TECHNICAL BRIEF

Enabling Technologies for Power and Cooling

Sponsored by: HP

Jed Scaramella
September 2006

Introduction

In June 2006 Hewlett-Packard (HP) announced its next-generation blade portfolio, the BladeSystem c-Class. The new BladeSystem was designed to address some of the key total cost of ownership (TCO) issues facing today's datacenter, including server management costs, interconnect complexity, and power and cooling. As part of the BladeSystem c-Class, HP introduced three new technologies:

- HP Insight Control management
- HP Virtual Connect architecture
- HP Thermal Logic technologies

These innovations play a central role in reducing overall datacenter operating expenses, and differentiate the BladeSystem c-Class system, both from competitive blade offerings and from rack-optimized server infrastructures. In this technology brief, IDC focuses on HP Thermal Logic, a complete system solution to current power and cooling challenges. It complements the IDC White Paper *Forecasting Total Cost of Ownership for Initial Deployments of Server Blades*, which describes the overall TCO advantages of the BladeSystem c-Class portfolio, as well as the technology briefs *Enabling Technology for Blade I/O Virtualization* and *Enabling Technologies for Blade Management* that focus on the operational and cost advantages of HP Virtual Connect architecture and HP Insight Control management, respectively.

Power and Cooling Emerge as Critical Factors for IT Executives

In recent years, the rate of server technology advancement has outpaced the datacenter's ability to support these systems, especially in terms of power and cooling. Historically, the objective of IT executives was to maximize their compute resource performance, and the associated expense of power and cooling was simply tied to the cost of doing business. But the dynamics are shifting, with processing becoming more commoditized, and power and cooling becoming the limiting factors that are increasingly necessary to optimize.

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

Several key factors have contributed to the increasing importance of managing power and cooling:

- **Increased system performance.** Requires greater power and cooling per server.
- **Shift toward high-density computing.** Increases the power and cooling burden at the rack level.
- **Server proliferation.** Increases the overall requirements for power and cooling throughout the datacenter.

Increased System Performance

Processors have held true to Moore's Law, the 1965 prediction by Intel cofounder Gordon Moore, that transistor density would double every 18 months. While this has been a boon to meeting the needs of demanding business users with complex applications, it has translated into an increase in the power consumption and cooling required for servers. Further, similar improvements have been seen in system components, such as internal memory and hard drives, which have further added to the burden of powering and cooling servers. To illustrate the problem, an average server today consumes 400W of power, compared to an average of 100W 10 years ago. With this quadruple increase, not only do datacenters have to supply the additional power, but they must also provide sufficient cooling capacity for the additional heat generated.

Shift Toward High-Density Computing

Growth in business workloads has put pressure on IT to deploy additional servers within their datacenters; unfortunately, datacenters are finite in their capacity, and building out new capacity is cost prohibitive to most companies. Large or small, all environments have a finite amount of power, cooling, and space they can support without costly upgrades. One widely adopted solution to floor space concerns is the deployment of servers in increasingly compact form factors. IDC estimates that server system density has increased by 15% annually over the last 10 years as companies have moved from pedestal servers to rack-optimized systems and now to widespread adoption of blade servers (Figure 1). In surveys of end users, IDC finds that on average companies deploy 14 servers per rack. This is up from an average of 7 servers deployed per rack only 10 years ago. Looking forward, the increased adoption of the blade form factor will push the average to 20 servers per rack. However, IDC has had discussions with many customers that anticipate more than 60 blades in a rack.

