

# Powering Down the Computing Infrastructure

George Lawton


**T**he amount of power used to drive the world's computing-related infrastructure—including PCs, servers, data centers, routers, and switches—has been growing rapidly.

Complicating this situation has been the increasing amount of heat that higher-performing systems generate. In response, manufacturers have had to add cooling systems, which use still more power.

The very largest computing complexes, such as data centers, now use more power than some large factories. For example, the five largest search companies now use about 2 million servers, which consume about 2.4 gigawatts, according to Ask.com vice president of operations Dayne Sampson. By comparison, the US's massive Hoover Dam generates a maximum of about 2 gigawatts.

This is a major reason why companies like Ask.com, Google, Microsoft, and Yahoo! are building facilities in the Pacific Northwest, where they can tap into relatively inexpensive hydroelectric power generated by dams constructed on the area's many rivers.

Computing-related power usage currently represents about 15 percent of US electrical consumption, according to Mark P. Mills, co-founding partner of the Digital Power Capital equity firm and also chief technology officer and board chair of ICx Technologies.



All this has yielded higher operating costs for companies.

The worldwide cost of powering and cooling computers in 2005 was \$26.1 billion, about half the \$54.9 billion spent on buying new servers, said Jed Scaramella, an analyst with IDC, a market research firm. And IDC expects server-related power costs to represent an increasing percentage of purchase expenditures, as Figure 1 illustrates.

Meanwhile, some government agencies are mandating energy efficiency in computer-related and other products.

Manufacturers, vendors, and users are thus looking for ways to reduce power consumption. For example, chip makers such as AMD and Intel are making more efficient microprocessors. However, said Mills, the growing number of increasingly powerful chips in use, even if more efficient, will still consume considerable energy.

Moreover, the IT infrastructure will continue adding computers and networking systems and will thus use more electricity.

Therefore, researchers are looking at options other than microprocessors—such as power-supply technology, data-center architectures, and thermal interfaces—for reducing energy consumption.

## ELECTRICAL SUPPLY

The most common computer power supplies are built to work with ATX (Advanced Technology Extended) motherboard/computer-case technology.

Computer power supplies convert the incoming alternating-current (AC) electricity—generally 120 volts in the US, 220 volts elsewhere, and even higher for data centers—to various levels of the low-voltage (from 1.25 to 12 volts) direct-current (DC) electricity that the processor and other subsystems require to operate efficiently.

Power supplies account for more than 2 percent of the US's electricity consumption, according to the Electric Power Research Institute (EPRI).

Most modern ATX power supplies are built with simple circuits that yield better economies of scale in production but that don't use power most efficiently. They thus typically have an average efficiency—in terms of the portion of the incoming energy that actually reaches the device—of only 50 percent at lighter loads and 67 percent for heavier loads, noted Ecos Consulting channel manager Jason Boehlke.

The EPRI says that more efficient design could cut power-supply energy consumption in half.

## Simpler power-supply standard

Most PC motherboards convert the different voltages of electricity coming out of the power supply to meet the varying needs of computer components. Manufacturers, therefore, must engineer ATX power supplies to support multiple voltages, rather than optimize them to output just one voltage. This creates power-use inefficiencies.




Figure 1. The cost of powering and cooling servers has increased during the last decade and will increase steadily during the next few years, according to market research firm IDC. Reflecting this is the growing percentage of server purchase expenditures that energy costs represent.

Google is thus calling on the computer industry to create a simpler and more efficient standard in which the power supply would convert incoming AC electricity only to 12 volts.

Optimized, and thus more efficient, transformers on the motherboard would then convert the electricity to the voltages required by the subsystems they support.

Using a single 12-volt output would let the power supply convert voltages at up to 92 percent efficiency and let the optimized motherboard transformers operate at up to 95 percent efficiency, said Bill Weihl, Google's director of energy strategy.

Weihl estimated that if deployed in 100 million PCs running for an average of eight hours per day, this new standard would save 40 billion kilowatt-hours over three years.

This would be effective only if the motherboard components convert the voltage efficiently, warned Chris Calwell, Ecos' vice president and director for policy and research.

### Improving data-center power distribution

To yield a given amount of supplied power, the current must be increased when the voltage is lowered. Thus, supplying power at lower voltages generates more heat, which wastes some of the power that could be used by servers. Therefore, researchers are studying the distribution of power at higher voltages.

