THE MAGAZINE OF 7x24 EXCHANGE INTERNATIONAL

# NewsLink

FALL 08

## MINIMIZING ENERGY COSTS IN DATA CENTERS: ENERGY MANAGEMENT BY DESIGN

## Improving Data Center Reliability through Commissioning and Energy Efficient Design

# Minimizing Energy Costs in Data Centers: ENERGY MANAGEMENT BY DESIGN

**Barry J. Needle**

The challenge for today's facility and data center managers is to reduce operating costs by minimizing energy consumption while maximizing power reliability, IT application performance, and availability.

Lately, in an effort to understand and improve IT energy efficiency, the industry has adopted the Power Usage Effectiveness (PUE) and Data Center Infrastructure Efficiency (DCiE) metrics supported by The Green Grid organization[1]. These metrics, derived via facility and IT equipment power ratios, facilitate more informed methods of reducing the total cost of operating data centers while managing increased service demands. Functions such as data center commissioning and system modification testing schemes address how much risk is prudent in the quest to save energy, a benefit found in the information revealed by the power ratios. Furthermore, the PUE and DCiE metrics allow individual operators to measure the effectiveness of efficiency improvement programs by comparing the efficiency of their own facilities to those of like organizations.

At the core of all actions to improve energy efficiency, reliability and availability is the answer to the question of cause and effect. Any change to the operation of a data center will have an effect on the amount of power the IT equipment uses, the quality and stability of the power distribution system and the health of the computing facilities output. Moreover, since every data center is unique in design and operation, the actions taken to manage energy will be most effective if formulated specifically for each data center site.

### The Data Center Powering Crisis

Data center IT equipment and the supporting power and cooling infrastructure are up to 40 times more energy intensive than a typical office building. From an energy usage perspective, a data center is more like a dense, power-packed industrial facility brimming with power-drawing equipment.

According to a recent EPA report[2], the power demand of the data centers in the U.S. is significant and growing.

- The energy consumption of servers (including cooling and auxiliary infrastructure) in U.S. data centers has doubled in the past five years and is expected to almost double again in the next five years [2011] to more than 100 billion kilowatt-hours (kWh), costing more than $7.4 billion annually (2005 dollars)
- "The peak load on the power grid from these servers and data centers is currently estimated to be approximately 7 gigawatts (GW), equivalent to the output of about 15 baseload power plants. If current trends continue, this demand would rise to 12 GW by 2011, which would require an additional 10 power plants."
- Data centers consumed about 60 billion kWh in 2006, roughly 1.5 percent of total U.S. electricity consumption.

It's no mystery to data center operators that most of the energy pushed to servers ends up as heat. The cost to power and cool racks of installed servers is quite significant and is forecasted to become even greater in relation to new server spending. Figure 1 graphically shows the increasing proportion of power and cooling relative to server spending worldwide, that for each dollar spent on a new server in 2005, forty-eight cents was spent on power and cooling. This is more than twice the ratio in 2000. In 2010, it is projected that this ratio will rise to $1:0.71

Unfortunately, expenses do not end with powering and cooling servers. There is

**Projected CO2 Emissions Associated with the Electricity Use of U. S. Servers and Data Centers (MMT -CO2/Year), All Scenarios, 2007 to 2011**

| Scenario | 2007 | 2008 | 2009 | 2010 | 2011 | 2007-2011 Total | % of current efficiency trends scenario |
|---|---|---|---|---|---|---|---|
| Historical Trends | 44.4 | 51.2 | 59.2 | 69.2 | 78.7 | 302.8 | 111% |
| Current Efficiency Trends | 42.8 | 47.9 | 53.6 | 60.5 | 67.9 | 272.8 | 100% |
| Improved Operation | 34.8 | 39 | 43.5 | 48.4 | 53.1 | 219 | 80% |
| Best Practice | 30.2 | 30 | 29.8 | 29.7 | 30.1 | 149.8 | 55% |
| State-of-the-Art | 28.1 | 25.7 | 23.5 | 21.4 | 21.2 | 119.9 | 44% |

- EPA, 2007

potentially an extreme price to pay for no power as well: idle servers. Unintended downtime is costly, at an average price of a million dollars an hour. That's what IT system downtime costs American business, according to a keynote address by the META group (now Gartner, Inc.) given six years ago — and the cost has yet to come down.

It should be noted that the forecasted power demand cited previously does not reflect unmitigated historical growth extrapolations. During the last several years, the industry's attention to the growing crisis of unbridled energy demand has fostered many of the ideas currently generating the positive effects of energy-efficiency trends. According to the EPA report, however, there remains significant potential for further improvement in reducing future energy demands and realizing notable environmentally "green" effects. Three energy efficiency scenarios were developed to explore the impact of new technological approaches that could be deployed without unacceptable risk to data center performance, reliability and availability.

