TRENDS

March 8, 2006

Power And Cooling Heat Up The Data Center

by Richard Fichera

TRENDS

Includes a cost model

March 8, 2006

Power And Cooling Heat Up The Data Center

Maximizing Efficiency Yields Both Pragmatic And Social Benefits

by Richard Fichera
with Laura Koetzle and Thomas Powell

EXECUTIVE SUMMARY

After a brief day in the sun in the late ’90s, concerns about power and cooling as critical limitations in the enterprise data center and corporate IT strategy faded into obscurity along with the dot-com economy. Over the past two years, however, as energy prices have soared and servers have grown denser and hotter, power and cooling have once again become critical issues. Improving energy efficiency and solving the problems generated by increasingly dense server form factors is a community effort that requires contributions from semiconductor, system, software, and data center operation vendors. Vendors — of whom some lead the charge and others are being dragged along — have collectively realized that there is money in efficiency; they will vie for energy-conscious customers over the next 24 months. But users shouldn’t wait for vendor products. You can take simple and effective steps today to both improve your overall operations and plan for future technology.

TABLE OF CONTENTS

NOTES & RESOURCES

Forrester interviewed multiple vendors, including: Advanced Micro Devices, American Power Conversion, IBM, Intel, Hewlett-Packard, Liebert, and Sun Microsystems. We also utilized material from conversations with multiple user companies.

Related Research Documents

“Add SWaP To The Server Benchmark Battlefield”
January 19, 2006, Quick Take

“Sun’s T1 Systems Signal A New Life For SPARC”
December 22, 2005, Quick Take

“Sun Announces Niagara Processor”
November 23, 2005, Quick Take

TARGET AUDIENCE

Heads of IT Operations and CIOs.

POWER AND COOLING — LOCKED TOGETHER IN A VICIOUS CYCLE

With the crash of the dot-com economy, visions of immense Internet data centers — and their concomitant cooling and power dissipation implications — evaporated along with much of the tech sector's valuation. However, the problems associated with data center heat and power management have come roaring back over the past two years, thanks to complex interactions between multiple factors, ranging from chip technology to data center design. A combination of rising processing requirements, potentially increased demand for a rapid scaling of processing capacity, energy costs, and the politics of conservation have also driven power and cooling back onto IT's priority list.[1]

At first glance, powering and cooling a data center sounds like a pretty basic proposition. After all, we've been cooling modern office buildings for decades, and modern heating, ventilation, and air conditioning (HVAC) design for offices has become a pretty cut-and-dried trade. The interplay of actors in a large data center, however, is a thornier problem. The generation and eventual disposal of heat is a multidomain problem — multiple levels in the technology stack produce heat, and each level's heat generation has a different impact on the total data center ecosystem. Additionally, the global solution to these problems requires the concentrated effort of a large community of stakeholders, including cooling solution and semiconductor vendors, system designers, software engineers, and data center architects and operators.

The management of heat in the data center is a chain, starting with the chips; moving through the board, enclosure, and rack; and ending with the data center as a whole (see Figure 1). At each step, the basic challenge is to keep that layer of the infrastructure operating at a temperature that is within its limits and to remove the excess heat to an external and cooler environment.

Figure 1 Contributors To Heat In The Data Center

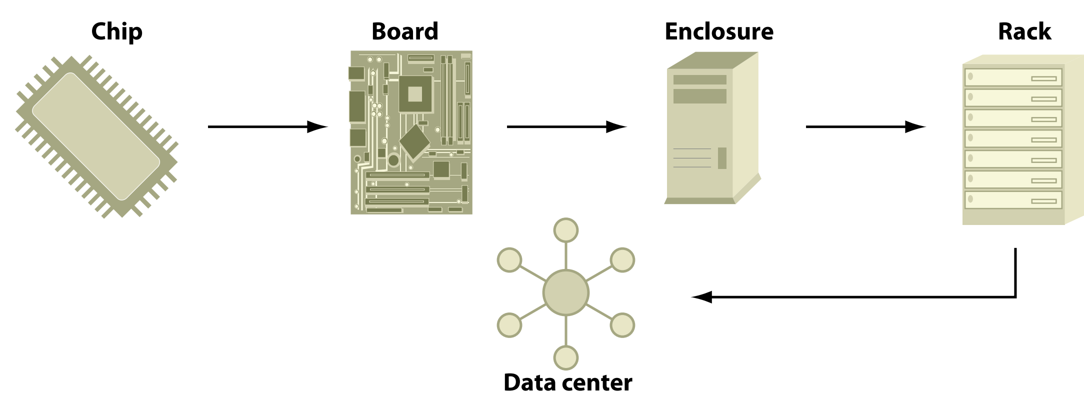

Source: Forrester Research, Inc.

