HP power and cooling technologies for the data center

technology brief

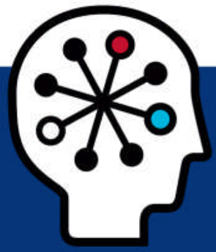

Abstract

The power demands of IT equipment—including servers, storage, and networking—have grown geometrically in the last decade. They are now straining the limits of the power and cooling infrastructure of many data centers. This trend, combined with the increasing cost of power itself, has put data center power and cooling at the forefront of IT issues. HP has been deeply involved in creating innovative products and solutions to address all aspects of this challenge – including improving power distribution efficiency and provisioning, lowering power consumption wherever possible, and monitoring and maximizing cooling efficiency in the data center. This technology brief provides a summary of the present power and cooling challenges in the data center as well as an overview of the products and solutions that HP has introduced to help address these important issues.

Emergence of the power and cooling problem

Rising power consumption

Over the past several years, server performance has increased dramatically. Along with this increase in performance has also come an increase in server power consumption. Although server performance per watt has been steadily increasing over this timeframe, absolute server power consumption has been increasing as well. According to some industry studies, in the recent past power consumption per server has been doubling approximately every three years. Power consumption per rack, which used to be budgeted at 5 kilowatts, is now routinely exceeding 15 kilowatts. As companies have continued to require greater levels of IT across their business functions, power density in the data center has also increased.

This increase in power density and consumption has led to a power and cooling crisis. Using conventional models and technologies, many administrators are running out of power and cooling capacity well before they run out of available space. Even when this is not the case, reducing power consumption in the data center has become an important part of Green initiatives at companies faced with rising energy costs. HP has invested significantly in technologies to ensure that ProLiant G6 servers use less power and deliver better performance than their G5 counterparts. However, the longer-term trend of increasing server power consumption is not something that is likely to reverse itself. Solving this challenge requires innovative solutions to effectively distribute provision, consume and monitor both power and cooling in the data center.

Cap Ex and Op Ex exceed IT costs

Along with the increase in server and IT power consumption, a related issue has also emerged—the dramatic increases in capital expenditures (cap ex) and operational expenditures (op ex) for the modern data center. Increases in operational costs are being driven primarily by energy costs while rising infrastructure costs (cap ex) are primarily a function of the complexities of building out data centers to handle steadily increasing power densities. Some analysts have calculated that capital expenditures in Tier IV data centers are now projected to approach $25k per kilowatt of IT power capacity.

The potential result of these trends can be seen in Figure 1. Without significant and ongoing efforts on the part of the IT industry, the combined op ex and cap ex for the data center were projected to be more than six times the cost of the IT equipment by the year 2012. This makes solving the power and cooling challenges in the data center even more imperative, and it is something that HP is committed to doing with innovative products and solutions across the product line.

Figure 1. Projected trend of op ex and cap ex versus IT equipment costs

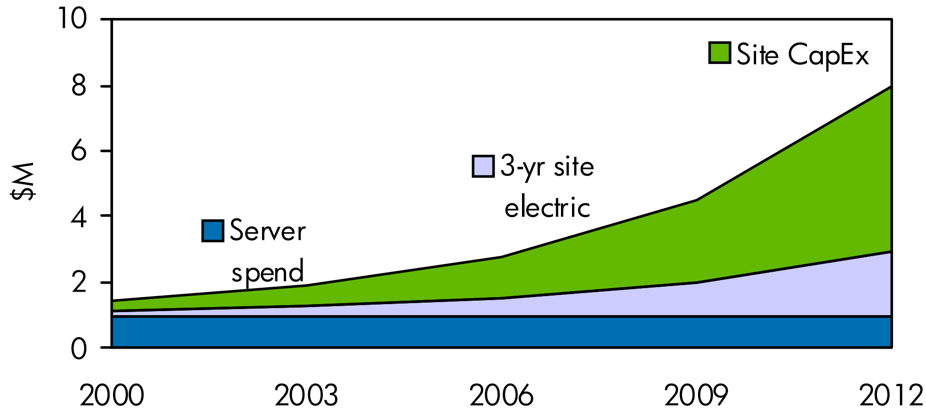

Source: Kenneth Brill,"The Invisible Crisis in the Data Center: The Economic Meltdown of Moore's Law," TUI3008A, Uptime Institute, 2007

Understanding the power life cycle

It is important to first understand where power is used in the data center. The power life-cycle consists of three distinct parts, each of which consumes power.

- Power distribution. Power distribution begins where the power is taken from the utility grid and extends to the point where it is delivered to the IT equipment for consumption. No distribution system is 100% efficient, and as a result some power is actually consumed during this process.
- IT power consumption. This is the power consumed by the IT equipment itself (servers, routers, disk storage systems, and so forth).
- Heat extraction. The power consumed by the IT equipment is converted almost entirely to heat. Additional power is needed to extract this heat and keep the data center temperature within the operational envelope of the IT equipment.

Data center power usage effectiveness

Each of the three parts of the power life cycle consumes power in the data center. Ideally, we would want all power to be used by the IT equipment, since power lost in distribution or to extract heat is power that could instead be used to support more compute resources. This ideal goal is reflected in the commonly used data center metric of power usage effectiveness, or PUE. In simplest terms, the PUE of a data center is the ratio of the total power consumed by the data center to the power consumed by the IT equipment alone. A perfect PUE would be 1.0, representing a data center in which all the power consumed was used by IT. Most data centers have PUEs between 1.5 and 3.0.

Figure 2 shows where power is consumed in a range of data centers. In the median data center in this study, the IT load uses 59% of the power consumed. However, cooling—or heat extraction—still uses over one-third of the power in the data center. Power delivery, along with other miscellaneous loads in the data center, consumes only 7% of the power. This includes the losses from conversion and transmission as power is delivered from the utility power grid through the UPS and to the servers in the racks.