Dubbed "improved operation", "best practice" and "state-of-the-art," the envisioned improvements in energy efficiency would be significant resulting in a potential dramatic reversal of energy demand trends. The accompanying reduction in the total carbon footprint (green house gas emissions) from the operation of IT facilities is noteworthy given the greater than 21% contribution[3] to total greenhouse gas emissions from power plants. The annual savings in 2011 ranging from approximately 23 to 74 billion kWh is compared to the current efficiency trends scenario. Annual electricity costs would be reduced by $1.6 billion to $5.1 billion.

The projected savings in electricity use corresponds to reductions in nationwide carbon dioxide (CO2) emissions of 15 to 47 million metric tons (MMT) in 2011.

### Breakthrough in Data Center Design?

A notion not often considered is that a data center's consumption of energy is **by design**. The design and specification process — from intended server loading and IT equipment selection to the power and cooling infrastructure — determines the center's power demand. All the engineering done for detailing the support infrastructure (including power distribution and backup, server and electrical room cooling and lighting systems) is to ensure the reliable operation and maximum uptime of the installed IT engine.

The data center's initial design is a snap-shot: once built, modifications are tested by trial and error. A company's computing services are far from being a static environment. Data center managers are called upon to implement new application loading schemes resulting in server consolidations and virtualization.

Growth means new energy-efficient equipment replacing old, thin-provisioning equipment and capacity expansion assessments. The "error" part of testing is risky and potentially costly, and today, all changes are being audited for energy efficiency, risk of system instability and cost impact. One way or another, minimizing the power demands (cost of operation) of data centers while maximizing reliability and availability is the goal.

The issue is how to ensure system reliability and uptime while managing power usage and risk of system upset cost-effectively. At a recent industry conference, the question of data center energy efficiency by design was considered with the hope that the IT industry might discover a technological breakthrough which could radically alter the design of data centers and economics of energy use. To the chagrin of the participants, no dazzling answers were forthcoming. Rather, the conclusion was that data center designers and owners needed to "tune up what they own" since no solutions were offered up.

### The Energy Efficiency Tune-up Tool — Energy Management by Design

The Uptime Institute Design Charrette 2007[4] found a number of conditions which were necessary for the tune up:

- A way to generate the metrics that show what performance levels can be reached and at what level systems are currently performing. "This granular benchmarking drives the tune-up process."
- Thorough knowledge and operational experience of a specific data center.
- Practical engineering and economic analysis training with implementation skills focused on reducing risk of unintended downtime or reliability issues.

### Power Analytics to the Data Center Efficiency Rescue

Power analytics is a design and simulation software methodology that allows electrical designers, facility managers and data center managers to easily perform highly-accurate simulations of their infrastructure design, under an almost unlimited range of operating conditions. Power analytics gives electrical engineering professionals the means to create a robust electrical "designbase" — a detailed design and knowledge base of the performance specifications of the entire electrical distribution system.

The key to an effective energy management program is accurate information regarding the consumption of energy. Based on the amount of IT equipment in racks, the power distribution and cooling equipment infrastructure, and the variations in application loading, power analytics can report accurate, real-time energy usage. This data can be compared to the "as-designed" energy usage calculated by the

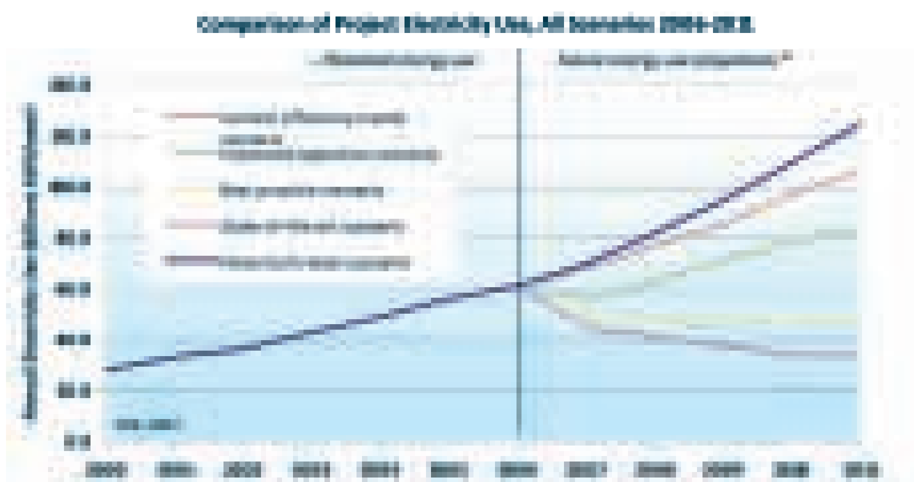

analytic system to give insight into system unbalances, capacity restraints, or overloads. The results of virtualization and other energy efficiency measures can be followed and assimilated. Intelligent power analytics can suggest scenarios for improved energy utilization based on its predictive diagnostics ability and by "what-if" simulation. At the current energy costs (~$0.089 kWh), a nominal realized annual savings of ten percent for even a relatively small, lightly-loaded data center is significant — greater than $100,000.

Today's facility engineers are generally focused on the reliability and capacity of the data center's power distribution system while the data center manager is concerned with server availability and service level agreements. While they may be preoccupied with different aspects of data center operation, they both are in agreement regarding taking risks: they do not want to take them. The adage, "no pain, no gain" is simply not part of their conversation.