Chip-Level Physics And Manufacturing Decisions Combine To Increase Heat

As chips have become faster, they have become hotter because of both the fundamental physics of switching and the compromises made in the manufacturing process to manufacture the faster gates:

- **Faster switching means more heat.** As you switch a junction at progressively higher speeds, it consumes more power, because you are charging and discharging a not-quite-perfect device more often over the same time period.

- **The manufacturing process is biased toward speed, not power efficiency.** The basic "recipe" for making chips has reinforced fundamental physical limitations.[2] Historically, Intel and most other semiconductor manufacturers have made tradeoffs in their processor designs in favor of allowing higher leakage currents to enable higher speeds. Leakage current is the current that flows through the junctions when the transistor is in an "off" state and in theory should be zero. Why does this matter? Because leakage current is totally wasted energy and has been increasing over time as a fraction of the total power dissipation of processors. Engineers estimate leakage current in the high-end Xeon processors at between 18% and 20% of total power usage and predict that if current trends were to continue, leakage current for a hypothetical 5 GHz chip would be approximately 50%.

- **Successively lower allowable temperatures aggravate the problem.** The final nail in the coffin for progressively higher power dissipation is the temperature limitations for operation of the chips. To protect the increasingly finer features of the chips, manufacturers have lowered the maximum acceptable operating temperatures for the chips — while raising their power dissipation. This has caused increasing problems for the engineers designing systems around those chips. For years, senior system designers have found it hard to lower the case temperature of the chip to keep up with Intel's road map, and Advanced Micro Devices' (AMD) step-function drop in power was a welcome relief.

  To put the problem in perspective, let's compare the heat dissipation of a Xeon chip to that of an electric stove: The Intel Xeon MP with an 8 MB cache server chip has a maximum heat dissipation of 130 watts, and a surface area of approximately half a square inch, for a power dissipation of approximately 260 watts per square inch. The 7-inch heating element on a standard Wolfe electric cooktop dissipates 1,800 watts, for a net dissipation of approximately 47 watts per square inch. Even in its final package, the Intel chip still has a heat dissipation per square inch approximately equal to the cooktop. To complicate matters, designers using the Xeon must keep the temperature of the chip well below boiling, and they can't put a big pot of water on top of it to dissipate the heat!

Board- And Enclosure-Level Challenges — Moving The Heat Around

Once the excess heat has been moved from the chip to the inside of the system enclosure, the next challenge is to remove it from the chassis to the outside environment. Today, all but the largest mainframes use air cooling, so the problem reduces to how to pass enough air across the surfaces of the chips inside the system to transfer the heat to the outside.[3] In addition, other system components — including power supplies, memory, and ironically, the cooling fans themselves — all contribute to power and thermal issues.[4] Designers nibble at the edges of the problem with more efficient power supplies and fans, but as system power dissipation has increased, thermal engineering — which was almost an afterthought for small systems a decade ago — has become a critical limitation for all systems. Increased volumes of very hot air have created a micro version of global warming in the data center.

The Data Center Rack Becomes The Bottleneck

The problems of the internal and the external environment come together in the data center rack. Because today's data centers provision most equipment in standard 19-inch racks, one of the real significant bottlenecks in thermal management occurs there. While we usually describe data centers in terms of aggregate power per net square foot, a critical operational limitation is how much power they can manage per rack (see Figure 2). Here's why:

Figure 2 Increased Power Has Caused Declining Rack Utilization

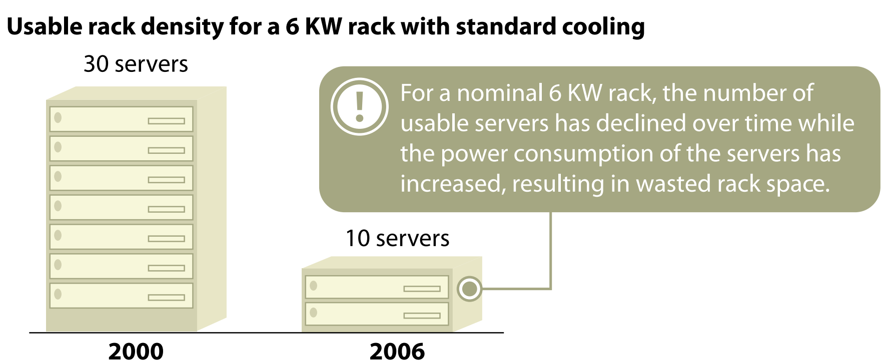

Source: 2006 specs are based on current vendor specifications; 2000 specs are based on observed historical data center practice.

Source: Forrester Research, Inc.

- **Standard rack cooling has limits.** In theory, a standard rack can mount enough equipment to draw in excess of 25 kilowatts (KW) per rack, if populated with dense servers or blade chassis. In practice, it is very difficult to cool a standard rack that dissipates more than about 8 KW of power because of the airflow limitations in most data centers.[5] Well designed and optimized data centers can support in excess of 8 KW per rack with ambient airflow. Generally, to go beyond approximately 12 to 15 KW per rack, either the rack has to have been specially designed or supplemental cooling is required.