Figure 2. Power consumption in the data center

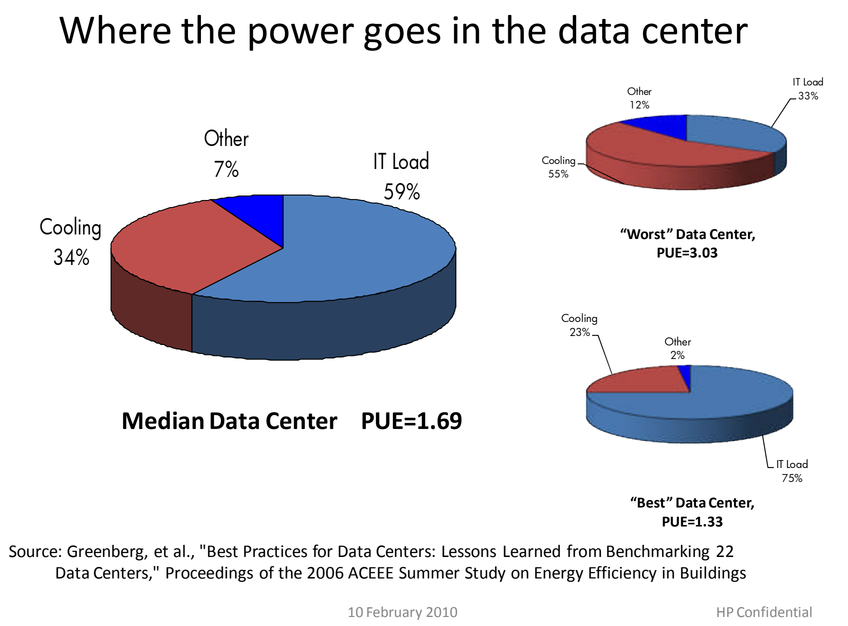

Data center infrastructure efficiency (DCiE) is another metric used to represent datacenter efficiency, and it is essentially the inverse of PUE expressed in percent. A data center with a PUE of 1.5 has a DCiE of 1/1.5 x 100, or 66 percent. For more information on the PUE and DCiE metrics, visit http://www.thegreengrid.org.

Addressing the power and cooling challenge

The key to addressing the power and cooling challenge is to understand how to improve power consumption as well as allocate power resources more efficiently. HP has identified three areas of innovation for solving the power and cooling challenge:

- Increased efficiency. Using innovative technologies to accomplish the same tasks using less energy.
- Managed provisioning. Creating solutions that ensure that the available power and cooling resources for the data center are fully utilized. Over-allocated resources waste capital investment.
- Monitoring and control. Measuring and reporting power and/or cooling metrics for HP products involved in every phase of the power life cycle. This data provides the information needed to increase efficiency and to implement effective provisioning solutions.

Solving the datacenter power and cooling challenge involves applying innovations in these areas to each part of the power life-cycle. Working within this framework, HP has introduced products and solutions that can decrease power consumption and allow administrators to dramatically increase the utilization of existing infrastructure.

Improving power distribution efficiency

As Figure 2 shows, between 2% and 12% of the power delivered to the datacenter is lost in transmission and distribution between the grid and the servers. Although this is not a large number, it represents a pure loss for which no meaningful work is accomplished. This lost power is converted into unwanted heat which requires additional power to extract.

Basic principles of efficient power distribution

Power is delivered to the datacenter from utility power grid as 480 volt 3-phase AC in North America and 380 volt to 415 volt 3-phase AC in the rest of the world. As it is distributed to the racks and then to the servers, it may be converted, or stepped-down, several times depending on the data center distribution infrastructure. 480 volt 3-phase AC can be converted to 208 volt 3-phase which is distributed to the servers as separate 208 volt or 120 volt single-phase circuits.

Each conversion point adds to infrastructure costs. Because no conversion is 100% efficient, each conversion also results in a power loss. A more efficient power distribution infrastructure can be created by observing some basic principles of power distribution.

- Keep power conversions to a minimum. Each conversion costs power and creates heat.
- Keep the voltage as high as possible as long as possible. Higher voltage circuits experience lower transmission losses than lower voltage circuits and also use less expensive wiring infrastructure.
- Minimize the total transmission distance.

With these principles in mind, HP engineering has been able to design and deliver power distribution products that improve the overall efficiency of the power distribution chain.

Improving the efficiency of the AC power distribution chain

Although other distribution schemes exist or are being considered, AC power distribution is by far the most prevalent in today's data center. Figure 3 illustrates a typical AC power distribution chain using standard UPS units and redundant AC power supplies in the servers. Including the AC to DC conversion inside the server power supply, this chain loses over 25% of its input power during the distribution process, delivering 72- 74% of the power to the ultimate consumer, the IT load. Although newer data centers may do better than this, there is still room for improvement. The HP strategy is to increase the conversion efficiency in this chain wherever possible.

Figure 3. Present AC distribution

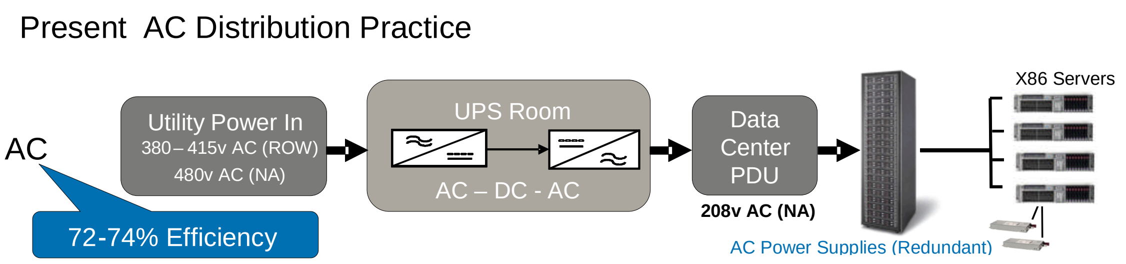

HP Uninterruptable Power Supplies

Traditional UPS designs use online double conversion and have often been the single least efficient component in the power distribution chain. In this design, incoming AC is converted into DC and then converted back to AC before it is distributed to the rack and the servers. Online double conversion provides the most pristine AC power downstream; however, power is lost during each of the two conversions. This results in a UPS efficiency that is typically in the 85% - 92% percent range.

The HP eco-mode UPS uses advanced line-interactive technology which monitors AC quality and passes it directly to the IT load as long as it is within specifications. This eliminates the two conversions and significantly increases overall power efficiency. The UPS will switch to double conversion only if the incoming power is out of specification. This prevents possible damage to the IT equipment downstream. This approach gives the HP eco-mode UPS an overall power efficiency of up to 97%. These efficiency levels are driving eco-mode UPS adoption across the industry, even for the larger systems.

HP monitored power distribution

The HP Power Distribution Rack (PDR) is designed to improve power management in the data center by moving power distribution to the row level. The HP PDR (Figure 4) uses dual 400 amp 208 volt 3-phase input feeds and distributes power to an entire row of racks through individual power cords, each protected with a single phase or 3-phase circuit breaker. This provides several benefits:

- Allows the 480 volt to 208 volt transformers to be off the data center floor, reducing cooling requirements
- Reduces the length of the power feeds needed from the remote power panel to the individual IT racks.

Figure 4. HP Power Distribution Rack and Monitored PDU's

Both the HP PDR and HP monitored PDUs support remote monitoring over the network. These capabilities allow administrators to monitor power consumption and load balancing at the row, rack, and branch circuit level. Remote power management can also be integrated with HP Systems Insight Manager, which supports PDR and PDU discovery as well as receipt of alert traps. This integrated power monitoring and measurement delivers the information that administrators need to help control their power distribution in the data center.

