A White Paper from the Experts
in Business-Critical Continuity™

Using Virtualization and Digital Control Technologies to Increase Data Center Operating Efficiency

Liebert®

EMERSON™
Network Power

Table of Contents

Executive Summary

Data center professionals continue to be squeezed on several fronts as they seek to optimize operational readiness, availability, and risk management for operations on a local, regional or global basis. These include battling legacy hardware/infrastructure transitions, increasing power and heat densities, escalating service-level availability agreements, risk management and expense control.

There are digital systems available today that provide a framework within which data center professionals can manage their entire suite of hardware and software assets in a real-time, cost/risk-optimized manner. These systems can immediately reduce costs and improve performance of the data center infrastructure.

This concept can be implemented across legacy hardware, infrastructure systems and data centers—or with additional OEM support, throughout the data center of the future—to create a dynamic data center where operating expenses can be balanced in real time with service-level agreements and forecasts to maximize ROI and minimize risk.

This data center will be capable of reacting to changing operational parameters, within certain guide bands, via a combination of power and OS/Applications management at the server level, along with digital control of the local power and cooling infrastructure.

The net result will be a reduction in operating expense on the order of 20 to 30 percent or more, depending upon the current utilization rate and risk balance ratio. This can be accomplished without any reduction in service level agreements while providing a higher degree of disaster recovery.

Current control system models indicate year-on-year operational improvements on the order of 25 percent with no reduction in system availability or resiliency. Even greater gains can be realized by analyzing data center utilization patterns.

Introduction

Data center, IT and facility professionals find themselves battling increasing IT power needs at a time of rising energy costs.

New chips and servers will enable IT organizations to do more with less; however, in most cases, improvements in efficiency will be more than offset by increased demand for power and capacity.

The data center is rapidly moving from the relatively stable, controlled environment with power densities of 50 to 75 Watts per square foot to large "Easy Bake™ ovens" running 200 to 300 Watts per square foot, with forecasts for 500 Watts per square foot around the corner. As data centers consume more energy and more resources, they have become a target for increased government regulation regarding environmental impact.

This paper focuses on tools that are being developed to balance work load in the data center to improve thermal management. It describes how real-time, historical and predictive information—combined with server virtualization, a service-oriented architecture and digital power and operating system controls—can optimize operating efficiencies.

This approach builds upon previous efforts with a digital power control scheme to manage the thermal envelope and improve data center efficiency. It proposes a framework that allows work loads to be shifted away from underutilized assets and uses digital power controls to cycle these assets on and off to improve operational efficiency.

This approach is enabled by emerging measurement and monitoring capabilities that provide the ability to monitor power consumption from the rack/chassis level to the system component level, including CPU power, memory power, system board power, system power and rack power.

When power consumption measurement is extended to storage systems, switches and hubs, and integrated with temperature data, power consumption can be mapped across the data center and used to deliver more efficient and effective cooling.

With a consistent data model from all systems (computer, power, HVAC, etc.), power requirements, cooling requirements and utilization can be predicted. Properly architected and instrumented, a control system could then:

- Monitor utilities
- Monitor data center impact on emissions based on energy supplied and energy consumed
- Audit compliance with "green" or "emissions" commitments
- Support improved decision making regarding energy risks

Current models indicate year-on-year operational improvements on the order of 25 percent with no reduction in system availability or resiliency. Even greater gains can be realized by analyzing data center utilization patterns.

Data Center Utilization

A 2004 study by Electrical Power Research Institute projects U.S. data center market power demands of 14.8 Terra Watt hours per year. With the continued replacement and refreshment of legacy hardware—typically in the 200 to 500 Watt class—with new multi-core, high-density servers in the 1,000 to 5,000 Watt class—total domestic data center demand can be expected to increase at an accelerated pace.