- **More power per system means fewer systems per rack — and space-inefficient data centers.** Overall rack throughput can be expressed as performance per watt multiplied by the number of watts per rack. Because the average rack dissipation has remained almost constant at 5 KW per rack or less in real-world data centers, latent problems are often not detected until an attempt is made to implement a dense rack. When additional capacity is added, it then becomes evident that the rack cannot be provisioned to the desired density, making rack power the critical variable for planning new data centers because it forces the deployment of more racks than would otherwise be required or the deployment of supplemental cooling.[6] With the next cycle of blades and 1U dual CPU servers pushing dissipation to approximately 25 to 30 KW per rack in 2006, and with likely increases in subsequent years, racks will remain a critical design point for data centers for the foreseeable future.

A Bird's Eye View Of The Data Center Energy Budget

In looking at the energy budget for the entire data center, users need to factor in the energy for cooling systems as well as the power for the equipment. In theory, the best cooling systems require about 0.3 watts to cool 1 watt of equipment, but real-world results from actual data centers, including all parts of the cooling cycle, run between 0.5 and 1 watt of cooling power for each watt of dissipation.[7] With data center electricity costs averaging approximately $0.11 per kilowatt hour (KWH) in North America, the costs of running a 1,000-server computer room are becoming significant, which makes energy-efficient servers much more attractive (see Figure 3). In addition to energy costs, users must consider the relative costs of building data centers to support dense servers.[8]

BREAKING THE POWER AND HEAT GRIDLOCK — MULTIPLE ACTORS, MULTIPLE LEVELS

Making data centers more power-efficient will require radical improvements in power production by multiple parties, including semiconductors vendors, systems vendors, and software vendors. As a result, users' scope of control will be more limited, but they can still help themselves by intelligently managing cooling and data center design, and most importantly, by making power efficiency a high-priority selection criteria in purchasing decisions.

