

APRIL 15, 2006 | CIO MAGAZINE

DATA CENTERS

Powering Down

Electricity-hungry equipment, combined with rising energy prices, are devouring data center budgets. Here's what you can do to get costs under control.

BY SUSANNAH PATTON

A typical 10,000-square-foot data center consumes enough juice to turn on more than 8,000 60-watt lightbulbs. That amount of electricity is six to 10 times the power needed to operate a typical office building at peak demand, according to scientists at Lawrence Berkeley National Laboratory. Given that most data centers run 24/7, the companies that own them could end up paying millions of dollars this year just to keep their computers turned on.

And it's getting more expensive. The price of oil (\$60 a barrel in February) may fluctuate, but the cost of energy to run the data center probably will continue to increase, energy experts say. This is because global demand for energy is on the rise, fueled in part by: the proliferation of more powerful computers. According to Sun Microsystems engineers, a rack of servers installed in data centers just two years ago might have consumed 2 kilowatts and emitted 40 watts of heat per square foot. Newer, "high-density" racks, which cram more servers into the same amount of space, are expected to consume as much as 25 kilowatts and give off as much as 500 watts of heat per square foot by the end of the decade. The dire predictions keep coming. Most recently, a Google engineer warned in a research paper that if the performance per watt of today's computers doesn't improve, the electrical costs of running them could ultimately exceed their initial price tag.

"As the demand for computing grows, the cost of power is a larger and larger concern," says Dewitt Latimer, CTO at University of Notre Dame. Latimer is grappling with finding the space and adequate power to handle a growing demand for cheaper and ever-more powerful high-performance computer clusters at Notre Dame. The problem comes not just from the computers themselves; Latimer is worried that the air-conditioning needed to keep the machines cool will also eat away at his budget.

Like Latimer, every CIO who is responsible for a data center—even those who outsource data center management to a hosting company—faces this conundrum: how to keep up with ever-increasing performance requirements while taming runaway power consumption. The problem is most pressing for companies on either coast and in large cities in between, where space is at a premium and companies compensate by putting more servers into their existing buildings. And there is no simple solution. Business demand for more applications results in companies adding more servers. According to market research company IDC (a sister company to CIO's publisher), server sales are growing by 10 percent to 15 percent annually.

The Solar-Powered Server Farm

How one company got its data center off the grid

[Read More](#)

Assess Your Energy Use

The U.S. Department of Energy is developing ways for companies to **evaluate how much energy their data centers consume.**

Nevertheless, some CIOs with huge energy bills are developing strategies for containing power costs by deploying more energy-conscious equipment and by using servers more efficiently.

"There's no question that the issue of power and cooling is a growing concern," says John Humphreys, an IDC analyst. "The assumptions used for building data centers have been blown away."

THE PROBLEM: IT HOGS ENERGY

IT's energy woes have a lot to do with market factors that affect everyone who drives a car or turns on a light switch; at the beginning of the year, the price of a barrel of oil was more than double what it was three years earlier. The price of natural gas, which fuels many of the country's electric power plants, has also shot up. And anyone who thinks the current energy crunch is going away need only look at global energy markets.

The oil shocks in the 1970s and '80s stemmed from large, sudden cuts in supply. This time, it's different. While it's true that some of today's high prices stem from supply shocks tied to the U.S. invasion of Iraq and hurricanes on the Gulf Coast, the world's thirst for oil over the past 25 years has grown faster than the energy industry has been producing it. And with rapid economic expansion in China and India, those countries are demanding more and more energy, putting further pressure on the world's energy markets.

Servers in corporate data centers may use less energy than manufacturing facilities for heavy industries, but within a company, IT is an energy guzzler. "We're pretty hoggish when it comes to power consumption in the data center," says Neal Tisdale, VP of software development at NewEnergy Associates, a wholly owned subsidiary of Siemens. NewEnergy's Atlanta data center performs simulations of the North American electric grid to help power companies with contingency planning. "We turn on the servers, and we just leave them on."

The exact amount of electricity used by data centers in the United States is hard to pin down, says Jon Koomey, staff scientist at Lawrence Berkeley National Laboratory. Koomey is working with experts from Sun and IDC to come up with such an estimate. Nevertheless, most experts agree that electricity consumption by data centers is going up. According to Afcom, an association for data center professionals, data center power requirements are increasing an average of 8 percent per year. The power requirements of the top 10 percent of data centers are growing at more than 20 percent per year.

At the same time, business demands for IT are increasing, forcing companies to expand their data centers. According to IDC, at least 12 million additional square feet of data center space will come online by 2009. By comparison, the Mall of America in Minnesota, the world's largest shopping mall, covers 2.5 million square feet.