Complicating this situation is the nature of the core power supply/chip/operating system relationship within enterprise-class servers. Unlike Energy Star desktops with sleep and hibernate modes, the typical enterprise server has two modes of operation: on

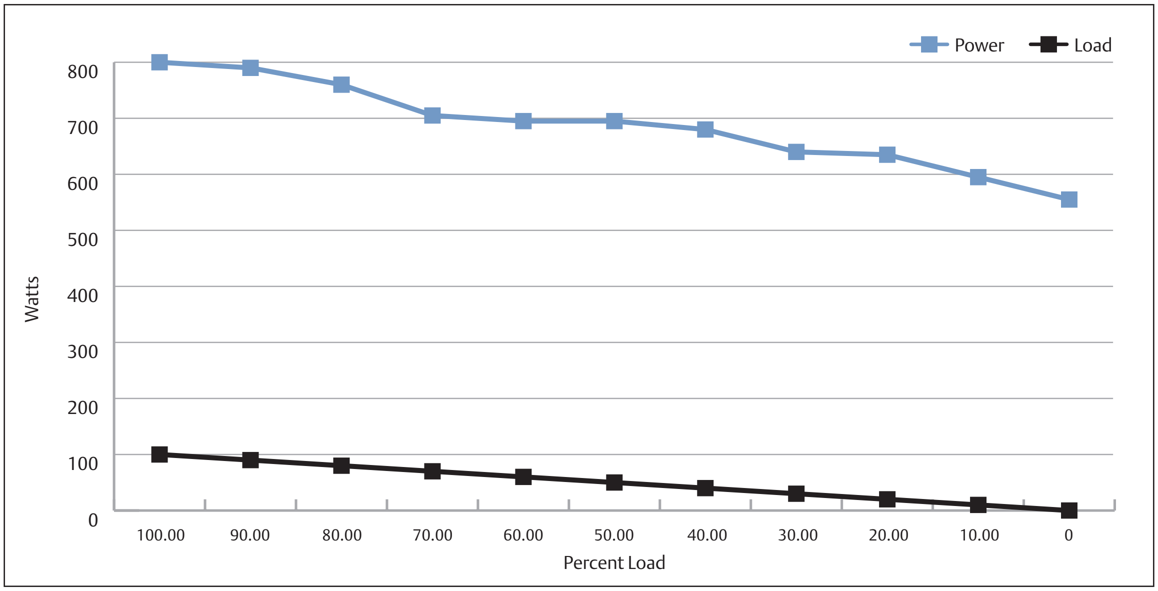

Figure 1. Server power consumption remains relatively high as load decreases.

Figures 2 through 8 show actual data center utilization measurements for different scenarios based on time and function. They illustrate the time-of-day, day-of-week, and seasonal-use characteristics that contribute to vast periods of under utilization.

and off. In idle mode, these systems consume between 70 and 85 percent of full operational power (see Figure 1).

Our research indicates that typical data center servers are being utilized at an average rate of between 15 and 20 percent annually. However, averages can be misleading as these are generated on a 7x24x365 basis, including periods of low demand. This is further complicated by 2N design considerations coupled with peak demand forecasts compounded by “head room” capacity.

Figures 2 through 8 show actual data center utilization measurements for different scenarios based on time and function. They illustrate the time-of-day, day-of-week, and seasonal-use characteristics that contribute to vast periods of under utilization. Increased control of utilization during these periods through digital power management can yield tremendous savings as well as overall reliability and availability improvements.

