

## TECHNICAL BRIEF

---

### Next-Generation Power and Cooling for Blade Environments

Sponsored by: HP

---

Jed Scaramella  
December 2008

---

#### Introduction

Hewlett-Packard (HP) designed the HP BladeSystem portfolio to address some of the key total cost of ownership (TCO) issues facing today's datacenter, including server management costs, utilization, and power and cooling. Since the launch of the HP BladeSystem c-Class family in 2006, HP has risen to claim the number 1 spot in the blade server market. As part of HP BladeSystem, HP features three fundamental technologies:

- ☒ HP Insight Software for infrastructure management
- ☒ HP Virtual Connect for virtual I/O networking
- ☒ HP Thermal Logic for power and cooling

These technologies play a central role in reducing overall datacenter operating expenses and improving flexibility and control of infrastructure resources. They also differentiate HP BladeSystem both from competitive blade offerings and from rack-optimized servers.

In this technical brief, IDC examines HP Thermal Logic, a holistic approach to improving power and cooling efficiency in the IT infrastructure, designed to help companies reduce energy consumption, reclaim unused datacenter capacity, and extend the life of their datacenter. It joins the IDC technical briefs *Next-Generation Management Software for Blade Environments* and *Next-Generation Technology for Virtual I/O and Blade Servers*, which focus on the operational and cost advantages of HP Insight Software and HP Virtual Connect, respectively.

---

#### Power and Cooling Emerges as a Critical Factor 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 performance and expand the availability of IT resources; the associated expense of power and cooling was simply seen as a cost of doing business. Today the dynamics have shifted, with energy budget having become the primary limiting factor in the datacenter environment in terms of both financial budgets and the expansion of IT capacity. Several factors have contributed to the increasing importance of managing power and cooling, including an increase in system performance, a shift toward high-density computing, and a shift within the server installed base.

### ***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, innovations such as virtualization and multicore have increased the need for larger amounts of internal memory, the second-largest consumer of power behind processors, further adding to the burden of powering and cooling servers. This all adds up to a drastic increase in the necessary power consumption of servers. Further complicating the issue is that datacenters not only have to supply the additional power but also must 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 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 past 10 years as companies have moved from pedestal servers to rack-optimized systems and now to widespread adoption of blade servers. 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.

This shift toward high-density computing has increased the pressure on power and cooling management at the rack level, to the point where power and cooling, and not physical available space, has become the limiting factor driving server capacity in many datacenters today.

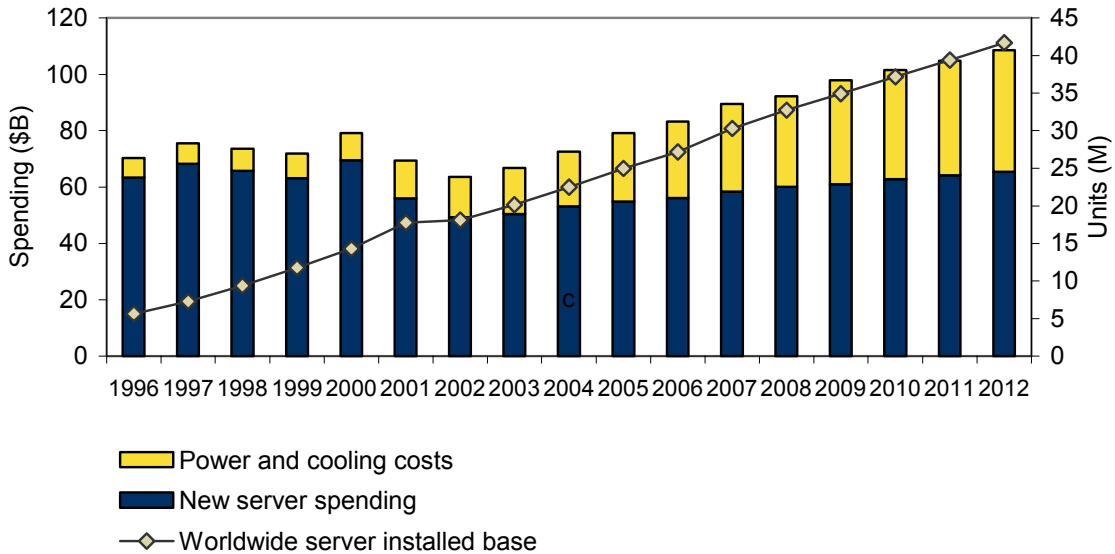
### ***Shift Within the Installed Base***