Improving power supply efficiency

Power supplies, which convert AC into the 12 volt DC that is distributed inside the server, are not the last stop in the power distribution chain. Within the server, this 12 volt DC is stepped down to the different voltages used by individual components, including processors, memory, the PCIe bus, and storage devices. Some calculations consider them to be the last step in the power distribution chain, while others (including server power benchmarks) treat them as the first element in server power consumption. Different examples in this paper include them in different parts of the power life cycle.

HP has continuously improved server power supply efficiency. The power supplies for the ProLiant G5 servers were the first intelligent power supplies, with their microprocessors and system management interface. Overall, these power supplies had a maximum efficiency of 89% (measured at 50% load). The power supplies used in the ProLiant G6 servers use newer designs and improved components to deliver a maximum efficiency of 92%.

Optimized AC Power Distribution

Using the HP technologies and products discussed, data center administrators can implement an AC power distribution chain that is 84 to 86% efficient, as shown in Figure 5. This improves power distribution efficiency ten percentage points over older distribution infrastructure shown in Figure 3.

Figure 5. Optimized AC Distribution

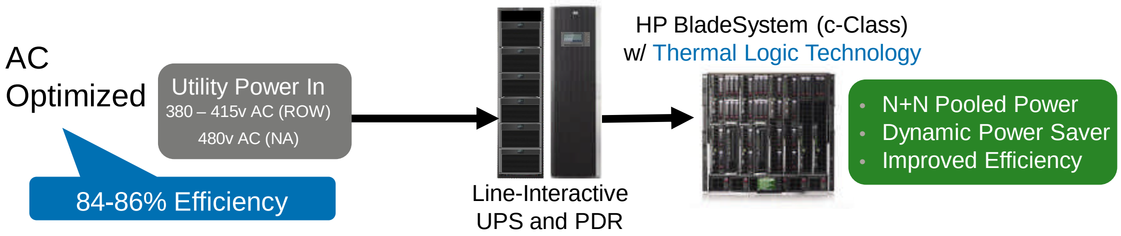

DC Power distribution

The industry continues to explore the possibilities of DC power distribution in the data center. With DC distribution, utility AC power into the data center is converted to High Voltage DC power (HVDC) which is then distributed in the data center until it is consumed by the servers and other IT equipment using DC-to-DC power supplies.

Proposed DC distribution models could potentially be 2 – 4 percentage points more efficient than optimized AC distribution. However, the following challenges still have to be addressed:

- DC power distribution will require replacing large parts of existing data center infrastructures.
- The industry needs to agree on standards for DC power.
- HVDC power has safety issues that need to be addressed.
- The return on investment (ROI) of DC distribution may not justify its adoption in many instances.

HP continues to monitor developments in DC power distribution and remains committed to providing the most cost-effective and efficient IT technologies available regardless of the power distribution method used.

Maximizing data center power utilization through provisioning

Allocating available power to the racks of IT equipment is one of the most important planning tasks in the data center. Typically, a data center has a fixed amount of power available from the utility feed. Once all of this power has been allocated to the different racks through the available circuits and PDU's, additional IT equipment cannot be effectively added to the data center without significant capital expenditures to increase the data center power capacity. The recent increase in data center power consumption has become one of the largest problems and most significant costs that many data centers now confront.

Trapped power capacity

With traditional power distribution planning, administrators provision enough power to a server to meet the maximum input of the server power supply (based on faceplate rating). While this ensures that all servers have sufficient power available at all times, it also uses up the data center power budget much faster than is necessary. Power supplies are commonly designed to fit a wide range of applications. Most real server configurations operating at 100% utilization never use more than 70% of their power supply's maximum output. If all servers are provisioned to their power supply maximum, the data center can potentially have all of its power capacity allocated; however, it may never draw significantly more than 70% of that the supplied power. This over-allocated power is referred to as trapped capacity.

This situation is made worse by the fact that most servers rarely operate at 100% utilization for any significant period of time. A server that uses 500 watts at 100% utilization may consume only 300 watts most of the time. This means that there is still more trapped power capacity that can be potentially reclaimed.

HP has developed server technologies and solutions that allow administrators to use this information to reclaim trapped capacity and more efficiently provision power in the data center. These include

- HP Power Advisor
- ILO and Insight Power Manager
- HP Dynamic Power Capping

HP Power Advisor

The HP Power Advisor is a standalone utility designed to calculate the power requirements for ProLiant servers and rack systems. Using the HP Power Advisor, an administrator can determine the power requirements for a particular configuration for a single server or for an entire rack of servers. These calculations use formulas that are based on data collected through extensive testing of HP ProLiant servers. It can be adjusted to determine server power requirements at different utilization levels. Figure 6 shows the calculated input power requirements for an HP ProLiant ML/DL 370 G5 at 100% utilization.

Figure 6. Power Advisor calculation for an HP ProLiant DL 370 G6 server

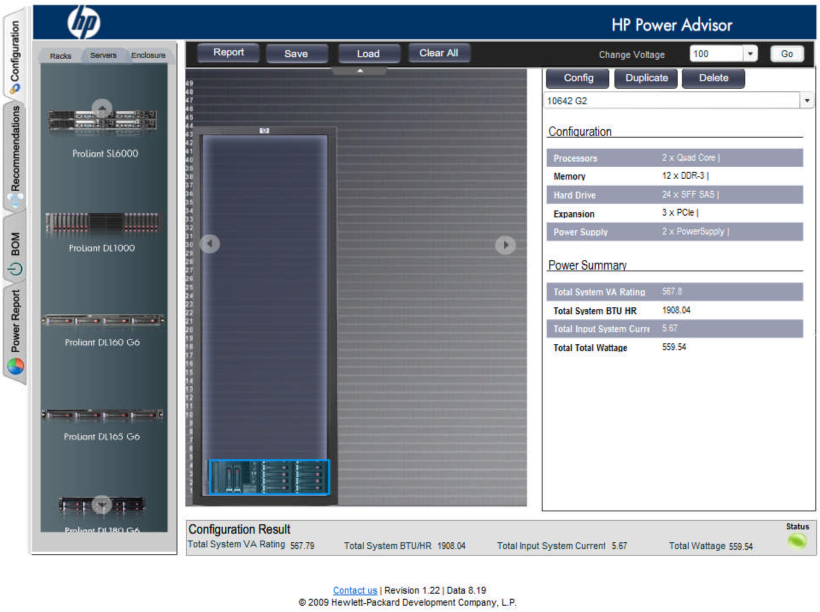

Monitoring server power use with iLO and Insight Power Manager

Several HP ProLiant G5 and later servers have been designed to monitor and report their own power consumption. Individual server power use is reported to the server's iLO management processor, where it can be viewed through the server management interface.