Figure 3 Model: Data Center Power Cost Calculator

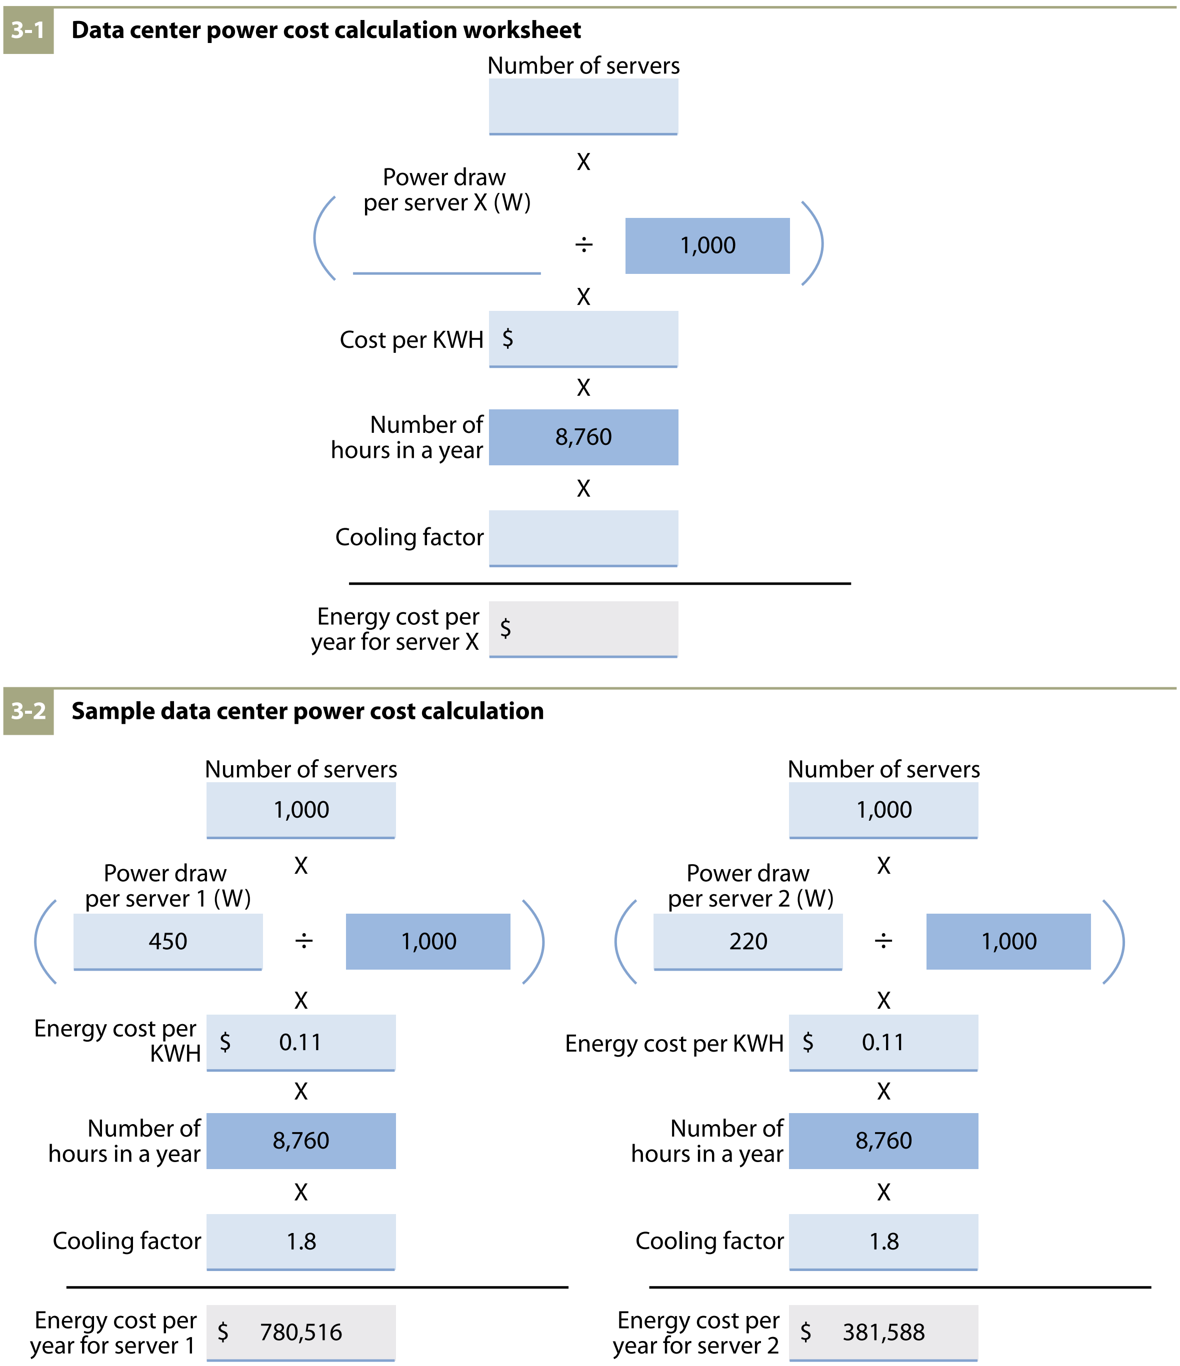

Source: Forrester Research, Inc.

Chip Vendors Will Improve Power Efficiency

The power management chain begins with low-level components. The most important of these is the processor, because in small systems the CPU generally consumes between 40% and 50% of the power. In the x86 world, AMD's Opteron has flourished in part because of its more efficient pipeline (in terms of instructions per clock cycle), its embedded memory controller, and its more power-efficient semiconductor process. To compete, Intel has indicated that it will improve its chip-level power efficiency substantially over the next 12 to 18 months through changes in fundamental design and in the production process. In addition, both Intel and AMD have added features in recent chips to allow systems to selectively change the clock speeds of the chips — meaning that lightly loaded or unused chips consume less power. The only catch is that OS and systems vendors must make design changes to allow users to exploit these features. Thus, system buyers must investigate their system vendors' power management utilities — don't assume that you get all the power management benefits just because you're buying a system with the latest AMD or Intel processor.

Simply making current architectures more power efficient can yield differences in power efficiency approaching 50%. In addition to those efforts, CPU architects are exploring radically different approaches to chip architectures. One good example is Sun's UltraSPARC T1 implementation, which combines multiple small cores, multiple threads, and a lower-speed, power-saving process to yield power efficiencies in the range of between three and five times that of x86 processors under ideal circumstances (see Figure 4).[9] Sun's UltraSPARC T1 and IBM's Cell microprocessor will inspire subsequent generations of multicore chips as an approach to power efficiency, and both Intel and AMD have hinted at multicore technology on tap for future product cycles.[10] Because IBM and Sun control all levels of the stack, their vertically integrated system, software, and CPU stacks will give them an edge over commodity suppliers in integrating power-efficient designs with software that can exploit those designs. Forrester expects AMD to retain a lead over Intel in power efficiency for at least the next 12 months.

Figure 4 Relative Power Efficiencies Of IBM, Intel, and Sun Processors

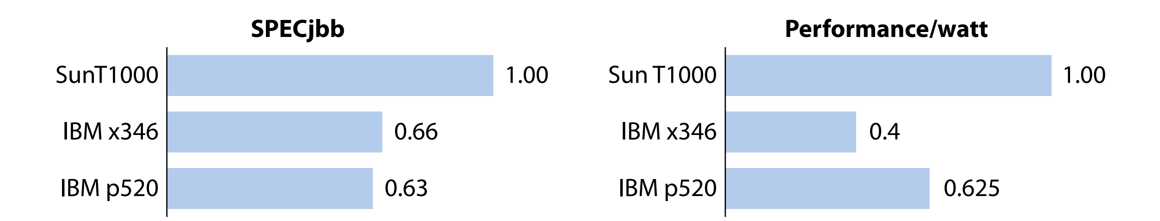

Source: Sun Microsystems' published Standard Performance Evaluation Corporation (SPEC) benchmark results

Source: Forrester Research, Inc.