In the time-of-day business scenario (Figure 2), demand is spooled-up as the day progresses. Some work leads the “normal” 8 a.m. to 5 p.m. workday, either through international business or time-of-use load

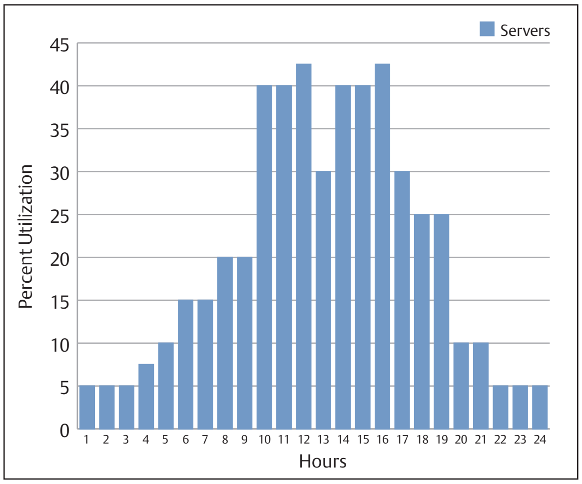

Figure 2. Time of Day Utilization – Business

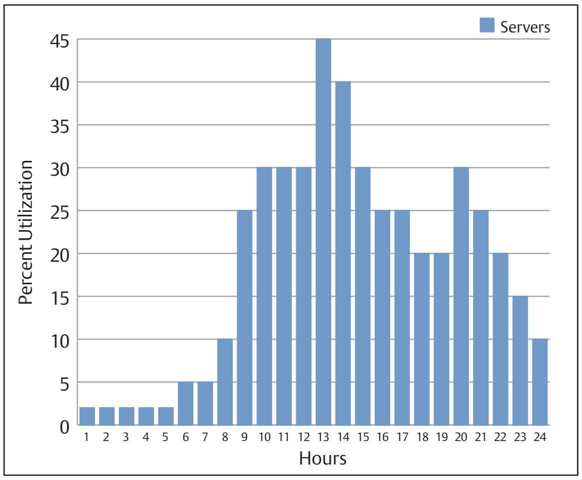

Figure 3. Time of Day Utilization – Consumer

Consumer-focused businesses present a somewhat different profile, ramping up later in the day and extending at higher levels well into the evening hours.

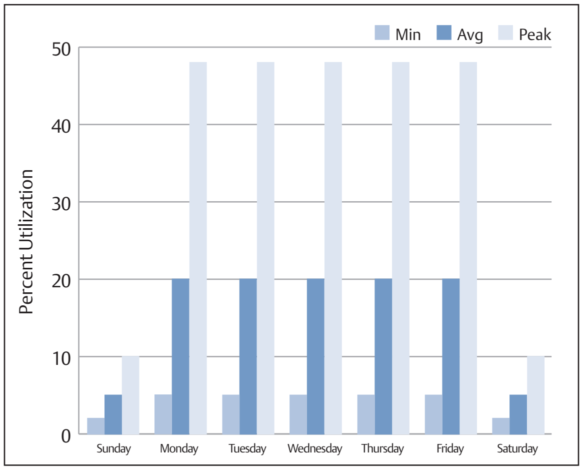

Figure 4. Day of Week – Business

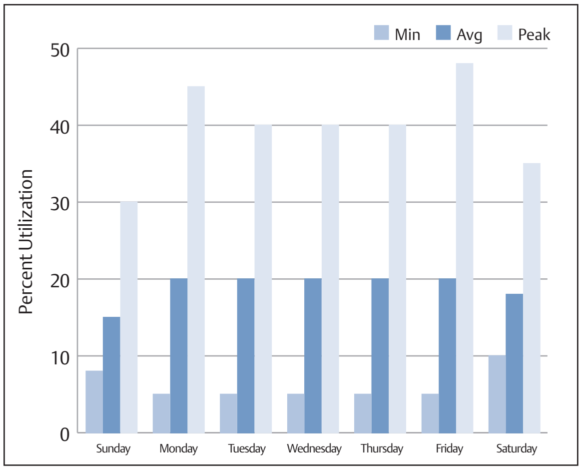

Figure 5. Day of Week – Consumer

shifting from another time zone. In addition, certain core business functions may be brought on line early by the IT staff to ensure immediate availability. Similarly, the day "end" extends well beyond normal business hours. Load shifting, time zones, and evening employee access, account for the bulk of above-baseline loads. Despite this shifting, there are several hours each day when server loads are at the average or far below average. These hours represent opportunities for operational improvement.