HP Insight Power Manager (IPM) provides another level of power monitoring and reporting functionality beyond that of iLO. Using IPM, which integrates with HP Systems Insight Manager, an administrator can monitor the power consumption of entire groups of servers and can maintain years of historical records of peak and average power consumption (Figure 7).

Figure 7. Monitoring power consumption of a group of servers with Insight Power Manager (IPM)

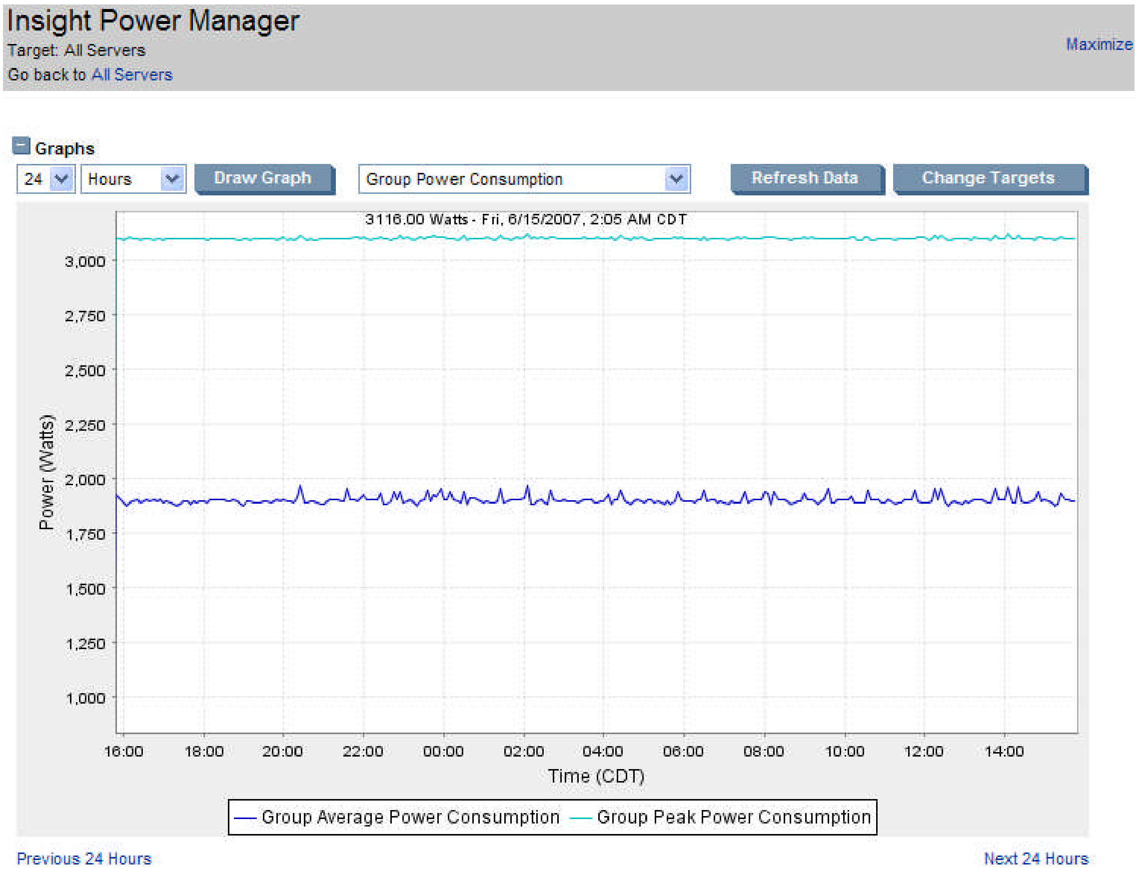

HP Dynamic Power Capping

HP Dynamic Power Capping technology is incorporated into ProLiant G6 and selected G5 servers. It uses sophisticated monitoring and control circuitry to safely limit maximum server power consumption to a preset level. Because it is hardware-based, Dynamic Power Capping can control server power consumption quickly enough to ensure that sudden surges in power demand from the servers are quickly managed and will not trip circuit breakers in the HP PDUs. Administrators can use the power consumption information gathered with IPM to set appropriate power caps to individual servers through iLO or to groups of servers through IPM. Setting power caps to the peak observed power consumption of the server(s) places a defined limit on power consumption without impacting performance. More importantly, it makes it possible to free up additional power that can potentially be reallocated to additional servers.

Using power capping to reclaim trapped power capacity

A facilities manager or server administrator can use the information provided by HP Power Advisor to reclaim trapped power capacity in the data center. Figure 8 shows the different provisioning possibilities for a rack of twenty ProLiant DL380 G6 servers equipped with 1200 watt common slot power supplies. Traditional power provisioning calls for provisioning this rack with 24000 watts of power. The HP Power Advisor shows that the maximum this rack of servers would consume at full utilization is 11000 watts. A manager can safely provision the rack with 11,000 watts of power, knowing that power consumption will not exceed this amount. The HP Power Advisor becomes the first step toward the goal of completely predictable and manageable power consumption.

Figure 8. Reclaiming trapped power capacity through managed provisioning of power to the rack

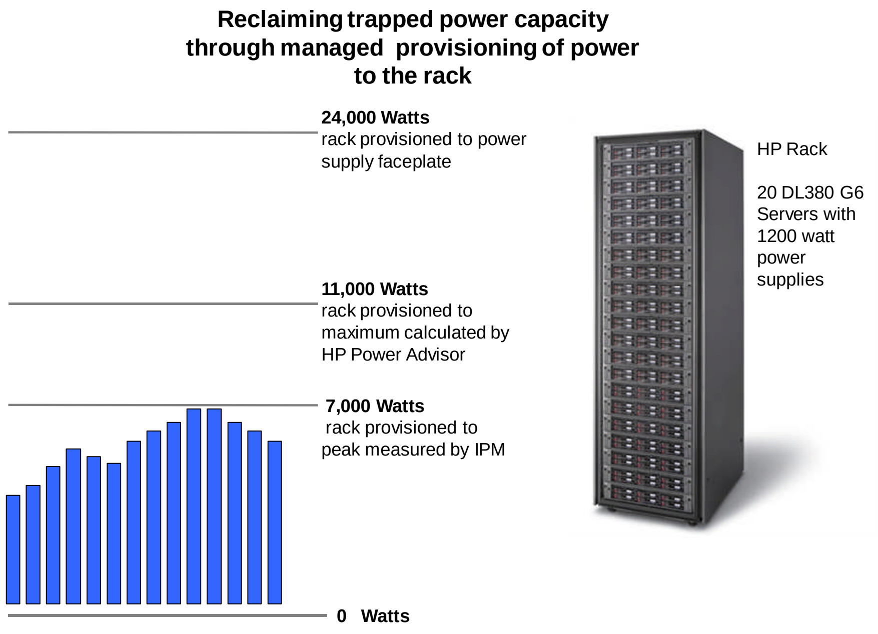

Using HP Dynamic Power Capping in conjunction with data gathered with HP Insight Power Manager, an administrator may be able to reduce power provisioning to the rack still further and potentially allocate it elsewhere in the data center. Figure 8 shows that the IPM measured the peak power consumption for this rack of servers at just below 7000 watts. Knowing this information, an administrator can use Dynamic Power Capping to cap the total power consumption at 7000 watts. This ensures that the rack of servers will never exceed this consumption level. Assuming relatively uniform workloads across all of the servers, the 7000 watt power cap should not adversely affect overall performance since it is slightly above the measured peak of the total rack power consumption. Under these conditions, the rack could safely be provisioned with 7000 watts of power, less than one third of that needed when provisioning to the power supply faceplate. HP Dynamic Power Capping provides completely predictable power consumption.

