IDC
Analyze the Future

TECHNICAL BRIEF

Improving Power and Cooling Efficiency in the Datacenter

Sponsored by: HP

Jed Scaramella

June 2010

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

Introduction

HP has been a pioneer in incorporating technology to improve energy efficiency and reduce total cost of ownership (TCO) for IT systems in the datacenter, with solutions designed to reduce server management costs, increase utilization, and improve power and cooling.

With the launch of its new Data Center Smart Grid in June 2010, HP is continuing to push the boundaries of what its technology can deliver. A central aspect of the Data Center Smart Grid is 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. Built into HP's ProLiant and BladeSystem offerings, Thermal Logic also plays a key role in differentiating HP offerings from competitive systems.

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.

IDC research reveals that datacenter power and cooling is a top-of-mind issue for IT and facilities executives. IDC estimates that the energy expense associated with powering and cooling the worldwide server installed base increased 31.2% over the past five years. During this same period, server energy expense has outpaced spending on servers themselves in terms of both revenue and units; server energy expense grew at a compound annual growth rate (CAGR) of 9.5% compared with a server revenue CAGR of -2.6% and a unit shipment CAGR of 1.0%.

Figure 1 displays the energy expense associated with powering and cooling the worldwide server installed base by type of datacenter for 2010. As Figure 1 shows, almost 70% of the server energy expense is generated by datacenters and the remaining 30% by smaller IT sites. Enterprise-class datacenters, which are the largest datacenters yet the fewest in number, account for over 40% of energy expense on their own. The energy efficiency of datacenters is gaining significant attention due to the concentrated energy consumption by a relatively small number of facilities.

FIGURE 1

Worldwide Server Power Consumption by Datacenter Size

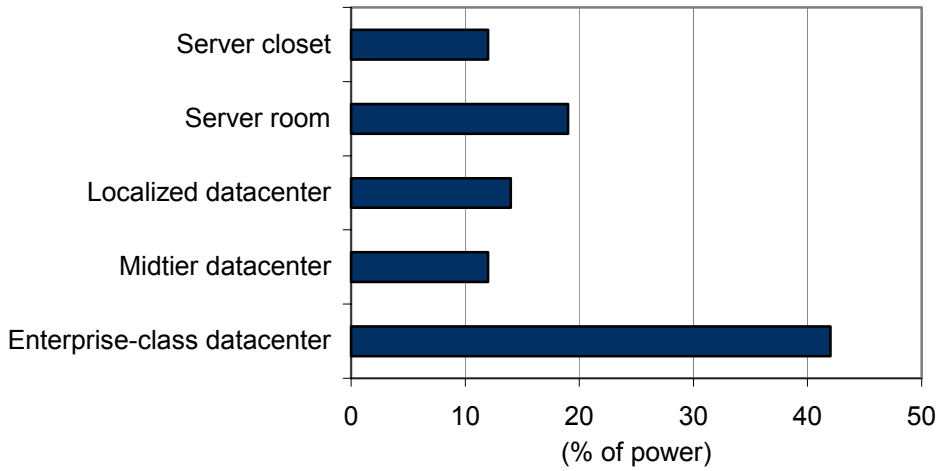

Source: IDC, 2010

IT Availability Is the Primary Driver of Energy Efficiency

Most of the conversation surrounding the energy expense of the datacenter is focused on the direct financial impact to companies; however, IDC finds that the primary driver behind companies' efforts to improve datacenter energy efficiency is to ensure the availability of IT to the business. Over the past decade, IT has become a strategic asset to businesses, and IT performance is very often tied to business results. Even though companies feel the direct financial impact of the escalating electricity bills, IT availability can have a much greater impact on business operations. While IT organizations' stated initiatives may include improving energy efficiency, reducing or controlling operating expense, or even looking to go "green," IT availability overrides all other factors and will not be compromised to achieve other efficiencies.

The Challenge: Power and Cooling 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 because 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 expediential 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

HP has continually enhanced its Thermal Logic technology since its launch in 2006 and has recently incorporated it as a key component of its new Data Center Smart Grid initiative to provide datacenterwide intelligent energy management across systems and facilities. Working at the server and rack level, 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 BladeSystem and ProLiant offerings, HP is helping its customers realize greater efficiencies in power and cooling in their computing environments. Leveraging innovations that range 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. Recent innovations include power-efficiency improvements built into its BladeSystem and ProLiant servers, as well as the release of its new HP Common Slot Platinum Power Supplies.

The latest HP ProLiant G7 servers are designed to provide more performance and better energy efficiency than previous generations of ProLiant servers. The G7 line includes rack-optimized servers, blade servers, and freestanding tower systems. Specific power-saving technologies built into the G7 line include:

- Dynamic Power Capping
- Sea of Sensors
- Common Slot Power Supplies
- Insight Control power management

Each of these technologies is described in greater detail in the following sections.