Even while servers are becoming more powerful and compact, the sheer number of servers in the datacenter has exploded (see Figure 1).

Ten or twenty years ago an organization might have purchased a handful of mainframes or Unix-based systems 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 under \$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, many still deployed as a single server per application.

**FIGURE 1**

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



Source: IDC, 2008

As Figure 1 illustrates, the cost to power and cool server systems has grown to a point where for every \$1.00 of new server spending, an additional \$0.48 is needed for power and cooling. This is a dramatic rise from the rate of \$0.24 in 2001, and IDC expects this ratio to rise to \$0.66 by 2012 under current conditions. IDC has found that power and cooling has become the determining factor limiting the capacity of most datacenters. This has added significantly to IT's burden, in terms of both management complexity and operating costs.

***The Goal: Measurement and Management***

In recent IDC surveys, IT executives consistently ranked power and cooling as their top item of concern. This is in stark contrast to several years ago, when the thermal characteristics of servers were often treated as an afterthought. Specific management challenges include:

- ☒ **Improved monitoring.** As the saying goes, "You cannot manage what you cannot measure"; surprisingly, however, many IT executives and datacenter managers still do not know how much power their IT equipment consumes. As server power costs have become more significant, datacenter managers need to be able to determine which servers, systems, or group of systems within their datacenters are consuming large proportions of energy.

- ☒ **Power efficiency.** Power efficiency is particularly important as savings are realized both in server power itself and in cooling. The electrical power a system consumes is converted into heat that must be handled by the cooling infrastructure, which requires its own power to operate. Improving the energy efficiency of a server can yield an exponential reduction in the cost of power and cooling.
- ☒ **Effective cooling.** 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**

Available since 2006, HP Thermal Logic is a core piece of HP's Green Business Technology Initiative, a concerted effort to help companies save money and expand their datacenter capacity through more energy-efficient solutions while reducing their environmental impact. Thermal Logic includes a range of innovations in power and cooling and provides a holistic approach to driving efficiency in power and cooling, spanning monitoring, reporting, and adaptive management functionality.

By incorporating Thermal Logic into its BladeSystems, HP is helping its customers realize greater efficiencies in power and cooling in their high-density computing environments. Leveraging innovations ranging from processor to enclosure design and from architecture to management, Thermal Logic allows server infrastructures 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. With Thermal Logic, HP aims to help datacenters address three primary goals:

- ☒ **Reduce** energy consumption
- ☒ **Reclaim** datacenter capacity
- ☒ **Extend** the life of the datacenter

This mantra of reduce, reclaim, and extend represents the heart of HP's Thermal Logic philosophy. Consisting of more than just systems with low-energy components, HP Thermal Logic represents an end-to-end approach to energy management supported by innovations across HP's breadth of offerings, touching on areas as diverse as processors, enclosure design, measurement and management, and datacenter design and optimization.

### ***Reducing Energy Consumption***

For years HP has been incorporating energy-efficient innovations designed to reduce energy consumption in its hardware products. In 2008 alone, HP announced multiple BladeSystem products that offer improved energy efficiency, including the BL260c,

HP's most energy-efficient socket blade to date; the BL2x220c, HP's first two-in-one blade; and the BL495c, a blade designed specifically for virtualization. Two of the most recent examples include the BL460c and the HP 2400W c7000 power supply, both launched in November 2008.