Using modular capacity to control data center capital costs

Because capital costs have become one of the largest parts of the IT budget, managing and controlling them has become increasingly important. As part of the overall provisioning solution, HP is delivering products that allow facilities managers to add power and cooling infrastructure in an incremental, or modular, fashion.

The HP Rack-mountable UPS is a good example of this. Available in sizes from 2U at 2880 VA to 6U at 12000 VA, they are designed to provide industry leading power density in terms of watts per-U-space. UPS capacity can be added incrementally to the data center as required.

The HP Modular Cooling System (MCS) G2, which is described in detail later, is another example of this modular infrastructure concept. The MCS can be used to provide additional data center cooling capacity on a rack-by-rack basis. This allows a manager to incrementally increase cooling capacity in those areas that require it without having to re-address the overall data center cooling infrastructure.

Increasing power efficiency with HP ProLiant G6 servers

In a well-designed data center, servers, storage, and networking devices used to perform the work that is the data center's output are the largest consumers of power. Any improvements in server power consumption will therefore pay large dividends. Lower server power consumption generates less heat and results in lower losses for power distribution, which lowers the cost of heat extraction at the end of the power life cycle.

Common slot power supplies

ProLiant G6 servers use the HP common slot power supplies, which feature a single mechanical form factor regardless of their power capacity. This allows them to be used interchangeably across multiple platforms and support hot plug redundancy. In addition to being more efficient at all power loads than previous generations of HP power supplies, common slot power supplies provide the opportunity to right-size the power supply capacity to the anticipated power consumption of a server.

Power supplies are typically most efficient when operating at 50% of their maximum output capacity. HP produces three common slot power supplies with maximum output power of 460 watts, 750 watts, and 1200 watts. Figure 9 shows the efficiency curves for these power supplies. Common slot power supplies allow planners to select a power supply that will operate close to its maximum efficiency for the planned server power load. A 750 watt power supply would be the optimal choice for a server that has an average power load of 350 watts, since it would be 92% efficient at that load. A 1200 watt power supply installed in that server would only operate at 88% efficiency.

Figure 9. Efficiency curves for common slot power supplies

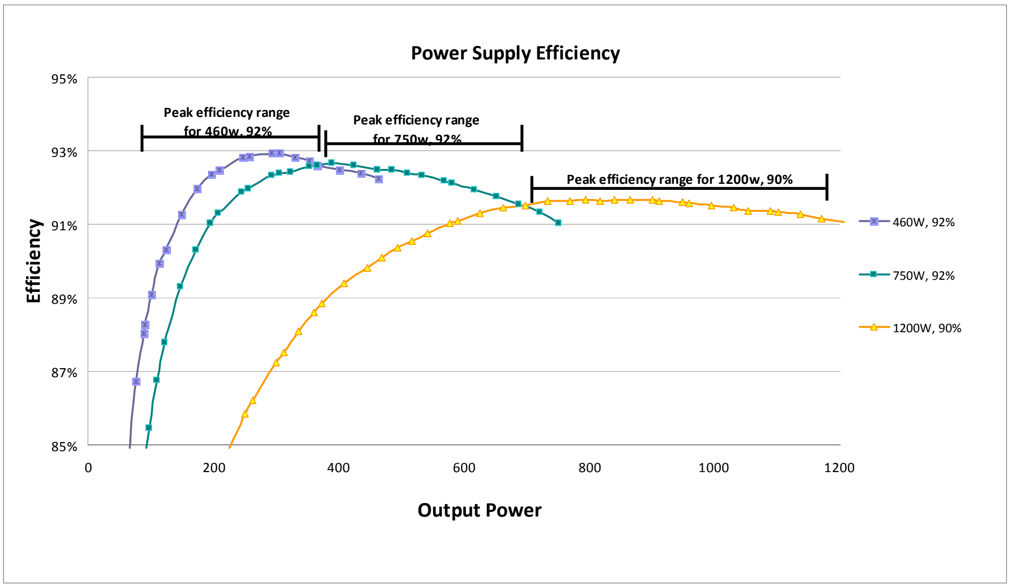

Increasing the efficiency of redundant power supply operation

Using redundant power supplies in servers delivers increased reliability; however, in the past it could result in decreased power efficiency. In G5 servers, both redundant power supplies were online simultaneously, lowering the power drawn from each supply and potentially decreasing the power efficiency of each.

The power supplies in ProLiant ML and DL G6 servers support a high-efficiency mode. In this mode, one of the redundant power supplies is kept in a standby state, allowing the remaining supply to support the full power load and maintain a higher operating efficiency. The additional power supply is only brought online if the primary supply fails. Special circuitry ensures that the standby power supply can be brought online quickly enough if the first power supply fails.

HP Dynamic Power Saver allows a BladeSystem enclosure to operate at the highest possible power efficiency. When Dynamic Power Saver is active for a c-Class blade enclosure, the Onboard Administrator for the enclosure increases or decreases the number of power supplies that are active in order to keep them working at their most efficient level while maintaining redundant operation.

More efficient DC power regulation

DC voltage regulators on the server system board convert the 12 volt DC supplied by the power supply into the 5v, 3v, and other feeds used by the various system components. Voltage regulators in ProLiant G6 servers use more efficient components to deliver peak efficiency and to maintain greater than 90% efficiency over a broader range of power demands. This results in an 8-point gain in DC power efficiency over previous generations of servers.

ProLiant G6 servers also use DC phase shedding to further increase power efficiency. With phase shedding, the voltage regulator automatically adjusts the number of DC power phases as the power demands of a subsystem or component change. This produces additional increases in power efficiency. In G6 servers, phase shedding efficiently increases power delivery to the processor and memory subsystems.

Integration with Intel® Xeon® 5500 series architecture

The Intel Xeon 5500 series processors used in many of the ProLiant G6 servers feature Intel® Intelligent Power Technology. This is a set of features that can be used to lower power consumption of the processor and related subsystems when they are not fully utilized, and they include:

- More processor operating and idling states
- Ability to disable under-utilized processor cores
- Management of memory and I/O power consumption

HP engineering has worked closely with Intel to optimize these features in ProLiant G6 servers. In addition, many power saving features are configurable through the ROM-based Setup Utility (RBSU). This allows an administrator to configure the power and performance profile to match the server's intended use.

