ZIFF DAVIS MEDIA CUSTOM PUBLISHING

A MARKET ANALYSIS

Data Center Power and Heat Management: Ready or Not?

From a survey of leading IT professionals

Executive Summary

With energy costs and hot, power-hungry silicon dominating the headlines, some observers have predicted an imminent crisis in the data center. Those fears are almost certainly overstated: IT managers report that they have power supply issues in hand at the moment, and that they are taking active steps to monitor and manage the heat generated by servers supplying the computing horsepower that drives business today.

Nevertheless, there are some early warning signs that it would be imprudent to ignore. Most significant is the number of businesses that are already using spot cooling techniques to deal with heat generated by increasing equipment densities. This is striking, given that server consolidation programs are gaining momentum and blade servers have only recently begun to ship in serious volumes—two trends that are likely to quickly and dramatically increase data center equipment density. In general, heat and power management strategy and tactics are still in their infancy at most companies, with most of the focus on room-level systems. Decision-makers are not yet thoroughly versed in the granular power-consumption data that they'll need as data center heat continues to rise, and may be underutilizing power-saving tools such as processor-level power management that are already widely available.

Survey rationale and respondent profile

The results presented here were obtained in a study conducted from June 21-26, 2006 by The Strategy Group, an independent research company located in Englewood, FL. An online survey was conducted among a representative sample of technology decision-makers selected from the Ziff Davis Media subscriber data base. In order to qualify, respondents demonstrated that they are knowledgeable about data center issues within their organizations. In total, 357 subscribers qualified for and completed the survey. Respondents reflect a broad cross-section of organizational sizes and industries.

Research sponsored by AMD

Plenty of power…for now

Managing the power input and heat dissipation needs of a data center has acquired a fair bit of attention over the past year. Rising energy expenditures have sent businesses diving into their utility bills searching for leaks, while equipment failures in ostensibly temperature-controlled rooms are leading IT and operations managers to spot-check temperatures.

Electricity prices have been on the rise, but the real culprit driving spiraling energy costs is an insatiable demand for computing horsepower. According to data from analyst firm IDC reported in recent news story on data center power consumption, there were six million servers worldwide 10 years ago. Today, there are 24 million, and IDC projects that there will be 35 million by 2010. A rack of gear that held an average of seven servers 10 years ago holds between 20 and 22 today, and servers that used to consume an average of 100 watts of power today consume an average of 400 watts. As a result, server racks that may have consumed only two kilowatts of power even two years ago will soon consume 10 kilowatts or more.

CURRENTLY HAVE SUFFICIENT ELECTRICAL POWER IN THE DATA CENTER FIGURE 1

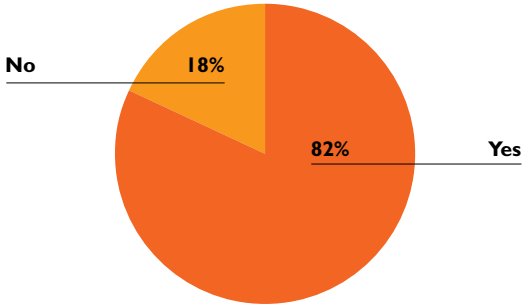

Our research supports these broad industry trends. Of the respondents who indicated that power consumption and cooling were significant issues, more than half indentified the need to add server capacity as one of the primary reasons driving the focus on power and cooling. For 36 percent, it is the primary reason.

Our research indicates that companies are aware of how much power their data centers are consuming, but for the time being, businesses feel that they have a handle on power availability and energy costs. Only 14 percent of survey respondents list the rising cost of electricity as the primary reason for their company's focus on power consumption issues. Also, 82 percent of companies report that they currently have sufficient electrical power in the data center (see Figure 1), and slightly more than half of the survey respondents report that their electrical power will meet projected growth for more than two years (Figure 2).