FIGURE 1

Worldwide Server Installed Base by Form Factor, 1996–2010

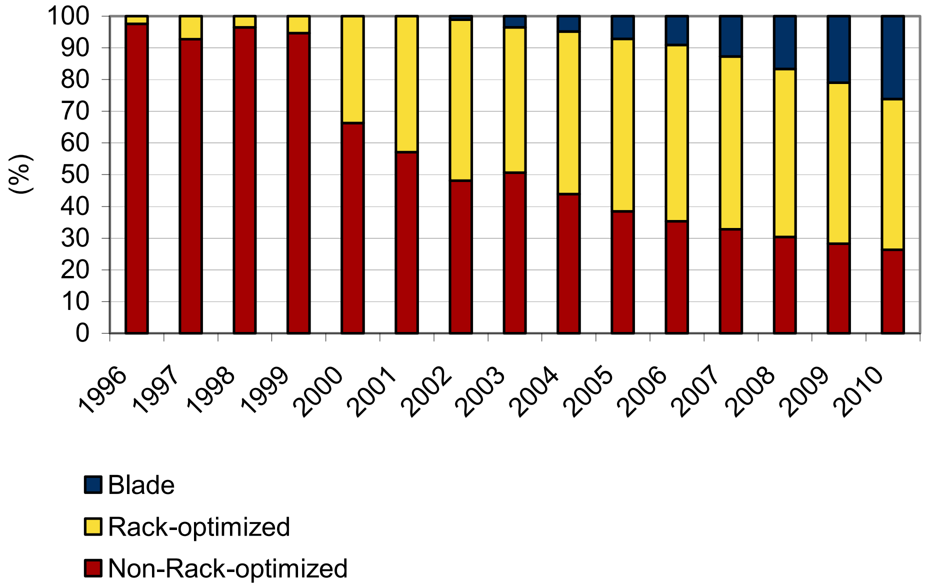

Source: IDC, 2006

This shift toward smaller form factors has increased the pressure on power and cooling management at the rack level. While the average power consumption per rack in 2000 was 1kW, datacenter managers today must account for 6–8kW per rack and must plan to manage over 20kW in the next five years. The trend toward high density has resulted in hot spots within the datacenter that are subject to failures and reliability concerns.

Server Proliferation

Even while servers are becoming more powerful and compact, the sheer number of them in the datacenter has exploded (Figure 2). There are two primary reasons for this server sprawl:

- **Expanding server footprint.** Most core business applications today are supported by multiple servers. Two, three, or even five or more servers are not uncommon. And with new workloads such as email and Web-based applications coming online every day, these only add to the number of servers an organization must support.

- **Shift in the server mix.** Twenty years ago an organization might have purchased a handful of mainframes and could have reasonably expected those systems to handle all of its IT needs. However, as pressure was building to reduce initial acquisition costs and new lower-priced technologies continued to be introduced into the market, customer buying patterns evolved accordingly. With each technology transition, the cost of acquisition came down by an order of magnitude, and today x86 systems average about $4,000 per system. These lower price points allow customers to distribute systems more widely throughout their organizations, to the point where it is not uncommon for datacenters to support 5,000 or more servers, most deployed as a single server per application. To compound the problem, businesses often request that IT observe certain maximum configurations of processors, memory, and drives per server to provide sufficient headroom for future application requirements.

FIGURE 2

Worldwide Server Installed Base, New Server Spending, and Power and Cooling Expense