SOLUTION #1: MORE EFFICIENT COMPUTERS

Just as automakers built SUVs when oil prices were low, computer manufacturers answered market demand for ever-faster and less expensive computers. Energy usage was considered less important than performance.

In a race to create the fastest processors, chip makers continually shrank the size of the transistors that make up the processors. The faster chips consumed more electricity, and at the same time allowed manufacturers to produce smaller servers that companies stacked in racks by the hundreds. In other words, companies could cram more computing power into smaller spaces.

Now that CIOs are beginning to care about energy costs, hardware makers are changing course. Silicon Valley equipment makers are now racing to capture the market for energy-efficient machines. Most chip makers are ramping up production of so-called dual-core processors, which are faster

than traditional chips and yet use less energy. Among these new chips is Advanced Micro Devices' Opteron processor, which runs on 95 watts of power compared with 150 watts for Intel's Xeon chips. In March, Intel unveiled a design for more energy-efficient chips. Dubbed Woodcrest, these dual-core chips, which Intel says will be available this fall, would require 35 percent less power while offering an 80 percent performance improvement over previous Intel chips. And last November, Sun Microsystems introduced its UltraSparc T1 chip, known as Niagara, which uses eight processors but requires only 70 watts to operate. Sun also markets its Galaxy line of servers as energy-saving equipment.

"The manufacturers are getting better now," says Paul Froutan, VP of product engineering for Rackspace, which manages servers for clients in its five data centers. With more than 18,000 servers to watch over, Froutan has been worrying about energy costs for years. He's seen the company's power consumption more than double in the past 36 months, and in the same period has seen his total monthly energy bill rise five times to nearly \$300,000.

Latimer, who oversees Notre Dame's Center for Research Computing, first appreciated the power consumption problem when the university decided to hire a hosting company to house its high-performance computers off campus. On-campus electrical costs associated with data centers have generally been rolled together with other facilities costs, and so the \$3,000 monthly utility bill from the hosting company—for running a 512-node cluster of Xenon servers—came as a shock.

Notre Dame's provost recently called Latimer and other leaders together to talk about how to handle the increasing demands that a growing research program was beginning to place on the campus utility systems and infrastructure. Faculty members are requiring more space, greater electrical capacity and dedicated cooling for high-powered computers and other equipment such as MRI machines. Latimer's recent conversations with Intel, AMD, Dell and Sun about his plans to buy new computer clusters "have been very focused on power consumption," he adds.

SOLUTION #2: THE LATEST IN COOLING

In September 2005, officials at Lawrence Livermore National Laboratory switched on one of the world's most powerful supercomputers. The system, designed to simulate nuclear reactions and dubbed ASC Purple, drew so much power (close to 4.8 megawatts) that the local utility, Pacific Gas & Electric, called to see what was going on. "They asked us to let them know when we turn it off," says Mark Seager, assistant deputy head for advanced technology at Lawrence Livermore.

What's more, ASC Purple generates a lot of heat. And so, Seager and his colleagues are working on ways to cool it down more efficiently than turning up the air-conditioning. The lab is trying out new cooling units for ASC Purple and the lab's second supercomputer, BlueGene/L (which was designed with lower-powered IBM chips, but is nevertheless hot). Lawrence Livermore recently invested in a spray cooling system, an experimental method in which heat emitted by the computer is vaporized and then condensed away from the hardware. Seager says this new method, which holds the promise of eliminating air-conditioning units, would allow the lab to save up to 70 percent on its cooling costs.

It's not only supercomputers that create supersized cooling headaches. Tisdale, with NewEnergy Associates, says maintaining adequate and efficient cooling is one of the hardest problems to solve in the data center. That's because as servers use more power, they produce more heat, forcing data center managers to use more power to cool down the data center. "You get hit with a double whammy on the cooling front," says Rackspace's Froutan.

To address the cooling dilemmas of more typical data centers, hardware makers such as Hewlett-Packard, IBM, Silicon Graphics and Egenera have offered or are coming out with liquid cooling options. Liquid cooling, which involves cooling air using chilled water, is an old method that is

making a comeback because it's more efficient than air-conditioning. HP's modular cooling system attaches to the side of a rack of HP computers and "provides a sealed chamber of cooled air" separate from the rest of the data center, says Paul Perez, vice president of storage, networking and infrastructure for HP's Industry Standard Server group.