Consumer-focused businesses present a somewhat different profile (Figure 3), ramping up later in the day and extending at higher levels well into the evening hours.

This is typical for sites such as Yahoo, iTunes and ESPN on a regional server basis. As with the business data center, an opportunity for improvement exists as there is at least one quarter of the typical day when server use is well below average.

The typical business data center is not weekend/holiday optimized (Figure 4). Though staff may be minimized and some facility managers allow for a slight change in environmental conditions during the weekend, server power consumption is not optimized.

The consumer data center does not have as dramatic a drop in weekend utilization as the business data center (Figure 5); however, it

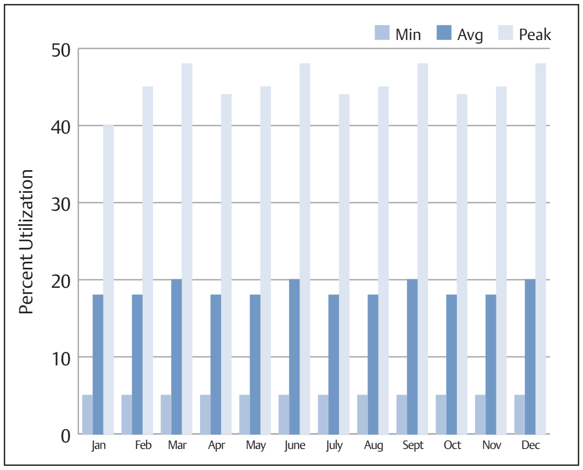

Figure 6. Monthly/Seasonal – Business

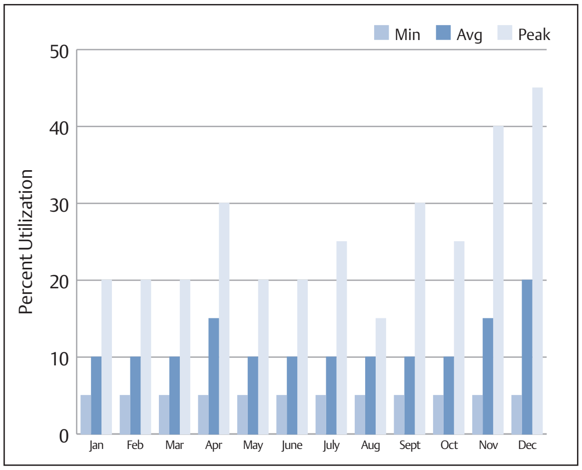

Figure 7. Monthly/Seasonal – Consumer

does have a historical profile and predictive element that can be used to reduce the number of servers on line.

In the business data center, there tends to be minor fluctuations with quarter and year-end activities (Figure 6). Specific months may vary based upon fiscal reporting periods.

The holiday season drives increased utilization in major consumer e-commerce sites, such as Amazon, Wal-Mart and Sears (Figure 7).

Government and academia have their unique seasonal profiles as illustrated by the example of the IRS (Figure 8). Short of the build up to April 15, the extensions in May, and the corporate quarterly filings, there are vast periods of the year when data center assets are sitting idle.

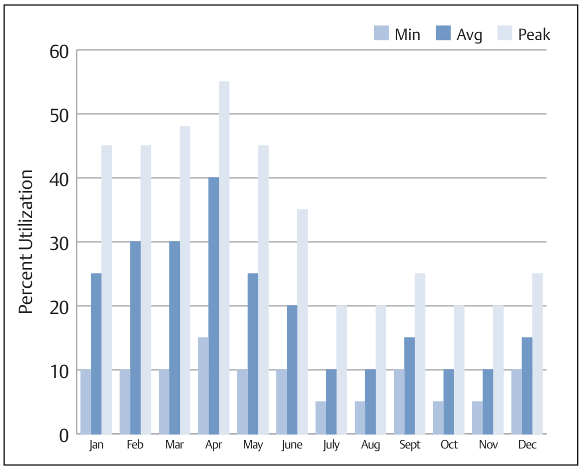

Figure 8. Monthly/Seasonal – IRS

Power Management Opportunities

Analysis of data center time-of-use issues reveals varying layers to approaching a total digital power solution.