However, there could be trouble on the horizon. Nearly half of IT and Operations decision-makers believe that they will need additional power capacity in their data

AMOUNT OF TIME ELECTRICAL POWER WILL BE SUFFICIENT TO MEET PROJECTED GROWTH FIGURE 2

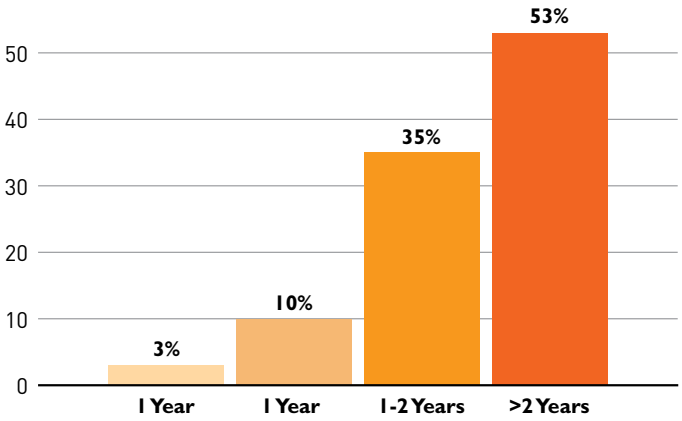

EQUALS 101% BECAUSE OF ROUNDING

centers in 24 months or less. Also, recent stories in the trade media indicate that executives outside the IT and Ops space are becoming increasingly aware of the costs of data center power consumption. It's clear that the days of energy spending as an unexamined cost of doing business are over.

Data center heat rising

Companies are slightly less sanguine about heat management. For years, they've been intimately familiar with Moore's Law, but they are now receiving an education in the First Law of Thermodynamics: energy can be neither created nor destroyed, so when you add energy to a system it has to go somewhere. Servers take in electrical power and convert a significant portion of it to heat.

It's a lesson that's likely to grow more pointed in the coming months as server racks become increasingly dense and active, causing hot spots that shorten equipment life. We've managed to hold off on the inevitable pun, but it will wait no longer: blade servers are hot. IDC expects the blade market to grow to $10 billion by 2009 and, according to the analyst firm, revenue in the second quarter of 2006 grew by 37.1% percent over the same period in 2005.

And yet, the blade server market is still in its infancy. According to IDC, bladed servers accounted for only 5.2 percent of the total server revenue in the second quarter of 2006. Given that context, the effect that blades are already having on the data center heat management equation is eye-opening. Companies appear to be

SERVER CONSOLIDATION OR BLADE SERVERS CREATING HOT SPOTS FIGURE 3

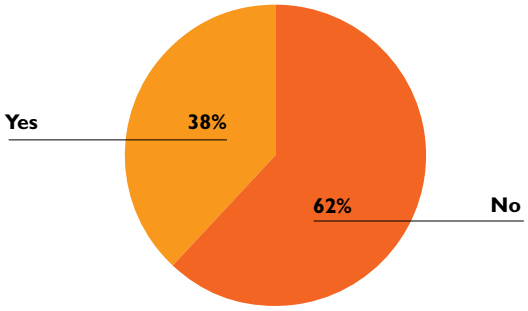

satisfied with established room-level cooling strategies such as raised floors and hot/cold aisle strategies. Only 30% of respondents indicated that they couldn't get enough cold air through a raised floor to keep systems within targeted temperature zones.

But 38 percent of companies report that their server consolidation initiatives or blade servers created hot spots that cannot be managed with ambient (room-level) cooling strategies (figure 3). These hot spots are more than an inconvenience, as a nearly identical percentage (43 percent) of survey respondents reported implementing stopgap or spot measures to cool specific locations within the data center (Figure 4).

Businesses looking for solutions

These results suggest that the emphasis on heat management within data centers is shifting and that companies are looking for solutions that extend beyond whole-room cooling.