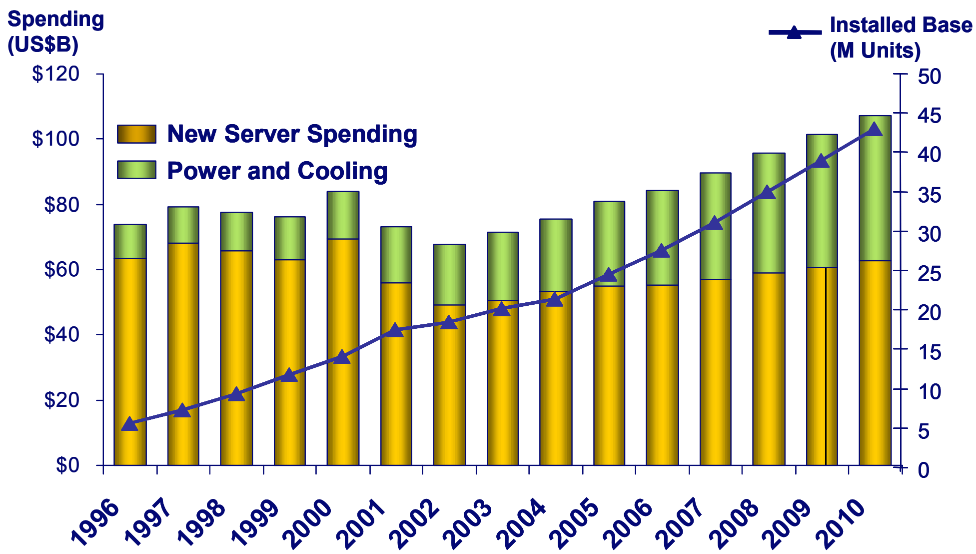

Source: IDC, 2006

As a result of these three factors — increased performance per server, greater server density, and increased full-featured server proliferation — the amount of power and cooling required to most organizations' installed base of servers has grown dramatically. This has added significantly to IT's burden, both in terms of management complexity and operating costs. In fact, by 2010 IDC expects that for every dollar of new server spending, an additional $0.70 will be needed for power and cooling. This is a dramatic rise from the rate of $0.50 in 2005 and $0.21 in 2000.

The Goal: Measurement and Management

In several surveys over the past six months, IT executives consistently rank power and cooling among the top items of concern. This is in stark contrast to even a few years ago, where the thermal characteristics of servers would hardly register a blip on the radar screen. Specific management challenges include:

- **Improved monitoring.** As the saying goes, "You can not manage what you can not measure," surprisingly, however, many IT executives and datacenter managers do not even know how much power their IT equipment consumes. The facilities manager receives a single electricity bill and is unable to break out IT consumption from that of other operations. As server power costs continue to become more significant, datacenter managers need the ability to determine which servers systems or group of systems within their datacenters are consuming large proportions of energy.

- **Power efficiency.** Once IT managers have an understanding of systems' thermal characteristics, they can begin to source more efficient server systems. Power efficiency is particularly important as savings are realized both in server power itself and in cooling. Every watt consumed by a server dissipates 1W of heat, which in turn requires an additional watt of energy to cool. Therefore, every watt reduced from a server's initial consumption will save the organization more than the cost of 1W when looking at both power and cooling. Furthermore, as new servers are deployed, they are required to fit within the thermal capacity of the existing datacenters. Deploying too many power-hungry systems could exceed the current limitations, forcing a disruptive datacenter retrofit or even expensive new facility construction.

- **Effective cooling.** Even as organizations strive to reduce their server power consumption, there will always be a need for effective cooling solutions. Many datacenter managers find that they can usually add more power as long as they are willing to pay the price, but that their datacenter's cooling capacity is often the limiting factor. Cooling constraints can restrict IT expansion, potentially even inhibiting the overall capabilities and growth of the business. The challenge for server manufacturers is to develop new designs and technologies that provide an integrated approach to cooling, from component to system architecture to fan design.

HP Thermal Logic

For the past decade, HP has been at the forefront of research and solutions for power and cooling initiatives. In 1996 HP formed the "Cool Team," a companywide, cross-departmental team to focus on the challenges of power and cooling. With an eye toward providing comprehensive power and cooling benefits for the datacenter, the team made a holistic effort to examine the issue at three different levels: the component level, the system level, and the overall datacenter level. The objective of the team is to work with HP Labs to integrate power efficiency expertise into new HP solutions at the initial design phase and to own thermal management throughout the development process.

The new HP BladeSystem c-Class incorporates HP Thermal Logic technology, which combines new innovations in monitoring, reporting, and adaptive management functionality for power and cooling. HP Thermal Logic takes a holistic approach to power and cooling issues, from processor to enclosure design and from architecture to management. Thermal Logic attempts to pool and share power and cooling resources, then use management and thermal design to efficiently deliver those resources based on the performance level required.

Management, Reporting, and Control

HP Thermal Logic technology enables IT managers to quantify power consumption and cooling needs across the datacenter at the level of individual servers, racks, or even groups of systems. Hundreds of sensors located throughout the blade enclosures provide views of previously unavailable power and cooling data, all easily available through HP Onboard Administrator and HP System Insight Manager. These data are aggregated in a graphical way to produce real-time or historical reports that can display heat output, air temperatures, and power consumption for each server, enclosure, or rack.

