Energy Management Driven By
Power Analytics™

AGO Research

The goals of the "greening" of the Data Center are to minimize energy consumption and reduce the emission of green house gases (carbon footprint) while maximizing IT performance. Energy efficiency metrics…applied in real-time can be used to shape energy conservation strategies, and to determine the effectiveness of the myriad measures to reduce energy consumption. Combining the Power Analytics platform with a virtualization environment system such as the EDSA Blackboard Simulator is the means to understand the unforeseen risks to performance accompanying a system change, and … to preview the total costs of operating the facility, economically and environmentally.

Introduction

Energy management is the hot topic of the ongoing discussion of the "greening" of the Data Center. White papers, articles, and symposia presentations have focused on the critical issues of upward spiraling energy consumption and production of greenhouse gases resulting from the proliferation of energy inefficient information factories. In the past, the energy and environmental costs of running an ICT (Information & Communication Technology) facility was not the main concern of the info-center manager - meeting Service Level Agreements was. Now, decision makers have to consider a holistic approach[1,2] i.e., profitability via maximizing business metrics for service performance and availability while minimizing the environmental (energy consumption and carbon footprint) costs of operating an ICT facility.

Holistic Energy Management and Metrics

The Green Grid organization, Uptime Institute, and McKinsey & Co have proposed frameworks with accompanying quantifying metrics for identifying opportunities for improving energy efficiency, and for providing top management with strategic guidelines on minimizing the impact of energy costs on their businesses. The frameworks present usable models of data center energy consumption upon which to base strategic planning, specific examples of how best practices can reduce a data center's energy consumption, and a set of metrics that be applied to plan and manage energy use in the data center.

[1] EPA, "Report to Congress on Server and Data Center Energy Efficiency", August 2, 2007
[2] Simon Mingay, Gartner, 2007, Green IT: A New Industry Shock Wave, [http://nl.sun.com/oversun/evenementen/superpowered/Gartner_on_Green_IT.pdf]

The Uptime Institute's whole-systems thinking model[3], with accompanying quantifying metrics, poses strategic categories to be addressed by ICT facilities owners which are,

- Relevant for understanding the growing energy use in a facility as well as in the industry, and the energy impact of systems architecture and platform selection, and network design.
- Valuable in characterizing the productivity and utilization of IT hardware assets.
- Helpful in selecting and justifying the purchase of equipment which delivers the most effective performance / watt at the plug.
- Important to maximizing the quantity of useful power/energy to the IT hardware for each unit of power/energy measured at the facility utility meter to reduce the amount of energy used to support the IT 'work" of a data center.
- Essential for understanding the total capital expenditures (CAPEX) required to accommodate the designed power capacity of the data center, and for verifying operational costs (OPEX) based on the total energy profile of the facility.

Measuring Energy Efficiency – the Metrics and More

About five years ago, the need to define (and measure) energy efficiency was recognized as it became apparent that the increasing amount of energy necessary to power and to cool the engines of ICT was not just as a technical, operational nuisance, but a potential threat to business profitability. The energy consumption of ICT was increasing at a prodigious rate due to a geometric rise in compute equipment performance [factor of three every two years] without an accompanying increase in effective energy efficiency [doubling every two years].

Energy efficiency metrics are mathematical representations of a dictionary definition of efficiency. Efficiency is "the ratio of the useful energy delivered by a dynamic system to the energy supplied to it." In the case of a data center, the easiest energy flow to measure between the systems is electricity.

In 2006, a number of foundational metrics were published.

- SI-EER (**S**ite **I**nfrastructure-**E**nergy **E**fficiency **R**atio), IT-EER (**I**nformation **T**echnology-**E**nergy **E**fficiency **R**atio), and DC-EER (**D**ata **C**enter-**E**nergy **E**fficiency **R**atio) by the Uptime Institute[4]
- PUE (**P**ower **U**sage **E**ffectiveness) by C. Malone and C. Belady[5]

Since then, Green Grid, Uptime Institute, and McKinsey & Co have expanded the inventory of metrics as "effectiveness" or "utilization" ratios have been created.