More efficient servers help too. Last spring, Tisdale discovered that his data centers had reached their air-conditioning limit. While he had always imagined that a lack of physical space would be his biggest constraint, he discovered that if he ever lost power, his main problem would be keeping the air-conditioning going. Tisdale had replaced all 22 of his company's Intel servers in its Houston data center with two dual-core Sun Fire X4200 servers, part of Sun's new Galaxy line. The new servers are more energy-efficient, according to Tisdale. And so when he proposed installing the servers in Atlanta, he justified the purchase by arguing that he could avoid having to buy a bigger air conditioner, which would have used even more power. Tisdale said that according to company projections, the move will save electricity and reduce heat output by 70 percent to 84 percent.

What's more, there are better ways to use traditional air-conditioning. Neil Rasmussen, CTO and cofounder of American Power Conversion (APC), a vendor of cooling and power management systems for data centers, says CIOs should consider redesigning their air-conditioning systems, particularly as they deploy newer, high-density equipment. "Instead of cooling 100 square feet, it makes sense to look for the hot spots," concurs Vernon Turner, group vice president and general manager of enterprise computing at IDC.

Traditional cooling units "sit off in the corner and try to blow air in the direction of the servers," Rasmussen says. "That's vastly inefficient and a huge waste of power." Rasmussen argues that the most efficient way to cool servers is with a modular approach that brings cooling units closer to each heat source. Meanwhile, he adds, CIOs who manage data centers in colder climates should use air conditioners that have "economizer" modes, which can reduce the power consumption in the dead of winter. Newer air conditioners have compressors, fans and pumps that can slow down or speed up depending on the outside temperature.

SOLUTION #3: A MORE EFFICIENT DATA CENTER

Just as aging cars are not as fuel-efficient as newer models, the majority of the country's data centers are using a lot more energy than they should. A survey of 19 data centers by the consultancy Uptime Institute found that 1.4 kilowatts of power are wasted for every kilowatt of power consumed in computing activities, more than double the expected energy loss.

However, like many people who aren't going to junk their older cars right away, many companies aren't ready to tear out their data centers to build new ones with a more efficient layout. "We haven't reached the point yet where it makes financial sense to rebuild most data centers from scratch," says Rackspace's Froutan. And so for most companies, the journey toward an energy-efficient data center will be a gradual one.

For NewEnergy Associate's Tisdale, that means retiring aging servers in one data center seven at a time and replacing them with more energy-efficient equipment. But redesigning your data center also means making the most of what you have through server consolidation and, more specifically, the use of virtualization software.

Virtualization is a technology that allows several operating systems to reside on the same server. Froutan says that virtualization will help his data centers make do with fewer servers by allowing them to perform more tasks on one machine. In addition, he says, energy can be saved by deferring lower-priority tasks and performing them at night, when the cost of power can be three times less expensive. IDC's Turner agrees that CIOs need to improve server utilization in order to cut both power and cooling costs. Instead of building one server farm for Web hosting and another for

application development, for example, they should use virtualization to share servers for different types of workloads.

Finally, advises APC's Rasmussen, if you are building a new data center, it's better to design it to accommodate the equipment that you need right now, rather than building facilities designed for what you might eventually need as you grow, as many companies have done. By using a more modular architecture for servers and storage—so capacity can be added when needed— a company can avoid such waste and still be prepared for growth.

HOW TO START SAVING

As CIOs search for more energy-efficient data center equipment and design, they need to educate themselves about which solutions will work best for them. As part of the information-gathering process, CIOs should establish metrics for power consumption in their data centers and measure how much electricity they consume.

There aren't many generally accepted metrics for keeping tabs on power consumption. But according to Turner, such metrics could include wattage used per square foot, calculated by multiplying the number of servers by the wattage each uses and dividing by the data center's total square footage. Sun has come up with a method called SWaP, which stands for Space, Wattage and Performance. The company says this method, which lets users calculate the energy consumption and performance of their servers, can be used to measure data center efficiency. John Fowler, executive VP of the network systems group at Sun, says sophisticated customers are installing power meters at their data centers to get more precise measurements.

It also pays to be an energy-aware buyer. When looking at how much energy a new server might use, "Don't just take the vendor's word for it," Fowler says. He suggests having a method for testing the server and its energy use before buying. However, the industry is still working on methods for comparing servers from different vendors in a live environment.

Ultimately, vendors' "eco-friendly" messages may resonate only slightly. NewEnergy's Tisdale, for example, still cares most about maintaining server performance. But he is impressed that new equipment will help him add more computing capability while maintaining current power usage levels. "Like a lot of people," he says, "I'm not interested in turning off the servers."

Senior Writer Susannah Patton can be reached at spatton@cio.com.

© 2006 CXO Media Inc.

Dated: April 15, 2006
<http://www.cio.com/archive/041506/energy.html>