Server Efficiency
To fully optimize the gains from this approach, it is necessary to start with the proper “light bulb” (server). It doesn’t make sense to deploy 100 Watt tungsten-filament light bulbs if exactly the same performance could be achieved with a 29-Watt compact fluorescent light (cfl).

While the number of compute-cycles-per-Watt of servers has dramatically increased, so too has the power density per server. The net result is faster, hotter processing. The server equivalent of the compact fluorescent bulb—one that could perform the same function using about 25 percent of the energy—has yet to be brought to market.

Were this type of server available today, it would create a reduction of 530 kWh for a facility with 1,000, 1 kW servers (1,000 x 1,000 x .8 actual load factor x .25). At $.08 a kWh, a data center operator would realize annual savings of approximately $370,000 for every 1,000 servers.

For the balance of this paper, a 1,000-server data center, with each server rated at 1 kW (.99 power factor), will be used as the basis for analysis of the potential for efficiency improvements of various technologies and strategies.

Power Supply Efficiency
An opportunity that presents great potential in the short term is power supply efficiency. Legacy server power supplies operate in the 75 to 80 percent efficiency range. However, the technology is now available to deliver digitally controlled power supplies with efficiencies in the range of 88 to 92 percent.

An opportunity that presents great potential in the short term is power supply efficiency. Legacy server power supplies operate in the 75 to 80 percent efficiency range. However, the technology is now available to deliver digitally controlled power supplies with efficiencies in the range of 88 to 92 percent.

The Electrical Power Research Institute has issued an 80/20 guideline for desktop-level power supplies that specifies a minimum operating efficiency of 80 percent from 20 percent of load to full load. Server manufacturers have the ability to meet or beat this guideline.

Moving from a 77 percent efficient power supply to one nominally rated at 88 percent (90 to 92 percent efficient models are available at power levels of less than 3 kW)

Power Supply Operating Costs

Power Supply Efficiency	Energy used by 1,000 1 kW servers at 80% Load	Energy costs at $.08 kWh
77%	(1,000 x (1,000 x .8)/.77 = 1,048 kWh	$83.84
88%	1,000 x (1,000 x .8))/.88 = 909 kWh	$72.72

Table 1. Hourly savings of $11 can be achieved by using high-efficiency power supplies on 1,000 servers. This translates into annual savings of approximately $97,000.

would generate savings that would more than pay for the cost of the power supply within the first year (see Table 1).

Capacity Management

Beyond simple power supply efficiency, there is the issue of the server-power supply-OS relationship and how predictable variations in time-of-day and seasonal usage can be used to reduce energy demands from servers at idle.

Unlike desktop and laptop computers, enterprise-class servers rarely “sleep” or hibernate. In the case of a 1 kW server at full rated CPU and data I/O throughput, peak power demand is on the order of 800 Watts. At idle, this same server still consumes between 560 and 680 Watts (see Figure 1).

Given the nature of networks, servers and service-oriented architectures, it is not likely that enterprise-class servers, especially those greater than 500 Watts, can reach a true “sleep” state of less than 10 Watts. However, it may be possible to achieve a “drowsy” rating of less than or equal to 25 percent of actual full-rated power. This 25 percent allows for network “presence” and connectivity while taking RAM and processor(s) offline beyond the rapid-on response state.

In looking at a 2N data center with 1,000 1 kW servers and a 50 percent peak utilization, applying a “drowsy” mode to the 2N half of the server population yields significant potential savings with both new and legacy power supplies (see Table 2).