Software Will Help Manage Chip And System Hardware

Currently, the physical information stored in most configuration management products is not sufficient for optimal power and cooling decisions, but that should change rapidly as vendors rush to add physical management information to their products. Already, both startup vendors and established players are integrating power and cooling information into their configuration management data base (CMDB) products.[11]

System-Level Designs Will Increasingly Focus On Thermal Issues

At the system level, engineers have been struggling for years to refine techniques for thermal management. Designers have borrowed design practices and sophisticated modeling tools from high-end design to use with even the smallest form factor servers. Blades have also served as a turning point for thermal engineering, thanks to the increased density of the blades in their enclosures. The fine points of thermal engineering are:

- **Utilizing airflow design.** Over the past several years, designers have become increasingly sophisticated in designing airflow, focusing on reducing obstructions to airflow, using baffles and plenums to direct the airflow, and designing increasingly efficient fans.[12] Indeed, the degree of sophistication being applied to these seemingly mundane elements is surprising. Fans, for example, don't leap to mind as a high-tech item, but we recently previewed an advanced intelligent fan design that had 25 patents associated with it. In 2006, these small-form-factor, variable-speed fans will be cooling a variety of next-generation systems, including advanced blade enclosures.

- **Managing the heat generation of the systems.** System designers must capitalize on vendor-supplied CPU management capabilities and manage other components, notably fans and disks. All of the major vendors are actively working both to add these power management features to their products and to integrate them into their management stacks for differential advantage. The combined benefits can be compelling. IBM, for example, cites an approximate 40% power savings on its late-model BladeCenter systems versus a similar system without active power management. Sun, with its new UltraSPARC T1 architecture, is now the exemplar of energy efficient processing for thread-rich workloads. In addition, the major system vendors, particularly HP, have been very aggressive in pursuing data-center-wide power management initiatives, which should show tangible results beginning in 2006.

- **Exploring exotic technologies like convective heat pipes.** Convective heat pipes periodically reappear as an answer to the problem of getting the heat off the actual CPU package.[13] Several major systems vendors are exploring variants on heat pipes that involve fluid-cooled chip packages with integrated heat sinks for better heat transfer from the chip die to the heat sink. While these technologies have proven themselves in larger systems, the mainstream design point for systems design will be air cooling of passive heat sinks via directed airflow.[14] Airflow management will become increasingly sophisticated and more tightly integrated with chassis- and rack-level power management.

Data Center Cooling Solutions Will Become More Sophisticated And Modular

The cooling system vendor community has recently begun to focus on providing supplemental cooling for very-high-density racks. Supplemental cooling can either take the form of an enclosed, specialized rack, or modular add-ons that enhance existing data center cooling solutions. Among the offerings currently on the market are systems from American Power Conversion (APC), Liebert, and others.

NERO FIDDLED WHILE HIS DATA CENTER BURNED — BUT YOU CAN DO SOMETHING NOW

Although some data centers show little evidence of evolution, intelligent layout and attention to simple high-level design issues can pay big dividends in data center operations. Users can also take definite steps with today's technologies to minimize their data center heating and cooling costs while waiting for new solutions from systems, chip, and software vendors. Optimize your overall data center architecture and layout, implement spot cooling for problem racks and problem sections, and prioritize heating and cooling issues when selecting new systems and new systems management tools. In tackling these issues, factors to consider include managing airflow, managing local hotspots, and managing server load.

Cooling Optimization Is The First Step

The cooling systems are second only to the systems themselves in their contribution to a data center's electrical load and cost. Therefore you optimizing your cooling systems will improve the energy efficiency of your data center considerably. The first element to grapple with is airflow, because it's both high-impact and cheap to fix if you've gotten it wrong. At a minimum, data center management should look at:

- **Alternating airflow between hot and cold aisles.** A surprising number of small- and medium-size data centers still follow a front-to-back layout (all servers facing the same way, so the hot exhaust of one server becomes the input for another), which results in unacceptable temperature gradients across the room. Simply following a basic alternating hot and cold aisle layout, which virtually everyone associated with data center cooling recommends, can radically improve the thermal performance of a conventional front-to-back airflow layout (see Figure 5).

- **Adding end baffles.** Baffles at the ends of rows can prevent turbulent mixing of hot and cold air at the ends of rows to avoid recycling hot exhaust into what is supposed to be cold intake air.

- **Checking under-the-floor airflow.** Installing underfloor plenums and removing any obstacles to airflow can improve overall balance of airflow across the data center.