IT can now "right size" and adapt its power consumption based on workload demands and set temperature or power thresholds. Thermal Logic provides an adaptive infrastructure whereby IT can control its power consumption and cooling-adaptive infrastructure either manually or through policy-based systems. This enables IT to finally gain control over the thermal dynamics of its server systems.

Power Efficiency

New technological advancements in Thermal Logic ensure optimal efficiency of power consumption across multiple systems by pooling and sharing power and adapting power draw based on demand. HP's Dynamic Power Saver and Power Regulator address the problem on multiple levels, including power supplies and component consumption.

By pooling and sharing fully redundant power distribution, blades have an inherent power-efficiency advantage over rack servers. This eliminates multiple power supplies per server without sacrificing reliability. The systems also hold a natural advantage by removing cables from the back of the rack to improve airflow.

The Dynamic Power Saver feature provides for efficient use of power in the server blade enclosure. Power supplies are placed in a standby mode when the power demand from the server blade enclosure is low. The standby power supplies incrementally activate to deliver the required power as demand increases, which enables the system to continuously operate at optimum efficiency.

Dynamic Power Saver recognizes that most power supplies operate less efficiently when lightly loaded and more efficiently when heavily loaded. A power supply installed in the enclosure running with a 10% load could have efficiency as low as 50%, but with a 50% load, efficiency increases to almost 90%, providing a significant savings in power consumption. Dynamic Power Saver drives power-supply efficiency by automatically shifting the load so there is a heavier load on fewer power supplies as opposed to a lighter load on all the installed power supplies. This feature could be useful for systems with dynamic workloads that change throughout the day, in high-utilization applications such as virtual machines, or where top performance is a requirement.

For workloads under low utilization, the Power Regulator increases power efficiencies at the processor level by adapting performance levels to workload in order to save power. More than any other component, processor chips have been responsible for a greater increase in server power consumption; therefore, addressing the problem at the processor level provides the greatest opportunity for power savings. The Power Regulator works to match processor speed to workload. When combined with the Dynamic Power Saver feature it reduces the overall power consumption of the enclosure and maximizes the enclosure efficiency. Reducing power consumption, both at the processor and server levels, results in energy savings that translate directly into financial savings for the organization.

Adaptive Cooling

HP drew upon the expertise of the Cool Team and HP Labs to optimize the design of the individual server blades and components, as well as the total cooling system technologies, to improve cooling effectiveness and energy efficiency across the system.

Each HP ProLiant server blade follows a design principle to direct the airflow over the "best path" within the server. For example, the coolest air is first directed over mechanical devices, such as hard drives, before being directed to processors or memory. New, smaller heat sink designs were created to improve cooling capacity, reduce airflow, and leave room for more server features such as hot-plug drives and extra memory slots.

The BladeSystem c7000 enclosure is designed to be airtight and to direct airflow from front to back in the most efficient manner without using excess cooling capacity. The Parallel Redundant Scalable Enclosure Cooling (PARSEC) architecture of the enclosure is a multizoned architecture that utilizes multiple sensors in each zone. Each enclosure is divided into four zones consisting of blades and fans. In each zone the fans directly cool the blades in that zone but also serve as backup cooling for the entire enclosure. Administrators can adjust the cooling configuration depending on the types of blades in the enclosures. For example, the amount of cooling provided to storage blades, which require less cooling than server blades, can be reduced to save on associated energy costs. Empty blade slots can also be isolated from the rest of the system in order not to waste cooling and power.

Thermal Logic employs innovative HP Active Cool Fans, leveraged from aircraft technology, that are able to provide dynamic levels of airflow to match heat loads. The fans generate high-pressure differential with good airflow at low power, which enables 10 fans to cool 16 blades using only 120–140W of power. Active Cool Fans deliver energy benefits by drawing less power than that required for comparable rackmounted server fans and use less airflow to achieve the same cooling capacity, reducing the demands on the datacenter cooling as a whole.