---

[3] Uptime Institute,2007, *Four Metrics Define Data Center "Greenness"*
[4] Uptime Institute, 2006, *High Density Computing: The Path Forward 2006*
[5] Malone, C., Belady, C., (2006), *Metrics to Characterize Data Center & IT Equipment Energy Use,* Proceedings of 2006 Digital Power Forum, Richardson, TX

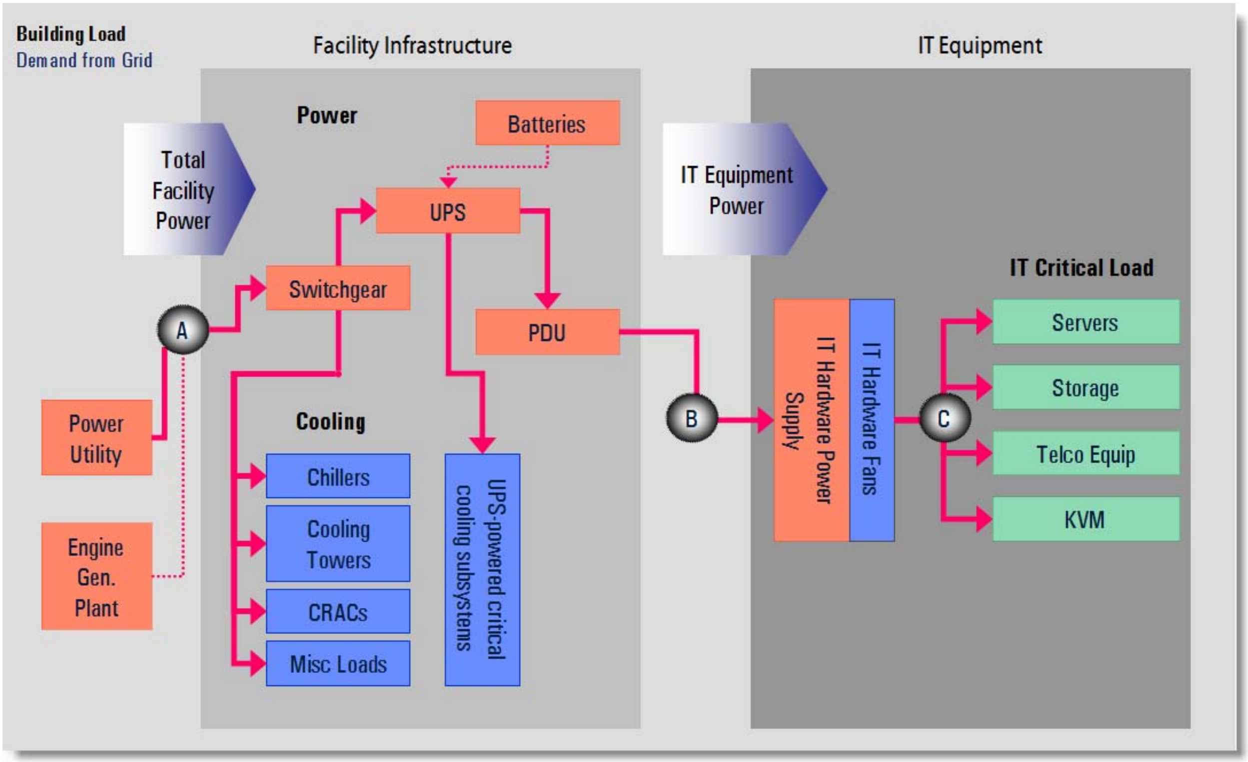

Figure 1: Data Center Electricity Flow and Metric Measurement Points

Figure 1 shows the electricity flow between the power, cooling, and compute systems of a typical data center and the basic points of power measurement. Total data center energy consumption is measured at **Point A** – the utility meter. **Point B** is downstream from the Power Distribution Units at the input to the IT Equipment load. Power measurements at **Point C** taken downstream from the equipment power supplies are typically Direct Current, represent the compute load.