IMPLEMENTED STOP GAP OR SPOT SOLUTIONS FOR COOLING SPECIFIC RACKS OR LOCATIONS IN THE DATA CENTER FIGURE 4

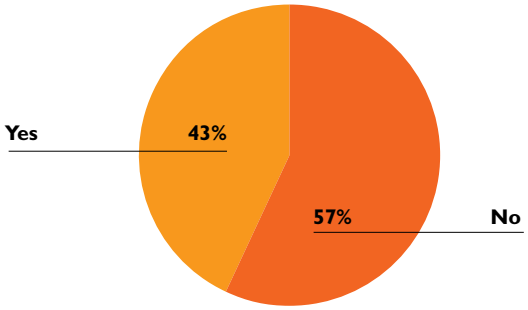

Research sponsored by AMD

The solutions to heat management run the gamut from tactical to the strategic. The first solution that many businesses have implemented is simply to arrange data center equipment in such a way that the existing temperature-maintenance hardware can work as efficiently as possible: for example, cool aisle/hot aisle layouts help ensure that the air intakes on server enclosures are actually drawing cool air (figure 4).

In addition, 52 percent of businesses are considering virtualization as a way to address power and cooling issues. Instead of assigning an application its own server box or blade, companies can use virtualization technology to divide the resources of a single, powerful server between several applications. Combined with other virtual computing enablers like storage area networks (SANs),

TAKE POWER MEASUREMENTS DURING THE EVALUATION OF NEW SERVER EQUIPMENT FIGURE 5

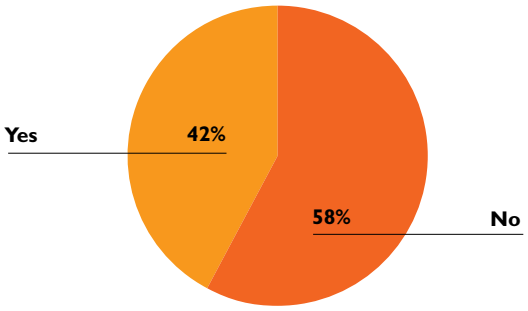

virtualization creates a pool of computing resources that can be tapped as necessary. As a result, companies can add applications without adding hardware, easing the load on heat management systems.

Companies don't appear to be ready to base their purchasing decisions on power and heat issues—according to research recently published by Ziff Davis, "Power Consumption and Cooling in the Data Center: A Survey", 62 percent of businesses

MOVED TO AN ELECTRICAL PLANNING MODEL THAT LOOKS AT RACK LEVEL INSTEAD OF ROOM LEVEL FIGURE 6

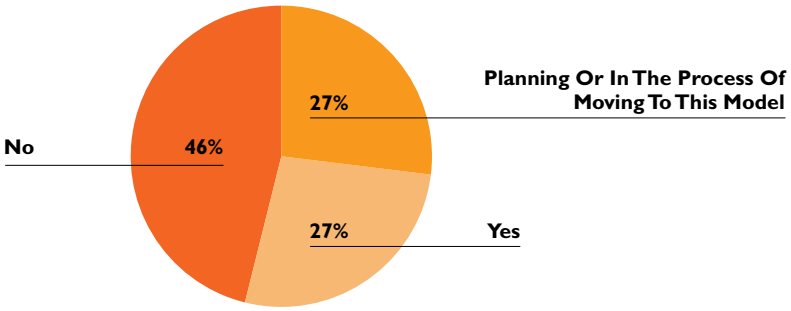

say that power consumption and cooling issues did not affect their server purchase decisions in the last twelve months. However, companies are absolutely considering the impact that new equipment will have on their data center: 42% of businesses say that they measure power consumption when they're evaluating new server equipment (Figure 5).

Companies are also refining what they do with that power consumption information. Already, 27 percent of businesses use an electrical model that lets them plan for consumption at the rack level instead of the room level. Another 27 percent are either planning for or in the process of moving to such a model (Figure 6).

Room for improvement remains

Despite the attention that's starting to be lavished on power and heat management, it's clear that much work remains to be done. For example, data center planners are only beginning to obtain and recognize the granular data that will help them deal with increased server densities. Watts per square foot is a key measure of power and heat distribution that can help a data center planner break down a room into hot and cool zones. However 84 percent of the businesses surveyed don't know the current planned level of watts per square foot in their data center (Figure 7). In all likelihood, this information is available or could be calculated at every business surveyed, but it's not yet a top-of-mind figure that IT or Ops managers can quickly quote.