Figure 5 Ideal Data Center Airflow

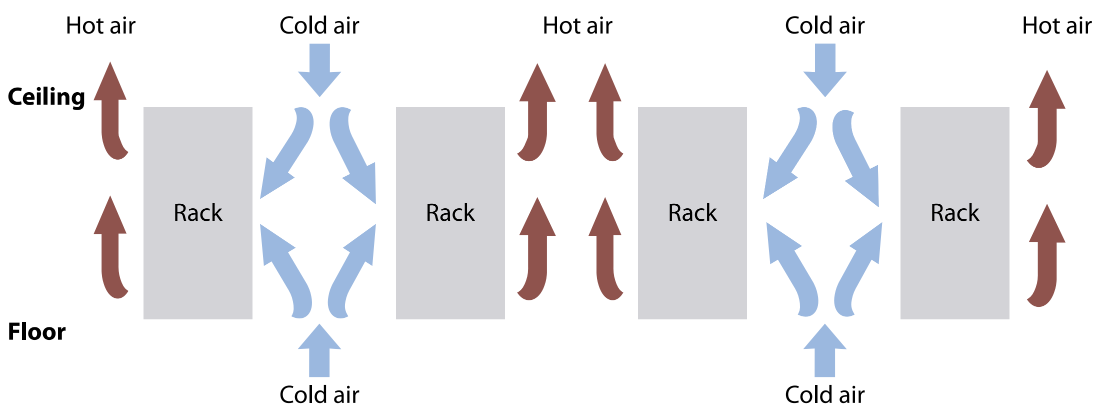

Source: Forrester Research, Inc.

These low-cost steps can help alleviate problems caused by poor distribution of otherwise adequate cooling capacity and can raise the potential for cooling dense racks with ambient air from 6 to 8 KW to well over 12 KW if done correctly. However, to cope with the problems posed by dense rack and blade systems, which can easily exceed 25 KW per rack, you'll need to employ more advanced measures.

Think About Managing Local Hotspots

If you have or anticipate local hotspots — such as dense racks of blades or other equipment with per-rack dissipation that might overload the limits of the data center airflow — plan for local cooling assists. Options include specialized, water-cooled enclosures; spot cooling; and "room within a room" technology from vendors such as APC and Liebert and from the systems vendors themselves. Currently Egenera, HP, and IBM offer specialized cooling assists.[15] For hot spots within a larger data center room, supplemental cooling like Liebert's XDV and XDO systems will work well, and using enclosed racks like Egenera's, HP's, or IBM's fluid-cooled racks can make a major difference in the local heat budget on the data center floor. Several sources report that local cooling and/or closed loop rack cooling can save in excess of 20% over cooling with conventional whole-room systems. The choice of technology will depend on whether you have chilled water available and the total BTUs of the rack, among other issues. Proper hotspot management will not decrease the overall data center cooling load but will prevent the hotspots from interfering with other areas. This will help you avoid overprovisioning globally with cooling equipment.

Another architectural approach to managing hotspots is to segregate high- and low-power regions of the data center and apply supplemental cooling to the entire row, as do the new systems from APC and Liebert. Although this solution may require you to rearrange more of the data center, it allows an entire dense row to be cooled efficiently without interfering with other portions of the data center.

Actively Manage Workloads For Power Efficiency

Because system power consumption is not linear with workload, there is an intersection between workload management and overall data center power and cooling. In advance of the next wave of advanced embedded power management technologies, the overall power requirements of installed systems can be actively managed by managing system loads. Because a typical x86 server consumes between 30% and 40% of its maximum power when it's idle, running systems with very light workloads wastes power. Thus, increasing the average utilization of the servers can yield significant benefits in overall operational efficiency.[16] You can increase utilization by consolidating workloads — which has become substantially easier with system virtualization technology — or by managing workloads more actively. Workload management can be manual, using either in-house scripts or vendor-supplied tools, or it can be dynamic and automated.[17]

In addition to load management tools, users should look at the first generation of data center power and cooling management software. This includes products such as the forthcoming APC InfraStruXure Manager (to be introduced in 2006), which will offer the ability to manage against available power and cooling resources either by rack or by row.

RECOMMENDATIONS

USERS SHOULD FOCUS ON EFFICIENCY

With an increasing supply of power-efficient technology and the prospect of long-term high energy costs, users should make power efficiency an integral part of their data center planning processes. Procurement specialists should:

1. Require information on power and cooling costs.
2. Include power and workload metrics like Sun's space, watts, and performance (SWaP) metric in their evaluation criteria.

Vendors will have to respond to your pressure with continued investment in power-efficient systems at multiple levels. Such efficiency will become expected, and users will demand it without a price premium.

Users with current data center cooling problems — or plans to implement dense next-generation racks — should investigate supplemental cooling from their system vendors or from independent vendors. At a minimum, ensure that your airflow and current cooling systems are operating as well as they possibly can. If you need help, cooling systems vendors, major systems vendors, and systems integrators can all provide consulting services for data center cooling.