Idle Energy Use – No Drowsy Mode		Idle Energy Use – Drowsy Mode	
77% Efficient Power Supply			
500 servers x 560 Watts per hour/.77 efficiency	363 kWh	500 servers x (800 Watts per hour x .25)/.77 efficiency	130 kWh
88% Efficient Power Supply			
500 servers x 560 Watts per hour/.88 efficiency	318 kWh	500 servers x (800 Watts per hour x .25)/.88 efficiency	114 kWh

Table 2. A drowsy mode on servers would save approximately 230 kWh, or $162,000 a year at $.08 kWh energy costs, with current power supply efficiencies. With new high-efficiency power supplies, savings would be 200 kWh, or $140,000 a year.

Individual server-power supply systems have their own unique efficiency curves and power profiles so users should analyze new server specifications to better understand their characteristics.

Real-World Power Management

While none of the above is truly in place, there is an interesting mix of available processor "hooks," power and temperature modeling programs, server power management tools and infrastructure controls that could be integrated with digital controls to provide a real-time data center analysis and control mechanism.

Servers
New monitoring tools coming on the market provide access to current temperature data, utilization rate, and various levels of state information. Through software-enabled management techniques, loads can be shifted around the data center to maximize use of targeted servers while leaving the largest statistically viable number either "drowsy" (when available) or powered off.

Turning servers off in a functioning data center requires meaningful data and statistical analysis—real time, historical, and forecasted—of a host of components. These can be viewed as a classic pyramid beginning with the CPU/memory device at the base and building up through the server/storage/communications device fans and power supplies, out to the integrated rack level. This is then layered with power (PDU, UPS, batteries), cooling (spot, supplemental, room), room infrastructure (lighting, fire suppression), and finally building/site level (utilities, generators, geo-political, weather, natural events) to produce an integrated view of data center operations.

> Note: A unifying communications mode or protocol does not exist at this time to monitor and control these various disparate hardware devices. Whether vendor-specific or consortium-standard, a unified approach to command and control via a digital power operating system would be highly beneficial.

Using internal server and storage temperature/state data, power and cooling systems data, and in-room environmental monitoring, combined with historical time of day/seasonal models, it may be determined that servers and storage are being underutilized. The data center is "cold" and power loadings are light.

In this case, the data can be examined for potential capacity-use benefits. Can alternate data center applications be shifted in real time to this site? A more likely scenario is that with the diminished actual and forecasted loads, 10 percent of all servers are turned on and the balance cycled off. This is 20 percent of N capacity, enabling an instant response eight times actual base-load conditions. As savings below 10 percent of rating are minimal, this provides a good floor of operation upon which the IT/data center manager can build availability well in advance of demand in a risk-free mode. With the server volume reduced, the associated power and cooling infrastructure can also be reduced in scope.

Power Infrastructure Control
Using the example just discussed, the UPS plant would be lightly loaded, operating in the 85 to 88 percent efficiency range, depending upon vendor and model.
With the data available through real-time measurement and monitoring, individual UPS units could be cycled off to increase the load on the remaining modules.

Using the 1,000 server data center, a typical 2N UPS plant would comprise four, 500 kVA UPS modules with a .9 Power Factor in a multi-module configuration. With only 10 percent of the servers cycled on, the UPS load is 80 kW or five percent of the total plant.

Turning servers off in a functioning data center requires meaningful data and statistical analysis—real time, historical, and forecasted—of a host of components.

CRACs are currently available with intelligent communications, such as the Liebert iCOM system, that enable CRAC units to recognize current state (humidification, dehumidification) as well as I/O temperatures and percent load.

