Managing Energy Costs in Data Centers

E SOURCE
Customer Direct
Pamphlet licensed for distribution to the customers of E SOURCE members

Energy costs in data centers are growing rapidly as servers have become more energy-intensive and more servers are being packed into tight spaces to satisfy increased computing demand. This has resulted in increased cooling loads and in some cases has strained the electricity distribution infrastructure required to meet data-center power needs. These trends are likely to continue over the long term as IT needs continue to expand. Another development that may make data centers even more energy-intensive is the possibility of "cloud computing"—moving the computing power of individual workstations onto central servers.

Improving the energy efficiency of your data center can help with each of these problems. Practical measures that can help reduce energy costs include the use of a strategy for managing server power, arrangement of server-room equipment for more effective cooling, and use of more energy-efficient IT equipment.

Power Management

Reducing energy used by existing IT equipment is often the best way to begin managing your energy bills. This is because it can be done for little or no cost and provides the additional benefit of reduced cooling loads. All of the power input to IT equipment eventually turns to heat, which must then be removed by the cooling system. If the IT equipment uses less energy, the accompanying reduction in the facility's cooling load will lead to additional energy savings. Although there is considerable variation among facilities, a typical data center that reduces its computer load power requirements by 1.0 kilowatt (kW) would also offset approximately 0.6 kW of air-conditioning (AC) power.

The easiest way to reduce energy use by servers is to use their built-in power-management system. The Rocky Mountain Institute has found that doing so could cut server energy use by 20 percent, but that many data centers do not take advantage of this savings opportunity. Chief scientist Amory Lovins told *Computerworld* magazine in June 2007 that "in a typical data center, the electricity usage hardly varies at all, but the IT load varies by a factor of three or more. That tells you that we're not properly implementing power management."

Power-management functions will idle servers in much the same way that personal computers can be set to standby when they are not being used. Going one step further, some servers can be turned completely off to match predictable workload decreases—for example, at night.

Equipment Layout for Effective Cooling

The most common means of increasing cooling effectiveness is to keep hot return air separate from chilled air. More advanced techniques include liquid cooling and chip-based cooling.

The Rocky Mountain Institute has found that with careful engineering, design, and operation, data-center managers may be able to achieve a 30 to 50 percent reduction in cooling energy usage and costs.

Hot/cold air separation. One of the simplest approaches is to keep hot return air separate from cold air. Chilling and moving the air used to cool servers is energy-intensive, frequently consuming 50 percent or more of the total energy usage in data centers. When hot air exhausted from servers is allowed to mix with cool air, some of that energy and money used to condition and move the air is wasted. In addition, hotter return air saves money by allowing AC units to run more efficiently. There are several techniques for keeping hot and cold air separate.

- *Create hot and cold aisles.* This is a fairly common and effective way of reducing energy costs. Servers are typically designed to draw cool air in from the front of the server and exhaust warm air out the back, making it easy to arrange hot and cold aisles (**Figure 1**). Arranging racks so that the servers face one another is a first step in keeping hot and cold air separate. Perforated tiles in raised-floor cooling systems are placed in cool aisles to direct chilled air to the server intakes.
- *Eliminate air exchange.* Once you have created separate hot and cold aisles, it is important to eliminate any air exchange between the two. Panels and cable grommets can be used to fill space between server racks. The Lawrence Berkeley National Laboratory found that hot-cold air isolation allowed AC fans to maintain the proper temperature while running at a lower speed, resulting in a 75 percent reduction in fan energy use.
- *Use high-density heat containment.* The software company Oracle applied an extension of the hot aisle/cold aisle concept known as a high-density heat containment strategy (**Figure 2**, next page). This approach uses sealed racks and hot air return ducts to completely separate the hot and cold air. Oracle found that this approach allowed the AC units to run more efficiently and that the sealed racks could be built up modularly. Although the return ducts require booster fans supplementing the server fans, the extra power is more than offset by the increased compressor efficiency and lower AC fan speeds. Oracle was able to increase the floor-space concentration of IT equipment by 58 percent and save more than $200,000 per year through a high-density heat containment strategy that completely separates the cool and warm air.

Liquid-cooled cabinets. Advances in direct cabinet cooling include placing coils with chilled water much closer to the servers themselves. This is a promising development because liquids have a greater capacity to absorb heat and are less expensive to move than air, allowing them to move more heat away from servers with less energy. Cooling the servers directly also reduces the need to cool the entire room. Typically, chilled water is run through tubes along the back or side of the server rack. Hot air from the back of the rack passes by the cool tubes, and as it is cooled it is either expelled into the data center or blown directly back into the front of the rack. Several systems are now commercially available, and many experts in the industry expect direct cabinet cooling to be the wave of the future.

FIGURE 1: Hot aisles and cold aisles

Creating separate hot and cold aisles provides more efficient airflow and increases cooling efficiency. It is important to block any openings that allow warm air to mix with the cool air.