HP Common Slot Platinum Power Supply

The power supply is naturally a critical component to HP's end-to-end approach to datacenter energy efficiency and management. That's why HP developed the new Common Slot Platinum Power Supply for its ProLiant and Integrity servers, incorporating a common electrical and physical design, a flexible range of power outputs, and its greatest efficiency rating yet at 94%.

Improvements in the Common Slot Platinum Power Supply over previous generations of HP power supplies include:

- **94% efficiency rating.** HP has continuously driven improved efficiency ratings into its power supplies, with previous generations rated at 90% and, more recently, at 92%. Rated at 94% by 80 Plus and the Electric Power Research Institute (EPRI), the Common Slot Platinum Power Supply has the highest energy-efficiency rating of any power supply in HP's history.
- **Common electrical and physical design.** The common design allows the power supply to be installed in many ProLiant G6 or G7 servers, as well as the BladeSystem c3000, that have a common slot power bay.
- **Flexible range of outputs.** This power supply comes with a choice of 460W, 750W, and 1200W outputs, enabling customers to "rightsize" the power option that is most appropriate for their own particular server configuration.

- **Configurable Load Balancing and High Efficiency modes.** When redundant power supplies are used, in Load Balancing mode, the power load is shared equally between those power supplies. In contrast, in High Efficiency mode, the entire server load is shifted to the designated primary power supply, which then operates at a higher load and therefore greater efficiency level (with the other power supplies placed in idle mode). This mode provides energy savings when servers are operating at loads that require less than 70% of their primary power supply capacity.

HP 2400W Platinum Power Supply

The HP 2400W Platinum Power Supply is a hot-plug power supply that incorporates innovations in energy efficiency for HP c7000 BladeSystem offerings. It 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 previous-generation 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. This broader curve of peak efficiency over a wider load range is the key characteristic that enables the 2400W Platinum Power Supply to deliver efficiency savings compared with previous generations of HP power supply technology.

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 because 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 2). 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 2

Reclaiming Trapped Energy Capacity Represents Significant Opportunity

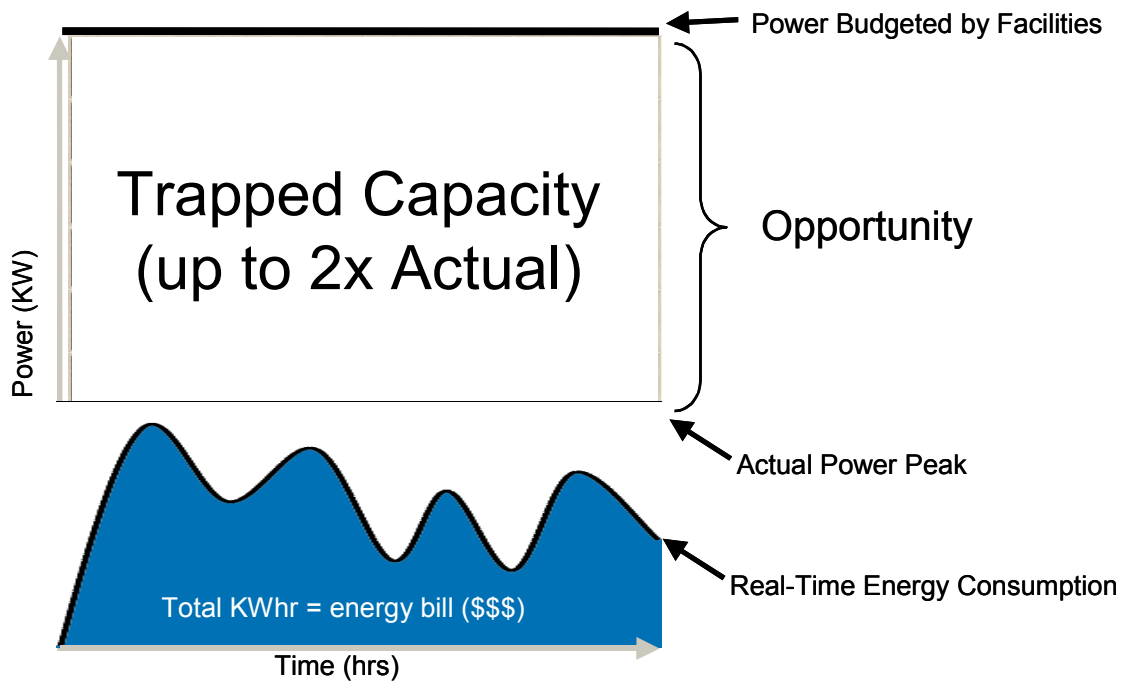

Source: HP, 2010

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 including Sea of Sensors, Intelligent Power Discovery, and Dynamic Power Capping, HP can actually help companies reclaim additional capacity in their datacenter.