The fact is, though, energy conservation schemes, thermal efficiency advances, capacity improvements, server loading rearrangements, new technology applications, and other energy management measures involve the risk of unintended consequences. The simulation of a system's performance in a virtual environment is the safest way to test a system modification and assess risk.

Power analytics aptly coined "virtual electrical expert" offer a technological richness needed to analyze systems from a variety of perspectives; from static to dynamic simulations including the ability to model and embed the detailed control logic of the intelligent electronic devices responsible for controlling how power flows and how it is directed throughout the system.

In addition, a robust library allows users to perform specialized forms of analysis and optimization, including Fault Analysis, Protection Coordination, Power Flow Analysis, Power Quality Analysis and Mitigation, Dynamic Behavior Simulation, Design Optimization, and Sizing Optimization.

When installed, the power and uniqueness of the platform is derived from the complete encoding of the design specifications from the original, as-built power infrastructure. All power system electrical parameters are calculated from the stored design specifications. During the data center's normal operation, the parameters are compared with the real-time power data. The intelligence of power analytics can accurately corroborate as-specified power parameters, determine if there are system anomalies, and predict when and where there are potential vulnerabilities for system and equipment failure.

More advanced power analytics systems on the market — AGO, LLC. recommends the Paladin® family from EDSA Corp. — allow users to actually capture current system state data, and run detailed "what if" simulations to verify system operations for the data center commissioning process, to investigate the effects of equipment rearrangement, configuration modifications, capacity expansion and other data room modifications that might have an impact on the live system without the risk of actually doing live testing.

Simulations of maintenance and repair actions can help discover unforeseen program vulnerabilities and guide optimum cost-effective scheduling. Facilities engineers can review powering schemes for reliability and capacity. IT managers, concerned with availability and service level agreements, can explore dynamic application loading scenarios in a virtual environment without the risk of unintended downtime.

The real-time capability can intelligently predict the timing and location of potential system upsets, and, in the case of a downtime episode,

can quickly apprise the right people as to the cause and solution. Since time is money, reducing overall downtime by as little as six minutes per year can mean a potential savings of about $100,000 if downtime is worth $1 million per hour.

Power analytics can formulate truly predictive diagnostics based on system design boundaries, and the implications of variable operating conditions from system aging. Intelligently scheduled system maintenance or repair based on a reliability assessment rather than a simple periodic basis can be less upsetting and costly.

**A closer look at Power Analytics**

When deployed, users will find that the power and uniqueness of the analytics platform is derived from the complete encoding of the design specifications from the original, as-built power infrastructure. All power system electrical parameters are calculated from the stored design specifications and, during the data center's normal operation, compared with the real-time power data. At any time, power analytics can accurately corroborate as-specified power parameters, determine if there are system anomalies and predict when and where there are potential vulnerabilities for system and equipment failure.

Furthermore, advanced features allow users to capture current system state data and run detailed "what if" simulations to verify system operations for the data center commissioning process, as well as investigate the effects of equipment rearrangement, configuration modifications, capacity expansion and other data room modifications might have on the live system without the risk of actually doing live testing. Simulations of maintenance and repair actions can help discover unforeseen program vulnerabilities and guide optimum cost-effective scheduling. Facilities engineers can review powering schemes for reliability and capacity. IT managers, concerned with availability and service level agreements, can explore dynamic application loading scenarios in a virtual environment without the risk of unintended downtime.

**Turning Red Ink "Green"**

Key to an effective energy management program is accurate information regarding the consumption of energy. Based on the amount of IT equipment in racks, the power distribution and cooling equipment infrastructure and the variations in application loading, power analytics can report accurate, real-time energy usage. This data can be compared to the "as-designed" energy usage to give insight into system unbalances, capacity restraints or overloads. The results of virtualization and other energy efficiency measures can be followed and assimilated. Power analytics can suggest scenarios for improved energy utilization based on its predictive diagnostics ability and by "what-if" simulation. At the current energy costs (~$0.089 kWh), a nominal realized annual savings of ten percent for even a relatively small, lightly-loaded data center is significant — greater than $100,000.

The added complexity of energy management will increasingly drive system and financial decisions. Power analytics addresses the continuum of energy management, from availability and performance to reliability and quality; a timely and powerful solution for the greener and more reliable enterprise.

**References**

[1] The Green Grid, "The Green Grid Data Center Power Efficiency Metrics: PUE and DCiE" December 2007

[2] EPA, "Report to Congress on Server and Data Center Energy Efficiency", August 2, 2007

[3] Emission Database for Global Atmospheric Research version 3.2, Fast Track 2000 Project. This value is intended to provide a snapshot of global annual greenhouse gas emissions in the year 2000

[4] Executive Director Report: The Findings of the 2007 Charrette, Kenneth G. Brill, DESIGN CHARRETTE 2007, Data Center Energy Efficiency By Design

*Barry J. Needle is the principal engineering consultant for AGO LLC. He can be reached at bjneedle@gmail.com.*