Users in the process of designing new data centers will need to examine the tradeoffs between denser floor space that requires more careful management and traditional low-density designs. Plan carefully for modular upgrades of cooling capacity in anticipation of increased loads.

WHAT IT MEANS

POWER AND COOLING WILL ASSUME A HIGH PROFILE

Power efficiency will remain a high-profile issue because:

1. Improving power efficiencies across the enterprise IT infrastructure means significantly lower energy costs.
2. Vendors that lead early will gain mind and market share.
3. Governments will produce more regulation, not less.

Thus, power efficiency will permeate vendor planning from chips through system design and operating software, and users will become increasingly sensitive to it as a purchasing criterion.

The need for optimized modular data centers will force closer cooperation between facilities design, data center operations, and line-of-business management, particularly because superoptimized designs may require periodic modular changes in floor space and cooling capacity to handle spikes. The computer systems, component and data center construction, and cooling vendors will collectively have to develop new practices for load forecasting and facilities management to provide that modularity.

ALTERNATIVE VIEW

THE PHOENIX CRASHES AND BURNS — MACROECONOMIC AND BEHAVIORAL PITFALLS

One obvious hazard to the entire grand vision of an increased focus on power and cooling efficiency is a major downturn in energy prices. If that were to happen, users would be very unlikely to alter their purchasing behavior simply for "green" reasons, and vendors who had invested heavily in power efficiency at the expense of raw performance and low prices would be left holding a fringe position, particularly if they had successfully positioned themselves as being the power efficiency standard-bearers. Additionally, if the total cost and complexity of really optimized global power management remains stubbornly high, confused and intimidated users will merely select incrementally better versions of what they currently have. However, even if energy prices decline, users will still have to cope with the demands of increased rack-level power dissipation.

SUPPLEMENTAL MATERIAL

Online Resource

The online version of Figure 3 is an interactive tool for comparing server energy costs given input for the number of servers, power draw per server, average hourly energy cost, and cooling factor coefficient.

ENDNOTES

1 Predicting the future of both the regulatory climate and the future cost of energy are inherently risky, but the cost of energy has been rising steadily, with occasional upward (never downward) shocks for decades. Reasonable macroeconomic forecasts seem to show that diminishing supplies and increasing third world demand will more than offset any increased conservation efforts by developed nations, keeping upward pressure on oil prices (the basic proxy for energy cost). And that's before we even think about little things like the seemingly eternal volatility of the oil producing regions. Prudence would dictate assuming that $60 per barrel is now a floor price for oil, not a ceiling.

2 The details are beyond the scope of this document, but semiconductor manufacturing is still not an exact process. Rather, it is one that yields, in each batch, a range of outcomes — rather like baking cookies and having them come out in a range of doneness and size. The details of the process — the combination of chemical, mechanical, and sequential operations needed to produce the chips — is the core intellectual property of the manufacturers. In designing a process, engineers determine tradeoffs between yield percentages (a major factor in cost), leakage current, and speed, which they fine-tune as the devices are manufactured. Once the process is mature, companies do not like to change it, preferring to maintain the same basic process over several generations of decreasing chip geometries.

3 Even something as seemingly mundane as the design of the innocuous-looking copper heat sinks that sit on top of processors and other hot chips has become a sophisticated specialty. A designer of these chips cited the need for 3D solids modeling, finite element analysis, and fluid dynamics models in addition to heat transfer models. Choices about the width of the fins, interruptions to deliberately induce turbulence, and other seemingly simple variables can make a difference in the efficiency of the heat transfer by a factor of two. At least one startup is pursuing "microfin" technology, in which the fins and the spaces between them are on the order of 100 nm wide (approximately the thickness of a human hair) and are fabricated like integrated circuits. Their goal is to improve the efficiency of the heat transfer by a factor of about 10.

4 The latest small-form-factor disks are actually relatively low-power, but their mass can complicate airflow design.

5 The average rack density of legacy equipment is substantially lower than 8 KW per rack, with some estimates placing the average data center at approximately 5 KW per rack. To even get to 8KW per rack requires proper layout of hot and cold aisles and good airflow management within the data center.

[6] IBM, in designing the Blue Gene/L, the undisputed leader in performance per rack for selected problems, focused very much on the overall power versus performance problem. The company has demonstrated systems dissipating in excess of 20 KW per rack with its proprietary rack and chassis, which has been heavily optimized for cooling efficiency.

[7] Both Liebert (a division of Emerson) and APC claim that their newest supplemental cooling modules can operate at best case overheads of less than 0.25, and that they will operate closer to their theoretical limits than conventional data center air conditioning systems. These are for just the local heat exchange, not the entire cooling load including building chiller loads.