### **Blade Technology Innovations and the BL460c G5**

The BL460c G5 incorporates HP's latest advancements in power efficiency technology to help reduce energy consumption. When designing the BL460c G5, HP went through every component to squeeze out energy efficiency wherever possible. Some of the improvements incorporated into the BL460c G5 include:

- ☒ Use of the latest Intel Xeon quad-core low-watt processor
- ☒ Use of low-power DIMMs
- ☒ Higher efficiency DC-to-DC power conversion modules that supply power to the CPU and memory
- ☒ The ability to automatically power off components such as network ports when they are not in use
- ☒ Incorporation of more power-efficient components across multiple areas of the blade

HP testing indicates that the BL460c G5 saves a minimum of 44W per blade compared with the previous version BL460c G1, representing 704W per enclosure at 100% utilization. For example, with a 60-Amp circuit running 32 460 G1 blades, adding both Dynamic Power Capping and the new 460 G5 blades, HP claims that the circuit could run four full enclosures or 64 blades, a full 100% improvement.

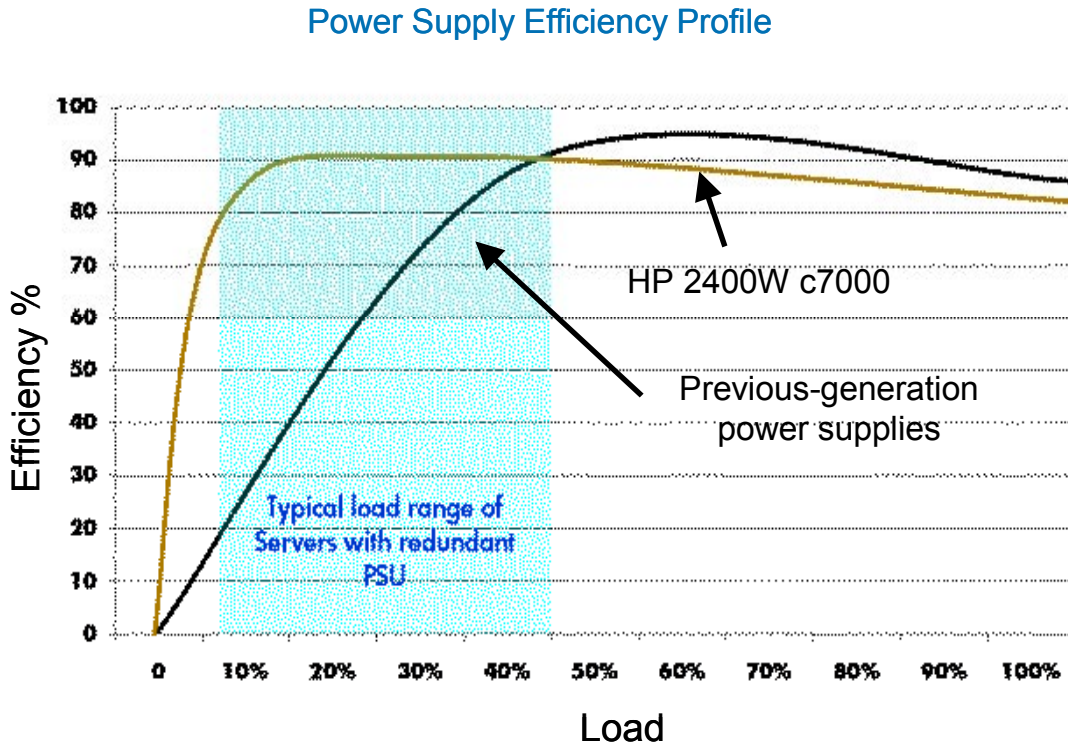
### **HP 2400W c7000 Power Supply**

The HP 2400W c7000 hot plug power supply improves on the energy efficiency of the previous-generation HP 2250W c7000 power supply. The 2400W c7000 power supply provides increased power output, allowing more blades to run on fewer power supply units, has lower standby power for reduced power consumption when servers are idle, and runs at higher peak efficiency than the HP 2250W power supply.