Currently, data centers commonly distribute power to their servers at 120 or 240 volts AC in the US and 415 volts AC in Europe.

Bill Tschudi, a principal investigator at the US's Lawrence Berkeley National Laboratory, said researchers there have experimented with transmitting power at 380 volts DC within the data center, which reduced power loss by 10 to 15 percent. This would not only reduce heat-based energy losses but would also eliminate inefficiencies caused by having to convert AC coming from the electrical grid into DC, he explained.

Using 480 volts AC could reduce heat-based energy losses by 8 percent, according to Bill Carlini, direc-

tor of product management at vendor American Power Conversion (APC). And, he said, 480-volt AC electrical-distribution equipment is already produced in greater quantities, and thus costs less, than DC-based products.

### Government incentives

In 2004, Ecos and the Northwest Energy Efficiency Alliance ([www.nwalliance.org](http://www.nwalliance.org)), a consortium of US electric utilities, launched the 80+ initiative to give incentives, such as rebates, to companies that deploy power supplies with efficiencies of 80 percent or more across a range of loads.

The efficient power supplies are currently produced in lower quantities and thus cost about 15 percent more than typical products, said Laurent Jenck, manager of the Power Supply System Engineering Group at power-supply vendor ON Semiconductor.

As of 20 July this year, the US Department of Energy will include the 80+ standards in its Energy Star certification requirements, said

Katharine Kaplan, an Energy Star staff member at the US Environmental Protection Agency (EPA). The DoE lets products that meet these requirements post Energy Star seals, designed to appeal to energy-conscious customers.

The EPA claims Energy Star-qualifying computer equipment is 65 percent more efficient than conventional models and that its use could save US businesses up to \$1.2 billion in energy costs over the machines' lifetimes.

Japan requires central and local governments to offer economic breaks to promote energy-use reduction. In addition, the European Commission, as part of its Energy Efficiency Action Plan, and the Chinese government plan to offer similar tax incentives to organizations.

### IMPROVING THE THERMAL INTERFACE

A high-end server consumes about 300 watts, according to Michael Bell, an analyst with market-research firm Gartner Inc. A rack of 40 blade servers thus uses about 12 kilowatts.

In the process, servers generate considerable heat, which requires cooling that uses additional energy. Most servers require 1 watt of cooling for every watt of power used in a moderately dense server system—one that uses less than 6 kilowatts per rack—and 2 watts of cooling for every watt used in the densest systems.

Reducing the amount of heat generated by servers is thus critical for reducing overall power usage.

**Dynamic cooling.** Hewlett-Packard recently announced the upcoming release of the Dynamic Smart Cooling system. This system uses localized temperature sensors to adjust the airflow from a distributed set of fans to optimize cooling for the components that are working the hardest and thus generating the most heat.

By not having to cool an entire system at the same high level, this approach could save up to 40 per-

cent of blade-server systems' energy costs, said Ron Mann, HP's director of engineering enterprise infrastructure solutions.

**Coolant-based systems.** Emerson Network Power has developed a system that runs coolant through channels directly to each blade server. The system uses sensors to determine which areas are the hottest. It thus can use cooling energy efficiently.

### Virtualization technology is reducing power usage.

SprayCool has developed a technology that uses fluorinert, a non-conductive liquid that can run directly on top of a server's electronics, said company marketing manager Patchen Noelke. This could be used with heat sensors to save energy by directing coolant only where most needed.

**Nanocarpet.** Powerful computers typically use heat sinks to passively draw heat away from the circuitry and reduce the need for cooling equipment. Purdue University researchers are developing a novel thermal interface material made of carbon nanotubes configured like a Velcro carpet that attaches to the chip and the heat sinks.

This material and its carpet-like structure promise to conduct heat away from the chip more efficiently than traditional heat-sink materials such as aluminum and copper and thus provide 10 degrees Celsius more cooling, said associate professor Timothy Fisher, who is leading the project.

Previous similar cooling approaches faced problems because the temperatures required to grow nanotubes on the chips damaged the processors, Fisher said. The Purdue technique works at lower temperatures.

### VIRTUALIZATION

The growing use of virtualization technology is also reducing power usage.