By incorporating Thermal Logic into the design of its new BladeSystem c-Class from the ground up, HP is able to optimize airflow, acoustics, power, and performance. This enables administrators to continue to increase server density in their datacenters, while lowering the costs to power and cool them and reducing the risk of systems overheating.

IDC Analysis

Power and cooling are top-of-mind issues for IT executives. As IT continues to support a growing number of servers, power and cooling have emerged as limiting factors to the number of servers that can be squeezed into the datacenter. Solutions to optimize server power and cooling can deliver real benefits and reduce overall IT TCO.

For all the innovation and benefits associated with Thermal Logic, however, barriers also exist. These include issues in power and cooling solution standards and the necessity to demonstrate a track record in power and cooling solutions.

Enabler for the Adaptive Infrastructure

HP is staking a claim with Thermal Logic as an enabler for the adaptive infrastructure, allowing IT managers to flexibly adjust their power consumption and cooling to meet the current needs of their users in an energy-efficient manner. HP plans to apply these principles and leverage the same technologies beyond blades to other system designs, management tools, and datacenter solutions and services. Just as server virtualization enables clients to increase utilization rates by pooling compute resources, HP Thermal Logic provides some of the same benefits, enabling administrators to provision power and cooling from a pooled supply based on the specific needs of supported systems. By providing the capability to adjust power and cooling resources and adapt them to the organization's changing requirements, Thermal Logic allows IT administrators to optimize infrastructure resources, match their power and cooling capacity to the workload, and improve overall TCO.

By providing more granular visibility into the thermal dynamics of server systems, Thermal Logic enables IT managers to better understand the dynamics of power use. "Power hog" servers can be isolated and tuned or replaced. Users responsible for more than their share of compute — and therefore power and cooling — resources can be billed appropriately.

Finally, Thermal Logic enables IT managers to tailor their power and cooling systems to the needs of their systems, eliminating wasteful over-provisioning of power and cooling capacity to individual systems. This saves energy costs and extends the life cycle of the existing infrastructure.

Power and Cooling Standards in the Datacenter

A challenge for HP is to not only prove the cost benefits of the BladeSystem c-Class but also demonstrate how Thermal Logic enables clients to deploy a standard, scalable solution for power and cooling throughout the datacenter. Even though power and cooling are priorities for IT executives, they reveal a relatively new market for server vendors. And unlike in a mature market, the standards are not clearly defined. Multiple server vendors are pushing their own metrics and standards that will show their products in the best light. This will create challenges for IT organizations that try to adopt and integrate power and cooling solutions from different vendors.

Need to Demonstrate Credibility in Power and Cooling Solutions

In entering the power and cooling arena, HP will also find itself facing non-traditional vendors such as facilities vendors (Liebert, APC). These vendors have a long-standing relationship with both datacenter and facilities managers. Additionally, Liebert and APC have ramped up their own power and cooling product offerings to include more sophisticated power management tools and modular cooling units. The question is where customers will look to solve their thermal concerns.

In this emerging and as yet undefined market, HP must establish its reputation and educate clients as to how Thermal Logic effectively solves the thermal concerns within datacenters, when compared to both traditional and non-traditional competitors.

Conclusion

As the next generation of HP blade server architecture and technology, the BladeSystem c-Class provides a step forward in increasing overall server efficiency and minimizing TCO. One of the factors enabling these efficiencies is Thermal Logic, an innovative set of technologies combining monitoring, reporting, and adaptive management functionality of power and cooling resources, combined with new power and cooling design innovations within the BladeSystem c-Class enclosure.

By allowing IT managers to measure, control, and optimize their use of power and cooling resources to match variable degrees of workload and utilization, HP Thermal Logic is an enabler for the adaptive infrastructure. By raising the ceiling on cooling as the limiting factor for the number of servers that can be fit into a datacenter, Thermal Logic allows IT to better scale compute resources to meet the growing needs of the organization, and by optimizing the power and cooling required per server, it enables IT to reduce its overall TCO.