Most importantly, the HP 2400W power supply is designed to run at high efficiency over a greater load range than HP's previous power supplies, and this load range is more representative of the power loads found in most datacenters (see Figure 2). This broader curve of peak efficiency over a wider load range is the key characteristic that enables the 2400W c7000 power supply to deliver efficiency savings compared with previous generations of HP power supply technology.

FIGURE 2

HP 2400W c7000 Power Supply Provides High Efficiency over a Wide Load Range



Source: HP, 2008

This power supply works hand in hand with the HP BladeSystem Dynamic Power Saver for more efficient use of power in the server blade enclosure. Dynamic Power Saver places power supplies in standby mode when the power demand from the server blade enclosure is low, incrementally activating them to deliver the required power as demand increases. Further, it automatically shifts power-supply load to maximize efficiency across the system; for example, by placing a heavier load on fewer power supplies as opposed to allowing a lighter load across the entire set of installed power supplies. The HP 2400W power supply enhances this feature by allowing one-half of the power supply to be placed in standby mode to provide greater granularity, increased efficiency at lower loads, and the ability to more closely match the number of active power units to the actual workload.

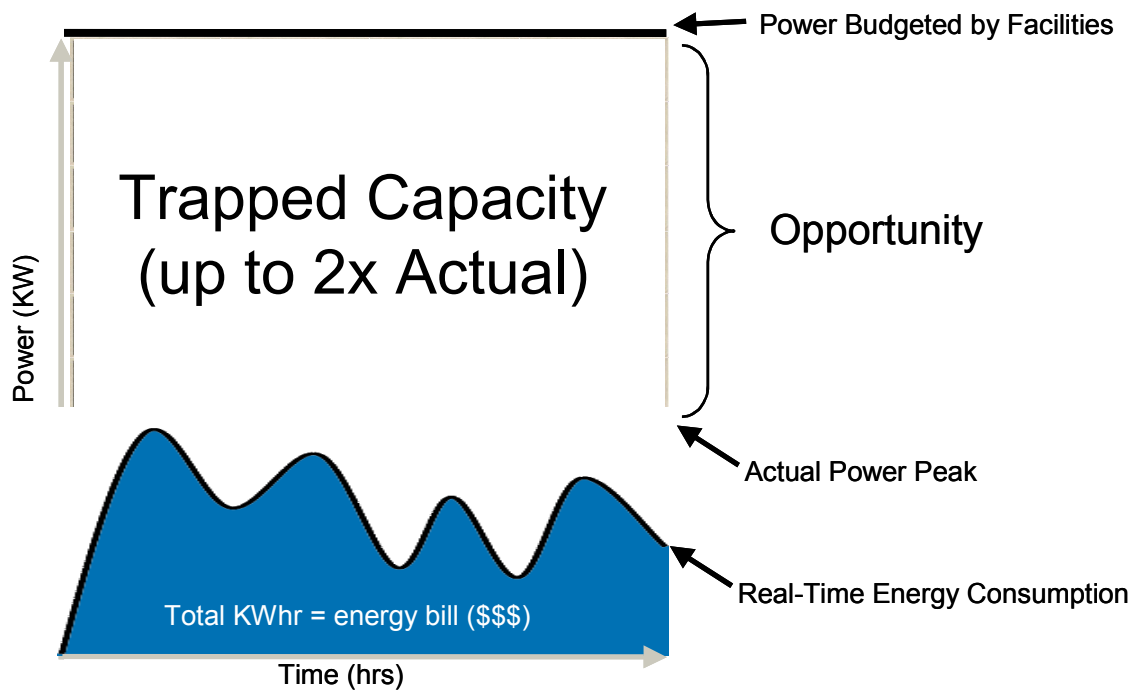
### ***Reclaiming Datacenter Capacity***

Datacenters are built with specific maximum power capacity budgets, and since power and cooling is often the limiting factor in datacenter design, energy budgets — and not floor or rack space — are the limiting factor determining the compute power in any given facility (see Figure 3). Unfortunately, power in the datacenter is rarely efficiently allocated and the full power budget envelope is rarely used. Additionally,

server loads fluctuate day by day, as well as throughout the day, so the amount of power actually used is very rarely equal to the maximum capacity; in fact, in most instances, the actual power peak is well below the maximum power budget available in the facility.

