institute of Standards and Technology Information Technology Laboratory, Dec. 2008; http://csrc.nist.gov/groups/SMA/ispab/documents/minutes/2008-12/cloud-computing-standards_ISPAB-Dec2008_P-Mell.pdf.
2. T. Chan, *Green Practices in the Real World: Rackspace Hosting's Take on Energy Efficiency*, Oct. 2008; <http://www.greentelecomlive.com/2008/10/15/green-practices-in-the-real-world-%E2%80%93-rackspace-hosting%E2%80%99s-take-on-energy-efficiency/>.
3. A. Hickey, *Interop: Cisco Takes On Virtualization, Cloud Computing*, ChannelWeb, Sept. 17, 2008; <http://www.crn.com/networking/210602222>.

portunity for greening. A best practice example is Google's data center, with its PUE of 1.19. Consequently, for every kWh drawn for the equipment, a little over 0.2 kWh is drawn for cooling and for the infrastructure to run the equipment.

Sophisticated cooling methods such as liquid cooling are available, but the industry favorite is the hot aisle/cool aisle approach recommended by the American Society of Heating, Refrigerating and Air-Conditioning Engineers. Figure 2 illustrates how this method works. The drawbacks are problems with recirculation, where the hot air gets mixed in with the cold air, and hot spots, where a certain area always tends to be hotter than the rest. Overheating certain components can lead to system crashes and could shorten the server's lifespan.




Figure 2. How the hot aisle/cool aisle method works. Racks are placed facing each other and only the fronts of the servers, which face the same aisle are cooled. This configuration builds on the idea that not all server components need to be kept cool. Figure from U.S. EPA Energy Star Program Report to Congress on Server and Data Center Energy Efficiency, Aug. 2007; http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf.

Even so, many small data center operations have no air management, and implementing this approach would be a huge step forward. These centers tend to blast cool air, not only on the equipment but also on all other areas outside the data center that are in the same zone. Also, these centers often keep the office temperature at around 70 degrees Fahrenheit with the lights on—even when no one is there. Keeping IT equipment at 80 degrees has no effect on its operation and can yield huge energy savings. Microsoft's data center in Silicon Valley, for example, raised its temperature just 2 to 4 degrees and was able to save \$250,000 annually in energy costs.⁴

Recycling

The image in Figure 3 is the dismal reality of where computers and servers go after their useful life. The Environmental Protection Agency (EPA) estimates that the United States produces more than 2 million tons of electronic waste (e-waste) each year.⁵ Most of the components in IT equipment are not recycled into other products; rather, the remains are exported overseas to de-



(a)



(b)

Figure 3. Consequences of exporting of e-waste. (a) Electronic mountain in Guiyu, China, in May 2008. Electronic villages are becoming too prevalent in developing countries, resulting in practices such as (b) burning and metal reclaiming in “burn houses,” which endanger the lives of nearby inhabitants. To prevent such practices, the United States must take immediate steps to deal responsibly with e-waste. Photos used with permission from the Basel Action Network, 2008.

Estimating Energy Use without Meters

Many data centers do not have electric meters, either because of size or cost. To address this problem, we devised a method for estimating a data center’s energy use and applied it to a small government agency data center in Virginia. To estimate the energy use down to the equipment level, we obtained a list of installed information technology (IT) equipment and a product sheet for each piece and assigned the center the average power use effectiveness (PUE) of 2.0. On the basis of that figure and the list, we estimated that the data center power consumed roughly 490 megawatt-hours (MWh) of electricity.

This center was small enough that we could obtain product sheets, but such specifications would not be possible for a large center. The estimation method is based on the 11 categories of equipment types in Table A. For each equipment type, we surveyed the models available and, assuming a PUE of 2.0, estimated the energy consumed per year with typical equipment use.

Table A. Equipment and energy use (including cooling) by category.

Equipment Category	Energy Use (MWh)
PC	2
Server (2 CPUs)	7
Server (4 CPUs)	19
Blades	2
Storage array (5 TB)	8
Storage array (5-50 TB)	17
Storage array (50-100 TB)	36
Storage array (100-200 TB)	66
Storage array (200 TB)	85
Enterprise storage switch	10
Network devices	3

In creating the categories, we found that the energy use of storage arrays and servers depends on the amount of storage and CPUs installed, respectively. For storage arrays, we sampled arrays of different sizes and fitted a linear line that best describes the energy use of all these systems. Figure A shows the results.

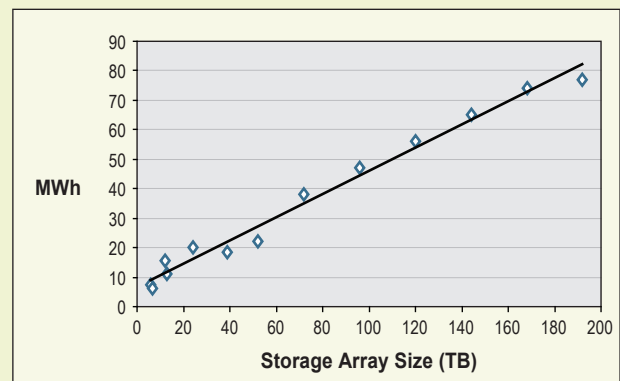


Figure A. Typical yearly storage array power use with cooling, assuming a power use effectiveness ratio of 2.0.