[8] Users need to be clear on the difference between building for a current load and building for future loads or consolidated loads. A dense data center to support a current load will cost more per square foot but can actually save money because it will be smaller and can potentially be cooled more efficiently if the cooling is designed correctly. However, most organizations that go to the trouble of building completely new space are also anticipating increased processing requirements and will build beyond their current requirements. A modular approach to provisioning and cooling the incremental space as it is needed can pay off with major gains in efficiency. According to several HVAC experts, one of the major defects in many data centers is that they are just too big, and if made smaller, the inventory could be cooled much more efficiently.

[9] On December 6, 2005, Sun Microsystems announced the first systems based on the newly announced Niagara CPU called the UltraSPARC T1, the first of Sun's new Chip Multithreading Technology (CMT) architecture products. The new systems — the Sun Fire CoolThreads T1000 and T2000 servers — deliver on the early promise of Sun's T1 technology, offering impressive performance compared to both legacy SPARC and current RISC and x86 servers. In addition to raw performance, the systems are extremely power efficient. Sun is attempting to capitalize on the systems' power attributes by introducing a new benchmark metric, SWaP (space, wattage, and performance), that focuses on the data center macro-environment. While we applaud the concept, users must carefully interpret the results to make sure that the metrics make sense in their environment. Taken as a whole, the new products embody one of the most disruptive technology shifts in the microprocessor arena in several decades and one that will strongly influence future design efforts across the industry. See the December 22, 2005, Quick Take "Sun's T1 Systems Signal A New Life For SPARC."

[10] The IBM Cell processor is a very specialized architecture, with one general purpose POWER CPU plus eight independent coprocessors connected by a high-speed fabric. Each of the coprocessors is a fully functional core with a modified POWER instruction set, and each can operate independently, making the chip align with the traditional MIMD (multiple instruction/multiple data) model. The independent processors do not have the overhead of memory coherence, and although they are very efficient, they are also difficult to program. The chip excels at parallel real-time applications such as gaming (its initial go to market segment) and will probably establish a niche presence in applications like seismic processing, signal processing, and real-time command and control.

[11] Actually, power management is relatively mature, with multiple products that can track the power budget of a rack and relate it to available power. The horizon step is integrating the power management information with predictive cooling management — forecasting thermal problems and recommending solutions.

[12] A plenum is a pressurized chamber used to feed air to buffer irregularities in the airflow as it is distributed to its point of use. A baffle is a barrier used to channel or concentrate airflow.

[13] A convective heat pipe is an enclosed pipe that uses internal convective fluid flow to transfer heat from the source end to the cool end. Heat pipes are many times more efficient than the same mass of solid copper for transmitting heat and are usually used to move heat from an internal location to a more easily cooled heat sink. Long a staple of military designs, they had a brief renaissance in small-form-factor servers, but they are being replaced by more efficient forced-air cooling for cost reasons.

[14] Experiments with liquid cooling will not go away, and as successively finer chip geometries further strain the limits of air cooling, liquid cooling will reappear. IBM, for example, has been experimenting with etched microchannels in the actual chip die, through which they would run coolant.

[15] In many cases, the vendors are in fact OEMing components and assemblies from the major suppliers, particularly Eaton Electrical and Liebert.

[16] Studies cited by Intel at the fall 2005 Intel Developer Forum suggested that aggressive load management could trim 20% from the operation of a server farm.

[17] Tools fall into several categories, including RISC/Unix server virtualization and workload management, such as HP's Virtual Server Environment, Workload Manager, Process Resource Manager, and Instant Capacity On Demand, as well as VMware's Virtual Center; Web site and app server load balancing; and emerging rapid provisioning and service-oriented architecture middleware from vendors like Cassatt, DataSynapse, and Scalent Systems. Other provisioning systems such as BladeLogic and Opsware can also be adapted to assist this task.

FORRESTER

Helping Business Thrive On Technology Change

Headquarters

Forrester Research, Inc.
400 Technology Square
Cambridge, MA 02139 USA
Tel: +1 617/613-6000
Fax: +1 617/613-5000
Email: forrester@forrester.com
Nasdaq symbol: FORR
www.forrester.com

Research and Sales Offices

Australia
Brazil
Canada
Denmark
France
Germany
Hong Kong
India
Israel
Japan
Korea
The Netherlands
Switzerland
United Kingdom
United States

For a complete list of worldwide locations, visit www.forrester.com/about.

For information on hard-copy or electronic reprints, please contact the Client Resource Center at +1 866/367-7378, +1 617/617-5730, or resourcecenter@forrester.com. We offer quantity discounts and special pricing for academic and nonprofit institutions.

Forrester Research (Nasdaq: FORR) is an independent technology and market research company that provides pragmatic and forward-thinking advice about technology's impact on business and consumers. For 22 years, Forrester has been a thought leader and trusted advisor, helping global clients lead in their markets through its research, consulting, events, and peer-to-peer executive programs. For more information, visit www.forrester.com.

FORRESTER