Under this scenario, the UPS units are well off their efficiency curve. Cycling two units off brings the loading up to 10 percent of rating. Though one or two points of additional efficiency may be gained when compared to the total power (80 kW), the net gain is minimal. Because time of day and seasonal data further indicates that this loading factor is only valid from midnight to 5 a.m., a safety factor can be built into the overall command program to prevent such cycling at the risk of losing UPS availability at critical periods.

By looking at a seasonal model, the overall plant efficiency may be improved by reducing the number of UPS units in service (provided they are alternated between duty in comparable time increments to improve overall life expectancy and component stress/wear factors). When the historical data and forecasted data indicate long periods of under use, such as in the consumer model or slow periods in a quarterly fiscal cycle, cycling UPS modules presents a minimal instant-demand risk while improving efficiency.

Cooling Infrastructure Control

From a cooling perspective, workload management can be extended to provide control conditions for the computer room air conditioners (CRACs) as well as any extreme-density, point-of-load cooling systems. In this mode under certain light-loading conditions, further overall savings can be realized from the cooling infrastructure. Further, by integrating with the building management system (BMS) it may be possible to dynamically vary the balance of plant via the chillers to achieve additional efficiencies.

Assuming a lightly loaded condition with a historical and forecasted model calling for extended periods of reduced load (greater than one hour as chiller and CRAC units have ramp up and latency issues), it would be possible to cycle down certain aspects of the cooling infrastructure.

To a degree, this occurs automatically in some current generation cooling systems. CRACs are currently available with intelligent communications, such as the Liebert iCOM system, that enable CRAC units to recognize current state (humidification, dehumidification) as well as I/O temperatures and percent load. These units can be configured to perform some degree of automatic system control. As they represent the slowest response mechanism to dynamically variable cooling requirements, and because of the nature of the wear mechanism on CRAC compressors, the best case is to equip CRACs with variable compressors or step-load, four-stage compressors.

In a high-density server application, the data center would require supplemental cooling from a point-of-load cooling system. This system too could be brought under control to maximize performance.

Supplemental cooling modules are located close to the thermal load, allowing them to operate 24 to 35 percent more efficiently than CRAC units on a per Watt of cooling capacity.

Therefore, the ideal system blends supplemental cooling and CRAC systems based upon real-time data and historical/forecast models to maximize supplemental cooling use while covering base-load cooling, humidity control and filtration requirements. Supplemental cooling systems would be kept running under all scenarios, even when forced to auto-cycle to a 10 percent duty cycle, as they represent the most efficient cooling solution and must be readily available when server loads increase above base-load cooling capabilities.

Savings could be realized on the supplemental cooling plant by reducing fan loading. Using server workload measurement, supplemental cooling fans could be shut off at localized racks either on a zone, row, or block basis, thereby balancing the room for cooling. At approximately 150 Watts saving per unit cycled off at 50 percent loading, an additional 6 kWh savings could be realized on top of the already reduced loading on the compressor.

Aggregating Savings

Table 3 shows energy costs with and without digital control of a 1,000 server facility with 20 percent utilization using low efficiency power supplies and high-efficiency power supplies.

These examples show real-world savings of more than one megawatt of power per hour. Capturing the 20 percent load portion of a typical data center on the weekends/holidays could lead to a savings of 1,770 megawatt

Table 3 shows energy costs with and without digital control of a 1,000 server facility with 20 percent utilization using low efficiency power supplies and high-efficiency power supplies.