Optimizing processor power use with HP Power Regulator

HP Power Regulator is a standard feature on most HP ProLiant servers. It can automatically optimize processor power consumption based on server activity. Power Regulator is implemented in the system firmware and directly monitors the instruction load of the server processor(s) to determine the level of system activity. Power Regulator uses this information to continuously adjust the performance states, or p-states, of the processor(s) to match processor power consumption to the current workload without noticeably impacting overall system performance.

Optimizing internal cooling with sea of sensors

Servers must exhaust heat into the data center to keep their internal components operating within a safe temperature range. This internal cooling is part of the server power consumption. Each fan inside the server can consume a maximum of six to eight watts of power. Previous generations of severs used a few temperature sensors and a "fan curve" to set the fan speed to a preset value based on the measured temperature. The "sea of sensors" technology in ProLiant G6 servers uses up to 32 separate sensors to more accurately map the temperature profile inside the server. Instead of using a simple fan curve, a proprietary feedback algorithm continuously adjusts individual fan speeds to maintain specific temperatures. This improved monitoring technology prevents overcooling within the server and lowers the overall power consumption of the fans.

Overall power efficiency for ProLiant G6 servers

All of these technologies contribute to making the ProLiant G6 servers extremely power efficient. Figure 10 shows the calculated idle and maximum power for "large configuration" ProLiant DL380 servers across the generations. Although the ProLiant DL380 G6 has significantly better performance than a similarly configured DL380 G5, it consumes almost 200 watts less at maximum power. This results in a maximum power savings of about 4000 watts per rack in a data center.

Figure 10. Reduced power consumption of the ProLiant DL380 G6 server

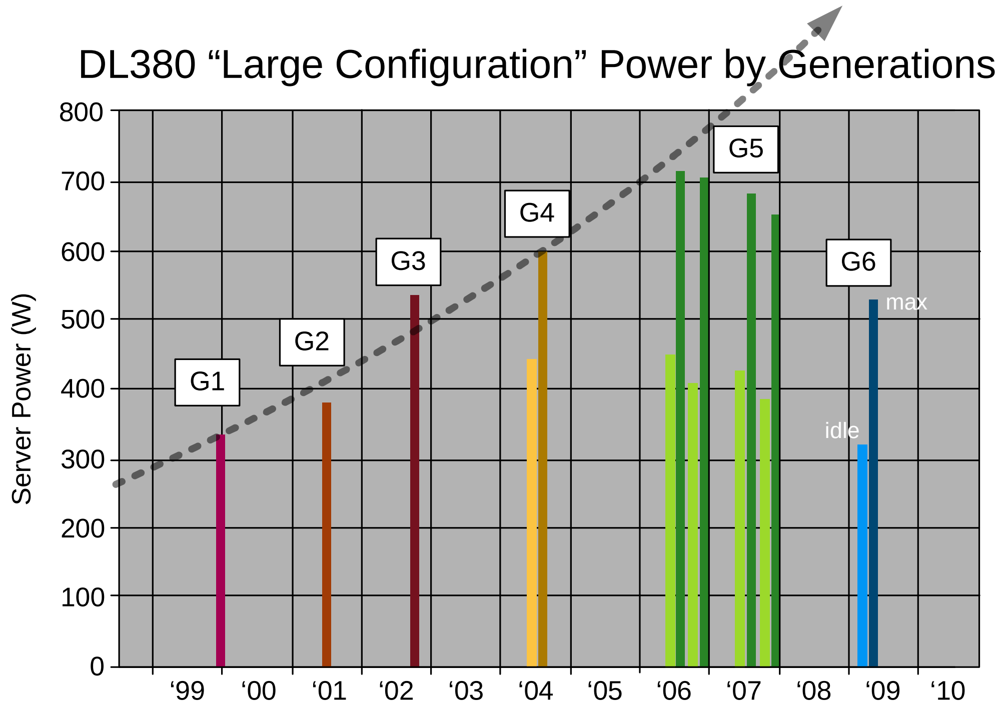

As a result of these efforts, the HP ProLiant DL360 and DL380 G6 servers are the first servers to receive the Energy Star® qualification.

Improving thermal management in the data center

Overview of Cooling

All data centers, no matter how efficient, consume significant electrical power. Power consumed is converted to heat. In order to maintain the data center at a temperature in which IT equipment can continue to operate, additional power must be expended to extract this heat. The goal is to remove this excess heat as efficiently as possible.

In general, cooling is accomplished by pumping cold air into the data center while simultaneously removing the exhaust air that has been heated by the IT equipment. In the past, facilities managers might resolve heat issues by over provisioning the cooling capacity of the data center, lowering the temperature of the entire data center to deal with hot spots. Increasing the cooling efficiency of the data center involves finding better ways to cool all areas of the data center optimally.

Improving cooling efficiency

As with power distribution, there are certain principles of cooling efficiency. Observing these principles can help an administrator increase the cooling efficiency of the data center:

- Promote efficient airflow. The more efficiently that cool air is delivered and hot air is removed, the more efficient the cooling. Keeping the cool air as separate as possible from the hot air is one of the main ways to increase cooling efficiency. Best practices such as having separate cool aisles and hot aisles in the data center help accomplish this.
- Maintain the highest possible "cool" temperature. The higher the temperature can be kept in the data center, the less cooling power is required. Using a higher "cool" temperature also increases the efficiency of the heat transfer that occurs outside the data center.
- Practice containment. The smaller the volume that needs to be cooled, the more efficiently the cooling can be accomplished.

Applying these principles through best practices, an administrator can increase cooling efficiency in the data center. HP provides several products that exploit these principles to provide additional gains in data center cooling efficiency beyond those offered through simple best practices.

Hot aisle containment using the HP Rack Air Duct

The HP 10000 G2 Rack Air Duct works with HP 10000 G2 racks to increase cooling efficiency by minimizing air recirculation around the rack. The rack air duct channels the hot exhaust air from the rear of the rack directly into the air return system, increasing overall cooling efficiency by eliminating hot spots and air mixing in the data center. Cooling efficiency is further increased since hotter air is returned directly to the Computer Room Air Handler (CRAH) units, allowing more efficient heat transfer.

The rack air duct has several distinct advantages:

- It is a completely passive system, with no moving parts.
- It can be retrofitted into existing data centers that use HP 10000 racks.
- It can increase efficiency enough to support cooling 15 kW rack configurations.