For example, we estimated the energy use of the Virginia data center to be 490 MWh, which was consistent with the 450 MWh we found by looking up product specifications.

Even with product sheets, however, a center has only an estimate of energy use. The most accurate way to find the exact amount of energy is to set up metering.

veloping countries. There, poor communities build a landfill and a small proportion is melted, boiled, and hammered back into basic chemical materials. In the process, workers inhale mercury, cadmium, barium, lead, and other dangerous toxic fumes. A better solution has to be found.

The Basel Action Network (BAN; www.ban.org) is an advocate for efforts to ban the exporting of e-waste. BAN was named for the Basel Convention, a multilateral environmental agreement, which in 1994 passed a landmark decision to reverse the deadly trend of exporting hazardous waste for any reason from rich countries to poorer ones. BAN also conducts field investigations in developing countries and provides photographic and video documentation of toxic trade. Of particular use to data centers is the organization's e-waste stewardship project, which aims to ensure that the exporting of hazardous e-waste to developing countries is replaced with producer responsibility and green design programs and legislation.

For federal agencies, the Electronic Product Environmental Assessment Tool (EPEAT) program (www.epeat.net) has become a key element in addressing e-waste, and others can use it as well. EPEAT, a non-government program to promote the use of equipment manufactured with disposal and greenness in mind, maintains a list of computers and monitors with a Bronze, Silver or Gold ranking. The rankings are based on environmental criteria specified in the IEEE 1680 standard. As of January 2009, Federal Acquisition Regulations require all federal agencies to buy 95 percent of their IT equipment from the EPEAT database. With more support not just from the federal government but also from industry and private purchasers, EPEAT can drive manufacturers to produce more green and recyclable products.

The Energy Star program is a green program wildly successful in the home appliance sector. The EPA is currently working with industry to come up with an Energy Star rating for enterprise servers. Once that rating is released, data centers will be better informed as to how much energy a piece of equipment will consume. Data centers often reach capacity not only because of space constraints but also because of overwhelming power requirements and an inability to cool all the equipment. As electricity rates continue to rise, the server energy rating has the potential to greatly influence data center practices.

Green building practices

Data centers are first and foremost buildings, and all buildings can benefit from environmentally responsible construction practices. The U.S. Green Building Council's Leadership in Energy and Environmental Design (LEED) is the gold standard in environment-friendly building practices, as "The Climate Is Right for Green" on p. 14 describes. The program is very detailed, and it is challenging to get the certification. Although not specifically designed for data centers, the program's detailed guidance and rigor makes LEED certification among the most prestigious achievements toward environmental responsibility.

Data centers have a unique opportunity to make a contribution by adopting green practices in cooling efficiency, energy use of IT equipment, and e-waste disposal—but all three areas must be addressed for a center to be considered green. Although progress is evident, there is a long way to go to entrench best practices and green guidelines. Smaller data centers, in particular, have fewer resources to experiment with green initiatives. They need monitoring methods that will confirm the promises of a green initiative so that they can choose the method with the highest return on investment.

Lack of energy-use awareness persists, not just in data centers, but in government, among private citizens, and in industry. Data centers store data, and those who produce that data often believe that unlimited storage exists for their business and personal e-mail. In reality, e-mail has to be stored somewhere at considerable cost to both data centers and the environment. With more awareness of this problem, citizens can use yet another approach to curb energy use in data centers—create less data.

As society worldwide produces more data, energy resources are dwindling. With global warming in the mix, finding a solution is imperative, but further study is critical in determining the effectiveness of the various approaches. As more data centers embrace green initiatives, more results will contribute to the overall picture. Clearly, when the promises of green become quantifiable, data centers should be quick to embrace such initiatives, not only for social and ethical reasons, but also to reduce operating costs and increase profitability. ■

References

1. U.S. EPA Energy Star Program Report to Congress on Server and Data Center Energy Efficiency, Aug. 2007; http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf.
2. List of countries by electricity production (between 2004–2007); http://en.wikipedia.org/wiki/List_of_countries_by_electricity_production.
3. Data Center Efficiency Measurements; <http://www.google.com/corporate/green/datacenters/measuring.html>.
4. R. Miller, Google: Raise Your Data Center Temperature, Oct. 14, 2008; <http://www.datacenterknowledge.com/archives/2008/10/14/google-raise-your-data-center-temperature/>.
5. S. Lezinski, Generation E-Waste: Confronting The Rising Tide of e-Waste In Maryland, June 1, 2008; http://wasteage.com/E-Waste/generation_ewaste_mde/.

David A. Garbin is a senior fellow at Noblis, where his interests include telecommunications technology, networking, network design and optimization, economic analysis, and voice and data communications. He received an MSEE from the Massachusetts Institute of Technology. Contact him at david.garbin@noblis.org.



Elizabeth W. Chang is a lead at Noblis, where her interests include geographical information systems. She received an ME in systems engineering from the University of Maryland. Contact her at elizabeth.chang@noblis.org.