Sea of Sensors

Most HP BladeSystem enclosures and ProLiant G6 and G7 servers feature what HP refers to as a "sea of sensors" to monitor each server's internal environment. These sensors track power requirements and manage power delivery to major components in real time. Information can be gathered for a group of servers over time, managed by HP Insight Control power management, and integrated into the rest of the HP Data Center Smart Grid offerings. It can help IT administrators manage their systems' energy use by setting power caps for HP ProLiant servers, as well as HP BladeSystem enclosures.

Intelligent Power Discovery

Intelligent Power Discovery helps to create an energy-aware network that shares data between IT systems and facility management systems. It takes advantage of data from the "sea of sensors" embedded in HP servers and feeds the data into both Insight Control power management software and (integrations with) third-party facility management offerings from Eaton and nlyte Software.

Intelligent Power Discovery combines HP Intelligent Power Distribution Unit (PDU), HP Platinum Power Supplies, and HP Insight Control software to automatically track new server installation and provide higher precision, control, and automation to power distribution. When new HP ProLiant servers are deployed, HP Intelligent Power Discovery automatically discovers those servers, maps them to the power source, verifies power redundancy, and makes sure everything is hooked up correctly.

Using power line communication (PLC) technology, HP ProLiant G6 and G7 servers now provide identification information such as server name, UUID number, and IP address to the Intelligent PDU and to HP Insight Control power management software. This reduces the amount of time need to configure the power distribution software and hardware. Additionally, HP Intelligent Power Discovery intuitively detects redundant power supplies and makes sure they are running on different PDUs. This helps reduce unscheduled downtime and the possibility of human error, the biggest challenges of datacenter management.

Insight Control Software

HP Insight Control is a component of the Data Center Smart Grid offering, which allows customers to continuously control and manage their HP infrastructure across ProLiant and BladeSystem. It provides server deployment, health and performance monitoring for both physical and virtual resources, remote control with iLO Advanced, and power management. Insight Control also integrates into enterprise management tools from HP Software, Microsoft, and VMware.

Another new feature within HP Insight Control power management software is Data Center Power Control. This feature offers a number of benefits to the power and cooling equation, such as the ability for businesses to keep critical applications up and running while gracefully shutting down noncritical workloads. Insight Control power management also offers power management at the rack level, continuous power readings, server power cap enforcement, and association of power requirements to workload demands.

Dynamic Power Capping

Dynamic Power Capping is 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. For example, within a 60-Amp circuit running 32 460 G1 blades, a 3-Phase PDU could run two full enclosures of 32 460c G1 blades. Now, with the addition of Dynamic Power Capping, the same 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 taken to heart the importance of measurement as a prerequisite to power management and applied that philosophy 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

IDC considers energy to be the key design point for the datacenter — both now and into the future. As noted above, energy is having a greater impact on businesses' financial performance and has become the primary limiting factor for expansion of IT capacity. The availability of sufficient power and adequate cooling determines where systems can be deployed within the datacenter and even within the rack. A strategy to improve energy efficiency is a primary factor to effectively operating a datacenter by reducing risk from failure due to hotspots and power overloads, as well as making full use of untapped power capacity.

The dynamics of operating a datacenter are evolving. IT organizations cannot focus solely on system performance and capital expenditure without considering the impact on IT availability and operation expenditures. Customers report to IDC that power and cooling remains their number 1 datacenter challenge; this implies that customers are placing greater emphasis on energy characteristics when considering purchasing criteria for IT systems and are placing greater emphasis on the efficiency, control, and power management aspects provided by datacenter equipment and management.

As part of HP's Data Center Smart Grid, Thermal Logic enables companies to realize greater synergies with energy efficiency, increased capacity, and reduced TCO. Thermal Logic's core capabilities include pooled power distribution and Dynamic Power Capping that increase power efficiency and optimize energy consumption by dynamically adjusting power levels across ProLiant servers and blades. HP has extended the reach of these solutions beyond its original blade focus to the entire datacenter. HP's Intelligent Power Discovery can help IT and facilities integrate datacenter operations and create a unified approach to datacenter energy management. The new capabilities will benefit customers by curbing the risk of unplanned downtime related to human error in power allocation and reducing the frequent underprovisioning or overprovisioning of power that occurs in datacenters.

Enabler for the Converged Infrastructure

Under the umbrella of HP Data Center Smart Grid, Thermal Logic is a key tenet of HP's Converged Infrastructure. Thermal Logic allows 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. Through its new software and partnerships with facilities vendors, HP is offering a more complete and integrated datacenter energy management solution by enabling IT systems and facilities to communicate automatically in order to optimize energy consumption to IT workload demand.

Conclusion

HP Thermal Logic represents a step forward in thinking about and managing the energy consumption of today's datacenter. It plays a key role in the HP Data Center Smart Grid initiative to deliver an end-to-end portfolio of energy-efficient technologies reaching from the server to the entire datacenter. Its holistic approach extends beyond simply squeezing additional energy efficiency out of individual servers or managing to specific benchmarks to include instead focusing on increasing the efficiency of the complete IT infrastructure, with an eye toward reducing the total amount of IT energy use.

A core technology available in HP ProLiant and 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.