The rack air duct product belongs to the general category of data center cooling solutions that goes by the term hot-aisle containment. As the name implies, hot-aisle containment solutions are designed to more completely separate, or contain, the exhaust air from a data center hot aisle and evacuate it before it can re-circulate and mix with incoming cooled air being directed into the front of the racks from the cold aisle.

HP Modular Cooling System G2

The HP Modular Cooling System G2 (MCS G2) is a closed-loop cooling system that provides a self-contained cooling solution for one or two racks of IT equipment. The cooling unit in the MCS G2 uses an air-to-water heat exchanger tied to a data center's chilled water system to produce cold air that is fed to the front of the racks(s). The servers in the rack expel the warm exhaust air to the rear where the MCS fan modules re-direct it back to the heat exchanger to be cooled and re-circulated (figure 11).

Figure 11. Heat exchange and airflow in the HP Modular Cooling System G2

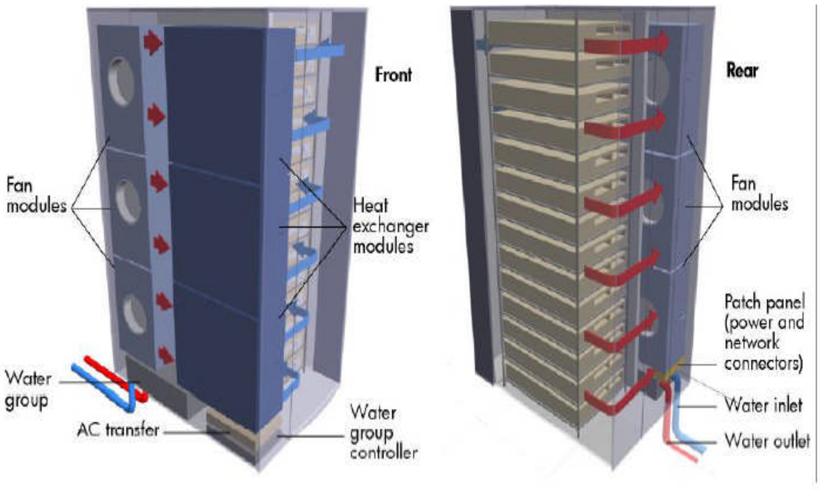

The MCS G2 is another example of a cooling containment solution. Isolating and reducing the volume of air that is being cooled increases the efficiency of the heat exchange. The MCS G2 can be used to cool up to 35 kW of total IT power consumption distributed across one or two racks.

The MCS G2 can be used to achieve different goals depending on the issues that need to be addressed. It can be used to increase capacity in an existing data center by supporting increased density. It can also be used to isolate an existing high density IT load in the data center that is causing a hot spot. By removing the hot spot, the overall data center cooling can be set to a more efficient temperature setting.

Data center monitoring using HP Datacenter Environmental Edge

As with the other phases of the power life-cycle, monitoring and managing data center cooling is an important part of the overall solution. To meet the requirements for managing heat and cooling in the modern data center, HP has introduced the HP Datacenter Environmental Edge, a remote monitoring system capable of providing a detailed picture of the environmental parameters of the datacenter.

HP Datacenter Environmental Edge uses arrays of environmental sensors to monitor temperature and humidity in the data center racks (Figure 12). Separate sensor assemblies are used to monitor air pressure in the air supply plenums. In addition, optional meters can measure the data center power, chilled water flow, and other parameters. Each sensor assembly is attached to an Environmental Base Station, a battery powered device that wirelessly sends the collected data to an HP Base Station Gateway where the data is then transferred over the network to the HP Datacenter Environmental Edge server to be stored and then displayed through a web-based user interface along with the data collected by additional sensor assemblies throughout the data center.

Figure 12. HP Datacenter Environmental Edge deployment

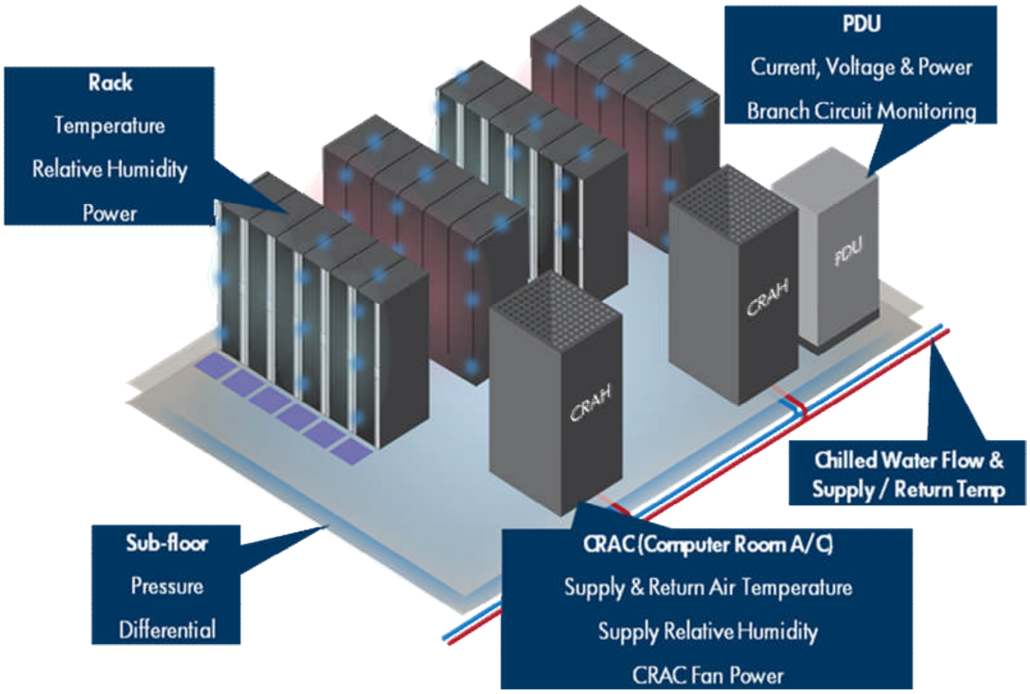

The architecture of the Datacenter Environmental Edge components provides several advantages:

- Fast and easy to deployment. Wireless battery powered base station reduces installation requirements for additional infrastructure.
- Deployable in any data center. Datacenter Environmental Edge can be installed in data centers with racks and CRAH or CRAC units from any vendor.
- Flexibility. It can be expanded or re-configured as required when data center layout or infrastructure changes.

HP Insight Environmental Observer, the user interface and reporting component of HP Datacenter Environmental Edge, provides a web-based software solution for monitoring and visualizing the environment of the data center based on the information collected by the Datacenter Environmental Edge. Using Environmental Observer, administrators can visualize the temperature, pressure, and humidity for the entire data center in real time or over a period of time (Figure 13). This allows them to visually locate hot spots and inefficiently cooled sections, as well as other trouble areas in the data center. This data helps them to plan appropriate corrective measures.