CURRENT PLANNING LEVEL FOR DATA CENTER POWER IN WATTS PER SQUARE FOOT FIGURE 7

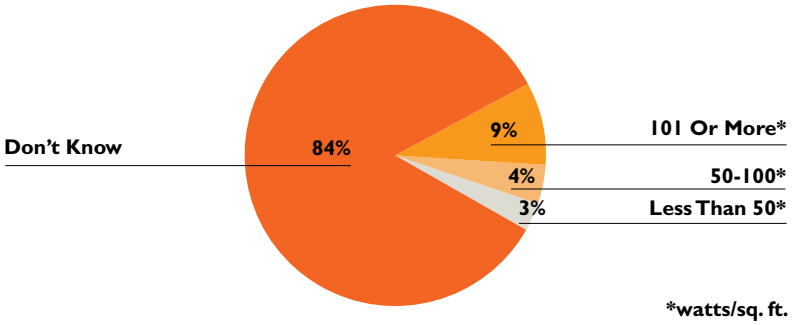

There are also some indications that businesses are not yet taking advantage of power and heat-management tools that they may already have available. One example is processor-level power management—only 14% of businesses indicated that they were taking advantage of processor-level power management if it was available (Figure 8).

Why is this an important omission? To put it bluntly, it's more effective to use power efficiently in the first place than to deal with the aftereffects of waste. And while a processor consumes only a fraction of a server's total power budget, processor-level power management can cut total system power usage by up to 20 percent with no measurable impact on performance.

How is this possible? New processors currently deliver a higher level of computing performance per watt of applied power than ever before. At the most basic level, advanced process technologies like silicon-on-insulator result in fast transistors with low power leakage, minimizing wasted power and heat, while native dual-core design adds performance without increased power. Smart decisions about processor support system placement and connections have also resulted in improved efficiency. For example, an integrated memory controller cuts significant amounts of power from the overall system budget, while a direct connect architecture allows for increased I/O throughput at lower power levels.

TAKING ADVANTAGE OF PROCESSOR-LEVEL POWER MANAGEMENT IF AVAILABLE ON SERVERS FIGURE 8

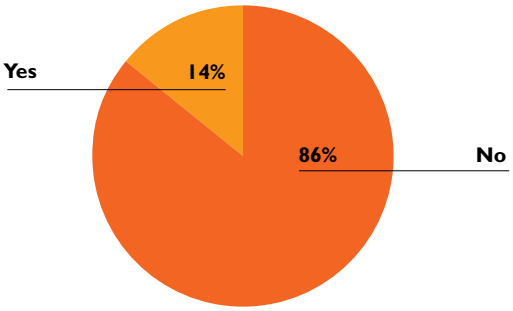

These architectural advantages deliver results automatically, but another potent power management tool—dynamic frequency and voltage adjustment (P-states)—depends on active involvement by the IT department to work to its full capacity. At its highest P-state, the processor runs at full clock speed and voltage, but during off-peak conditions it can be stepped down to "idle," consuming as little as 25 percent of the full-speed power level. Processors like the AMD Opteron™ processor feature multiple P-states, and can be throttled up or down in stages, instantaneously meeting performance needs while preventing unnecessary power use.

Processor power management is usually handled by the operating system, which automatically makes decisions based on performance demands and existing thermal conditions. However, many operating systems offer interfaces through which many of the performance parameters can be adjusted. With a small amount of effort IT managers can tailor the thermal management of their servers' processors to the needs of their data center, and even down to the specific requirements of a single room, rack, or aisle.

While they are not yet all-consuming issues, power and heat management have risen in prominence over the past few years. If the current trends in hardware purchasing and data center construction continue, savvy businesses will need to refine and improve their approaches to power and heat in the data center, taking advantage of more granular planning tools, previously overlooked power-saving opportunities, and more powerful heat management technologies. ■

Research sponsored by AMD