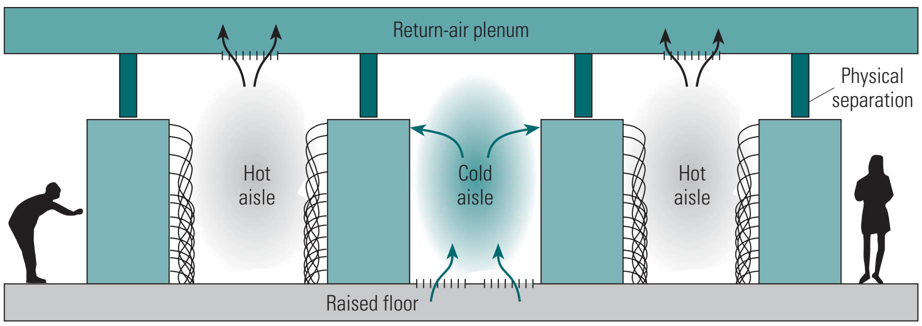

Courtesy: William Tschudi, Evan Mills, and Steve Greenberg, Lawrence Berkeley National Laboratory, and Peter Rumsey, Rumsey Engineers

The ultimate extension of this concept is known as chip-based cooling, in which an inert liquid cooling medium is brought into direct contact with chips inside of the servers. Heat is then removed from the liquid with a heat exchanger. This technology is commercially available, but requires replacing servers and a completely redesigned cooling system. Applications in the northwestern U.S. demonstrated that spray cooling removed about one-quarter of the cooling load, so conventional cooling techniques still do the majority of the work. One utility company has found that chip-based cooling achieved a payback of 9.4 years, and this may improve if costs decline and rebates for these systems become more common.

Efficient Equipment

Energy-efficient IT equipment is the holy grail of data-center efficiency because it promises to directly reduce energy consumption as well as cooling loads. New concepts and emerging standards will aid in the quest.

Data-center virtualization reduces the number of servers. Virtualization refers to running multiple computing functions on a single server, rather than requiring a separate machine for each computing task. Many servers operate at processor utilizations below 15 percent while still consuming 60 to 90 percent of the energy required for them to run at maximum capacity. Virtualization allows data centers to accomplish the same work using fewer servers operating at higher processor utilizations. Although virtual servers are typically more powerful and energy-intensive than a single traditional server, reducing the number of servers and improving utilization saves energy.

Continuous power management saves energy. Servers with the capability of scaling microprocessor voltage or frequency can save energy by closely matching server capacity to workload demands. This has the potential to follow workloads more closely than more common power-management functions; however, not all servers have this ability. Still, an estimate by the U.S. Department of Energy suggests that it may be possible to reduce server annual energy consumption by 30 percent through the use of continuous power management.

FIGURE 2: High-density heat containment keeps warm air and cool air separate in Oracle's data center

The duct return collects warm air from the rear of the servers and physically separates it from cool air supplied by the air-conditioning system.

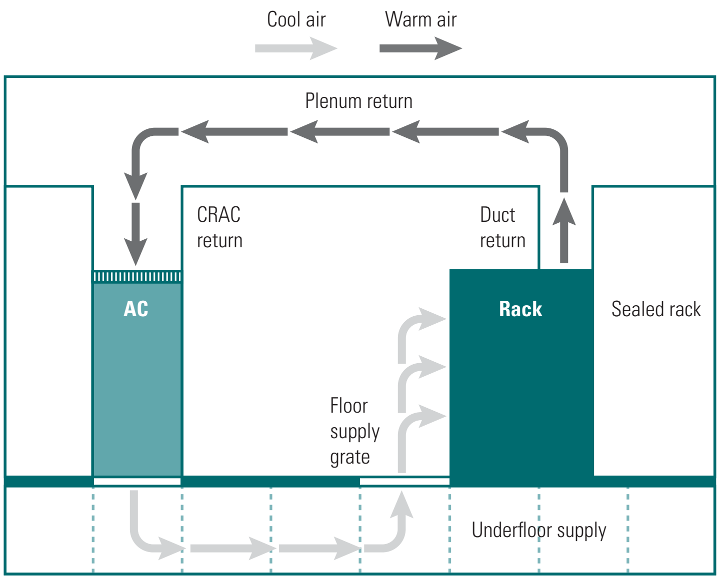

Note: CRAC = computer room air conditioning.

© E SOURCE

Server efficiency standards are being developed. It is currently difficult to compare the energy efficiency of different servers because there is a lack of equipment standards. Fortunately, several efforts are underway to develop equipment standards, and accurate comparisons may soon be possible.

The U.S. Environmental Protection Agency (EPA) is working to create an Energy Star specification for servers and other IT products. And the Standard Performance Evaluation Corporation (SPEC) is currently developing a standard of server efficiency at performance levels ranging from full capacity down to idle. SPEC has released an initial benchmark for one type of server function, and benchmarks of a broader range of workloads are planned.

MAID systems are a promising new development. Data is typically stored on disks that must remain spinning—and consume energy—for their information to be retrieved. Massive Array of Idle Disks (MAID) systems catalog information according to how often it is retrieved and place seldom-accessed data on disks that are kept idle until the data is needed. According to the manufacturer, MAID can save up to 85 percent of storage power and cooling costs. Still, MAID can be difficult to implement and has thus far seen limited applications and testing.

Data-Center Benchmarking

There are several whole-facility data-center initiatives underway for benchmarking. Being able to compare energy consumption across a number of facilities and applications will help managers define data-center targets and best practices.

The EPA is developing efficiency benchmarks for stand-alone data centers as well as those located in offices and other commercial applications.

The Green Grid organization recently released its first energy-efficiency guidelines for data centers. The guidelines use high-level benchmarks comparing the proportion of IT electricity consumption to total facility consumption and also provide a standard for greenhouse gas emissions.

Data-center benchmarks are still being developed, and facility managers can further these initiatives by providing energy data from their own data centers at the EPA and Green Grid web sites listed below.

Resources

The Uptime Institute, www.uptimeinstitute.org

Data Center Energy Management Reading Room, Lawrence Berkeley National Laboratory, http://hightech.lbl.gov/DCTraining/reading-room.html

Data Center White Papers, Emerson Network Power, www.liebert.com/servicesupport_pages/WhitePapers.aspx?id=15

Enterprise Server and Data Center Energy Efficiency Initiatives, EPA Energy Star Program, www.energystar.gov/index.cfm?c=prod_development.server_efficiency

Standard Performance Evaluation Corporation, www.spec.org

The Green Grid, www.thegreengrid.org