Figure 13. Visualization of the data center environment using HP Insight Environmental Observer

HP Insight Environmental Observer can also provide notifications of any critical changes in the data center environment based on parameters that the administrator defines. These alerts can be delivered by text, e-mail, or SNMP.

The Datacenter Environmental Edge is solely a monitoring and reporting solution. However, the data that it provides can be used to identify and rectify cooling inefficiencies in the data center. Based on HP case studies, administrators can use the information provided by Environmental Edge to lower their energy costs by ten to fifteen percent and reduce the typical datacenter PUE from 1.8 to 1.5.

HP Performance - Optimized Data center

The HP Performance-Optimized Data center (POD) is a self-contained data center that can be custom-configured and then deployed within a matter of weeks to provide HP customers with large scale-out capabilities. The 4000 square-foot POD (Figure 11) has up to twelve 50U racks and is capable of supporting up to 3520 compute nodes, 12000 large form factor (LFF) hard drives, or any combination of the two. The POD comes with its own fully integrated power and cooling infrastructure and can be hooked directly into 480 volt utility power distribution and into existing chilled water cooling.

Figure 11. HP Performance - Optimized Data center (POD)

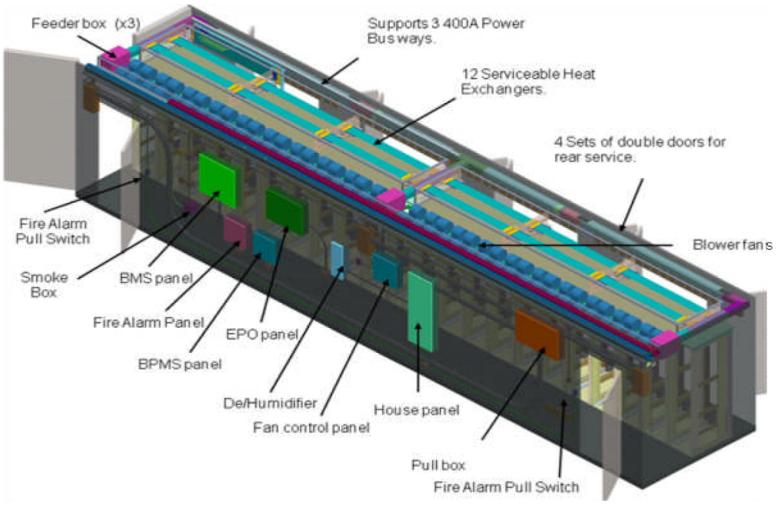

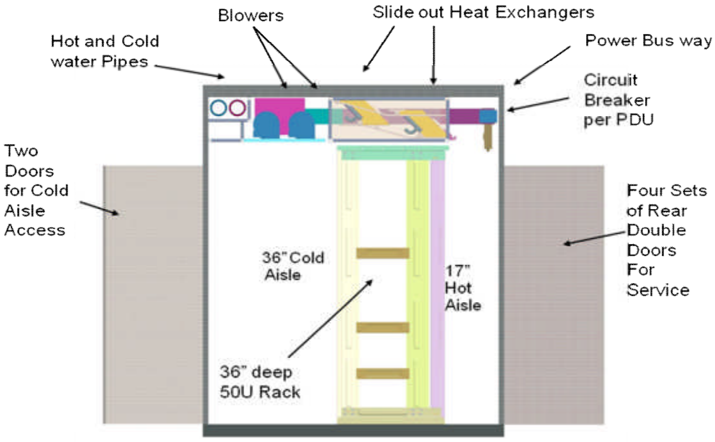

In many ways the POD is the ultimate power and cooling containment solution. The POD design takes maximum advantage of the efficiency principles that have been previously outlined. It minimizes power losses by decreasing the number of power conversions. Incoming 480v 3-phase is converted to 415v 3-phase AC and then delivered directly to the racks as separate 240v single phase feeds for distribution to the servers. Cooling efficiency is maximized though the engineering of a highly efficient airflow in a tightly contained space and a cold aisle that can be run at 90 degrees Fahrenheit (32 degrees Celsius). As a result, the HP POD can have a PUE as low as 1.25 depending on the configuration of the IT equipment.

HP Data Center Smart Grid

The HP Data Center Smart Grid is the framework around which HP will shape future development in power and cooling management. The Data Center Smart grid builds on the current data center energy efficiency portfolio to create an intelligent, energy-aware environment across IT and facilities to further optimize and adapt energy use, reclaim facility capacity and reduce energy costs.

HP Data Center Smart Grid will add more smart technologies that communicate the real-time time status of the IT equipment on power, cooling, utilization, available capacity, and more. It will also add capabilities to environmental monitoring to track facilities capacity and activity. Most importantly, its goal is to integrate these solutions together into an intelligent energy-aware system capable dynamically responding to changes in the combined IT and data center facility infrastructure.

Conclusion

Power and cooling have become the primary issues to be addressed in the data center and will probably remain so for the foreseeable future. Improving the basic power efficiency of the data center requires making improvements in all parts of the power life cycle—power distribution, power consumption, and cooling. Maximizing data center power and cooling resources requires technologies and solutions for monitoring and provisioning. HP is striving to deliver the most power efficient IT products possible while providing integrated monitoring and provisioning technologies as well as solutions to address the data center power and cooling challenge now and in the future.

For more information

For additional information, refer to the resources listed below.

Resource description	Web address
Technology brief Power basics for IT professionals	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01234421/c01234421.pdf
Technology brief Increasing energy efficiency with modular HP three-phase power distribution	ftp://ftp.compaq.com/pub/products/servers/proliantstorage/power-protection/3Phase_Tech_Brief.pdf
HP Datacenter Environmental Edge	http://h18004.www1.hp.com/products/servers/solutions/datacentersolutions/environmentaledge/index.html?jumpid=reg_R1002_USEN
HP rack and power infrastructure products	http://h18004.www1.hp.com/products/servers/platforms/rackandpower.html
Technology brief Optimizing facility operation in high density data center environments	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00064724/c00064724.pdf
HP Rack Air Duct	http://h18004.www1.hp.com/products/servers/proliantstorage/rack-options/rack-air-duct/index.html
Technology brief HP Modular Cooling System Generation 2 water cooling technology	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01490034/c01490034.pdf
HP Dynamic Power Capping	http://www.hp.com/go/powercapping
Technology brief HP Power Capping and Dynamic Power Capping for ProLiant Servers	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01549455/c01549455.pdf
Technology brief Power Regulator for ProLiant servers	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00300430/c00300430.pdf
Technology Profile The highlights of HP10000 G2 Series Racks	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01036998/c01036998.pdf

Resource description	Web address
HP Power Advisor	http://www.hp.com/go/hppoweradvisor

Call to action

Send comments about this paper to TechCom@HP.com.

TC091891TB, February, 2010