Traditionally, companies have used one server for one task, such as e-mail or Web hosting, to simplify configuration, management, and operating-system use. Each server in a data center thus frequently sits idle and, in fact, operates at only 5 to 20 percent of its processing capacity, noted Bogomil Balkansky, director of product marketing at virtualization vendor VMware.

Nonetheless, the servers are still running and consuming about 80 percent of their peak energy-use levels, noted Dave Douglas, Sun Microsystems' vice president of eco-responsibility.

Virtualization lets a single computer run multiple operating systems or several versions of the same operating system, thereby enabling a server to run more than one task. Users could thus shut down numerous servers and thereby reduce power usage significantly.

According to Balkansky, some virtualization customers have already significantly reduced the number of servers they use. He said a data center could save \$560 per year in power and cooling for every server that is shut down.

### OTHER APPROACHES

Various companies and groups are beginning to address significant data-center energy-reduction challenges.

**Industry and governmental organizations.** Several companies, including APC, Dell, HP, IBM, and Sun, founded the Green Grid ([www.thegreengrid.org](http://www.thegreengrid.org)) to focus on the best practices and management approaches for lowering data centers' energy consumption.

The DoE recently released the Server Energy Measurement Protocol ([www.energy.gov/index.cfm?c=products.pr\\_servers\\_datacenters](http://www.energy.gov/index.cfm?c=products.pr_servers_datacenters)) and is working with companies to test and adopt it.

"The protocol establishes a procedure for attaching an energy-usage measurement to existing performance measurements for servers,"

said Jonathon Koomey, a Lawrence Berkeley National Laboratory staff scientist who is working on the project.

This is important because there currently is no standard way of measuring servers' energy-related performance, he noted. The ability to accurately gauge energy usage is critical to determining conservation efforts' effectiveness, he added.

The nonprofit Standard Performance Evaluation Corp. ([www.spec.org](http://www.spec.org))—which includes companies such as AMD, IBM, Intel, Microsoft, and database vendor Sybase—is working on a set of benchmarks related to servers' energy usage, which could be finalized in the near future.

**Modular systems.** APC and Emerson Network Power are developing modular power and cooling systems. These products would let companies add modules—and thus relatively small amounts of power and cooling capacity—as data centers grow, explained Emerson vice president of power engineering Peter Panfil. Businesses would thus effi-

ciently use only what is needed at a specific time, he noted.

APC calls its system InfraStruxure. Emerson named its product the Adaptive Architecture.

**Project Blackbox.** Sun recently created Project Blackbox, in which a complete data center can operate in a standard 20-foot-long shipping container equipped with a cooling system and multiple power and high-speed-networking connectors.

Sun engineered these servers to reduce heat and distribute power more effectively and thus use energy 20 percent more efficiently than standard 10,000-square-foot data centers.

**B**efore many companies spend time and money on new power-saving equipment and measures, they could take less expensive steps with their existing technology, such as turning off unused servers and improving the airflow through data centers to reduce heat buildup, said Ken Brill, founder of the Uptime Institute, which provides information on best practices for reducing data-center downtime.

In the long run, though, predicted IDC's Scaramella, companies will look at new approaches as their power consumption and energy bills increase.

Paradoxically, said Digital Power Capital's Mills, improved computer power efficiency might lead to even greater energy consumption by the entire computer-related infrastructure.

Energy efficiency lowers usage costs and thus makes computers and other devices more attractive to a wider audience, thereby increasing demand, he explained. Thus, he said, even though the new equipment might be energy efficient, companies will buy so much of it that the overall amount of power consumed will increase.

"Maybe one day," Mills concluded, "efficiency will grow faster than demand, and then we will end up with a reduction." ■

*George Lawton is a freelance technology writer based in San Francisco, California. Contact him at [glawton@glawton.com](mailto:glawton@glawton.com).*

## IEEE Software Engineering Standards Support for the CMMI Project Planning Process Area

By Susan K. Land  
Northrup Grumman

Software process definition, documentation, and improvement are integral parts of a software engineering organization. This ReadyNote gives engineers practical support for such work by analyzing the specific documentation requirements that support the CMMI Project Planning process area. \$19

[www.computer.org/ReadyNotes](http://www.computer.org/ReadyNotes)

# IEEE ReadyNotes