**FIGURE 3**

Reclaiming Trapped Energy Capacity Represents Significant Opportunity



Source: HP, 2008

This unused, or "trapped," capacity represents wasted computing resources and, if it could be reclaimed, could enable datacenters to increase their overall capacity without increasing their footprint, which in turn could lead to significant infrastructure cost savings. By offering technologies to better manage energy efficiency, HP can actually help companies reclaim additional capacity in their datacenter. One key technology is Dynamic Power Capping.

#### **Dynamic Power Capping**

Dynamic Power Capping is a recent Thermal Logic innovation designed to allow companies to control server energy consumption and thereby reclaim unused datacenter capacity. Dynamic Power Capping supports these goals by letting datacenter managers:

- ☒ **Measure peak power usage.** The built-in Onboard Administrator continuously monitors the power usage of each blade and the enclosure, enabling datacenter managers to observe actual (versus face plate estimates of) maximum power usage and determine a lower maximum energy budget that is still sufficient to handle maximum anticipated loads without impacting system performance.
- ☒ **Adjust deployment levels to reclaim capacity.** Based on these adjusted energy budgets, administrators can deploy additional blade servers to reclaim the unused energy capacity.
- ☒ **Set energy caps.** To guard against peak periods that may exceed the new, lower energy budgets, administrators can set maximum energy caps beyond which the servers are forbidden to operate. With the BladeSystem c-Class, these caps can be set at the enclosure level, and if any given blade exceeds its cap, the enclosure redirects power from blades operating well under their caps, enabling the blade to increase its performance and complete the work.

According to HP analysis, Dynamic Power Capping can significantly increase the capacity of a BladeSystem c-Class infrastructure. Referring back to the example of the 60-Amp circuit running 32 460 G1 blades, 3-Phase PDU could run two full enclosures of 32 460c G1 blades, but with the addition of Dynamic Power Capping, the circuit could run 45 465 G1 blades, a 50% capacity improvement.

### ***Extending the Life of the Datacenter***

One of HP's core goals with Thermal Logic is to help companies extend the life of their datacenter. By helping companies optimize their energy use across the IT ecosystem, Thermal Logic can help companies realize greater efficiency with their compute resources. By allowing companies to reclaim trapped capacity, it can help them delay the need to build out additional capacity and extend the life of their current datacenters, saving costs and resources and at the same time reducing environmental impact.

This is supported by a broad range of HP products and services, including innovative technologies designed to help companies better measure and manage their datacenter's energy consumption. Companies can take advantage of these capabilities to manage their datacenter energy profile directly, or they can leverage the expertise of HP Energy Efficiency Services.

### **Energy Measurement and Control Across the Datacenter**

HP has applied the saying "You cannot manage what you cannot measure" to its Thermal Logic offerings. Thermal Logic is designed with a number of management and measurement features that allow IT managers to better control and ultimately reduce their power and cooling needs throughout their datacenter.

Measurement is realized with hundreds of sensors located throughout the blade enclosures to provide views of previously unavailable power and cooling data, all easily available through HP Onboard Administrator and HP System Insight Manager. This data is 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 also allows companies to forecast potential power savings through what-if scenarios when used in conjunction with HP Insight Dynamics – VSE.

## **HP Energy Efficiency Services**

Offered as part of the HP Datacenter Services offering, HP Energy Efficiency Services are intended to help companies better utilize and manage energy, which in turn allows them to maximize capacity, control costs, and extend the life of their datacenter. These services are designed to complement the innovations in energy savings in HP products and are an important part of HP's goal of providing products and services to help datacenters improve their end-to-end energy efficiency and management.