77% Efficient Power Supply			
20% Load – No Control		**20% Load – With Control**	
20% of servers at full load 200 x (1,000 x .8 LF)/.77	208 kWh	20% of servers at full load (200 x (1,000 x .8 LF))/.77	208 kWh
80% of servers at idle 800 x (1,000 x .7)/.77	727 kWh	80% of servers powered down	0 kWh
Cooling .5 kWh per kW load @ 80% to 100%	395 kWh	Cooling .30 kWh per kW load @ 10% to 40%	62 kWh
Power .1 kWh per kW load @ 70% to 100%	79 kWh	Power .11 kWh per kW load @ 10% to 70%	23 kWh
TOTAL	**1,409 kWh**	**TOTAL**	**293 kWh**
88% Efficient Power Supply			
20% of servers at full load (200 x (1,000 x .8 LF))/.88	181 kWh	20% of servers at full load (200 x (1,000 x .8 LF))/.88	181 kWh
80% of servers at idle 800 x (1,000 x .7)/.88	636 kWh	80% of servers powered down	0 kWh
Cooling .5 kWh per kW load @ 80% to 100% (790 kW)	345 kWh	Cooling .30 kWh per kW load @ 10% to 40%	55 kWh
Power .1 kWh per kW load @ 70% to 100%	69 kWh	Power .11 kWh per kW load @ 10% – to 70%	23 kWh
TOTAL	**1,231 kWh**	**TOTAL**	**259 kWh**

Table 3. Savings between 970 kWh and 1,170 kWh can be achieved with digital power control, depending upon power supply efficiency.

Current data center practices with their focus on uptime and availability are coming under pressure from increasing power densities, escalating energy costs, and looming environmental legislation to improve operating efficiencies.

hours or $141,000 per year. Using the same scenario with 88 percent efficient power supplies yields annual savings of 1,510 megawatt hours or $120,000 per year.

Data center professionals can establish a significant ROI for deploying an automated intelligent real-time control architecture to their data center operations by extrapolating other periods of savings across an entire year through examination of 30, 40 and even 50 percent loading periods along with extended periods of low seasonal activity.

Tools of Engagement

Those starting from scratch with a greenfield data center or a build-out are in the best position to explore this solution by evaluating and implementing the following capabilities where feasible:

- High-efficiency power supplies
- Server/storage-level power, utilization, and temperature monitoring
- Row/room level environmental monitoring
- Intelligent power-control systems from individually monitored and controlled power strips at the rack level to managed UPS systems
- Load-variable CRACs with intelligent communications
- Monitored/managed supplemental cooling system with local rack-level fan controls

These tools can also be applied to legacy data centers in concert with the next server refresh/augmentation to take advantage of multi-core, multi-threaded technologies, service-oriented architectures, and the software tools and reporting hooks that are available in new generation platforms. There is reason to believe that between rack and row-level environmental monitoring units, real-time server utilization profiles from network management systems, such as Tivoli™ or OpenView™, and managed power strips, an existing data center could be retrofitted to achieve 50 to 60 percent of the previously referenced savings.

Conclusion

Current data center practices with their focus on uptime and availability are coming under pressure from increasing power densities, escalating energy costs, and looming environmental legislation to improve operating efficiencies. The time is right and the technology is available to consider the case for a dynamic data center as described in this paper.

The first avenue of exploration begins with specifying the proper server, software, and infrastructure suite to properly position the data center for monitoring, communicating, adapting and controlling. To that end, data center professionals should start by asking for the most efficient power supply available.

With the proper server, power and cooling infrastructure in place, the deployment of server power monitoring capability can be implemented across the data center, reaching out to the associated power and cooling equipment rooms. This system, once tuned and backed with sufficient real-world time-of-day, day-of-week/month, and seasonal utilization models, can be used to maximize asset utilization and reduce operating expenses.

Emerson Network Power
1050 Dearborn Drive
P.O. Box 29186
Columbus, Ohio 43229
800.877.9222 (U.S. & Canada Only)
614.888.0246 (Outside U.S.)
Fax: 614.841.6022

EmersonNetworkPower.com
Liebert.com

WP150-47

Emerson Network Power.
The global leader in enabling Business-Critical Continuity™.

EmersonNetworkPower. com

- AC Power
- Connectivity
- DC Power
- Embedded Computing
- Embedded Power
- Monitoring
- Outside Plant
- Power Switching & Controls
- Precision Cooling
- Racks & Integrated Cabinets
- Services
- Surge Protection