The portfolio of Energy Efficiency Services includes:

- ☒ Facility and Technology Assessment Services
- ☒ Energy Efficiency Design
- ☒ Thermal Assessment Service for Blades

HP offers a range of facility and technology assessment services, including power and cooling analysis, computational fluid dynamics (CFD) analysis, thermal zone mapping, and energy efficiency analysis. CFD modeling enables HP to model the probable effects of potential cooling improvements or the addition of new equipment to the facility. Thermal Zone Mapping is a technology that provides a multicolored 3D graphical thermal map of the datacenter to help IT identify air flows and identify opportunities to adjust air conditioning output settings to improve energy savings and efficiency.

HP energy efficiency design services help companies design datacenters with the cooling infrastructure to increase efficiency for high-density environments. With these services, HP can help companies design to environmentally accredited standards.

Thermal assessment services for customers come in two flavors. The HP Thermal Quick Assessment is an entry-level assessment of the datacenter power and cooling environment that provides a gap analysis and written report designed to highlight best practices and quick wins tailored for the environment. It examines subjects such as cooling load versus capacity, airflow management, heat recirculation, racking practices, and airflow obstructions. The Comprehensive Assessment Service builds on the Thermal Quick Assessment and includes CFD modeling and Thermal Zone Mapping. It includes a more comprehensive set of data collection, including 3D under and above floor modeling, and provides a more comprehensive report and recommendations.

---

## **IDC Analysis**

According to a recent IDC customer survey, power and cooling ranks as the top concern for datacenter operators. Approximately half of the customers surveyed reported server or storage downtime as a direct result of power and cooling issues. While rising energy costs may put the focus on the significant electricity expense associated with running a datacenter, the impact of power and cooling on IT availability causes just as much, if not more, concern among IT managers. In addition to unplanned downtime, power and cooling has emerged as a limiting factor to the

number of servers that can be squeezed into the datacenter. The same IDC study reports that even though customers are experiencing issues with power and cooling, they are utilizing on average only 52% of their datacenter floor space. The point is that datacenters are running out of power long before they are running out of space. The fact of the matter is there is no single solution, or "silver bullet," that can solve the challenge of power and cooling. A holistic approach is required to alleviate the challenges to fully maximize the datacenter's capacity. HP's "chip-to-chiller" approach involves a set of solutions to optimize server power and cooling to deliver the IT benefits of reducing overall TCO, improving availability, and expanding capacity.

Thermal Logic technologies within HP BladeSystem enable companies to realize greater synergies with energy efficiency, increased capacity, and reduced TCO. 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 Capping feature provides the BladeSystem enclosure with the capability of dynamically adjusting power levels across the blades to optimize energy consumption. This feature enables IT to improve energy efficiency without impacting performance levels. For all the innovation and benefits associated with Thermal Logic, however, barriers also exist. They include issues in power and cooling solution standards and the necessity to compete against facilities vendors with their own long-standing relationships with datacenter managers.

### ***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. 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 overprovisioning 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 is a priority for IT executives, it represents 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 Compete Against Facilities Vendors***

In the power and cooling arena, HP faces nontraditional vendors such as facilities vendors Liebert and 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 with both traditional and nontraditional competitors.

---

## **Conclusion**

HP Thermal Logic represents a step forward in thinking about and managing the energy consumption of today's datacenter. It extends beyond simply squeezing additional energy efficiency out of individual servers or managing to specific benchmarks, instead focusing on increasing the efficiency of the complete IT infrastructure, across all workloads, with an eye toward reducing the total amount of IT energy use.

A core technology available in HP BladeSystem offerings, Thermal Logic allows IT managers to measure, control, and optimize their use of power and cooling resources to match variable degrees of workload and utilization. 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 companies to reduce energy consumption in their IT infrastructure, reclaim unused datacenter capacity, and extend the life of the datacenter. Thus, it can help companies better scale compute resources to meet their changing needs, and by optimizing the power and cooling required per server, it can help companies reduce overall TCO.

---

## **Copyright Notice**

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2008 IDC. Reproduction without written permission is completely forbidden.