Green Grid

Metric	Definition
DCiE [Data Center Infrastructure Efficiency]	IT Equipment Power / Total Facility Power *100% *[Point **B** / Point **A**]*
PUE [Power Usage Effectiveness]	1 / DCiE *[Point **A** / Point **B**]*
CPE [Compute Power Efficiency]	IT Equipment Utilization * DCiE
UDC [Data Center Utilization]	IT Equipment Power / Actual power capacity of the data center
U_{server} [Server Utilization]	Activity of the server's processor / Maximum ability in the highest frequency state

Metric	Definition
$U_{storage}$ [Storage Utilization]	Percent storage used / Total storage capacity of data center
$U_{network}$ [Network Utilization]	Percent network bandwidth used / Total bandwidth capacity of data center

Uptime Institute

Metric	Definition
SI-POM [Site Infrastructure Power Overhead Multiplier]	Data center power consumption at the utility meter / Total hardware AC power consumption at the plug for all IT equipment
H-POM [IT Hardware Power Overhead Multiplier]	AC Hardware Load at the plug / DC Hardware Compute Load *[Point **C** / Point **B**]*
$DH\text{-}UR_{servers}$ [Deployed Hardware Utilization Ratio]	Number of servers running live applications / Total number of servers actually deployed
$DH\text{-}UR_{storage}$ [Deployed Hardware Utilization Ratio]	Number of terabytes of storage holding important, frequently accessed data (within last 90 days) / Total terabytes of storage actually deployed
$DH\text{-}UE_{servers}$ [Deployed Hardware Utilization Efficiency]	Minimum number of servers necessary to handle peak compute load / Total number of servers deployed

McKinsey & Co

CADE [Corporate Average Data Efficiency] = **Facility Efficiency * IT Asset Efficiency**

Where,

- **Facility Efficiency** = Facility Energy Efficiency (%) * Facility Utilization (%)
- **IT Asset Efficiency** = IT Utilization (%) * IT Energy Efficiency (%)

And,

Metric	Definition
Facility Energy Efficiency (%)	Actual IT load / Total power consumed by the data center *[Point **B** / Point **A**]*
Facility Utilization (%)	Actual IT load (servers, storage, network equipment) used / Facility capacity
IT Utilization (%)	Average CPU utilization
IT Energy Efficiency (%)	CPU Loading / Total CPU power

Subsequent metric development has yielded variations on the same efficiency theme, that of minimizing the expense of providing services, and has focused attention on productivity. While related to efficiency, data center productivity is maximizing the work of providing services for a given consumption of a resource. When the resource is energy the metric is Data Center Energy Productivity, a notion proposed by the Green Grid organization, which is expressed as,

DCeP *= Useful IT work produced /*
Total Data center Energy Consumed Producing the Work

Where the valued IT "work" produced by a data center is related to the amount of energy it uses…in essence the holistic mantra, get the most IT "work" for the least energy (environmental) cost." While this productivity metric is tantalizing simple to express, the value for IT "work" still needs development before it can be used effectively.

Real-Time Energy Efficiency Measurement – Power Analytics Applied

The fundamental energy efficiency of a data center is revealed when the electrical parameters of the power measured at the input of the ICT equipment is proportioned to those measured at the utility meter recording the total energy used to power the ICT infrastructure. (See Figure 1) This value, **DCiE**, **PUE**, **SI-POM**, or **Facility Energy Efficiency** (all equivalent), indicates how much power is used by the facility infrastructure to power and cool the IT equipment, and to power the redundant distribution systems required for maintaining the expected availability and reliability of the "info-factory's" services. The **DCiE** metric guides and reports on the effectiveness of any program to reduce power distribution losses, cooling load, and IT loading, and enables straightforward benchmarking comparison between data centers.

The Green Grid concludes that to gain real insight and to successfully manage an IT facility's energy efficiency continuous real-time monitoring should be employed to do the historical trending and statistical analysis, this is especially important to understand in light of the dynamic nature of an ICT facility.

"…the initial design of the data center is obsolete the day after the installation and commissioning is complete. Often, energy efficiency calculations are based on a static

design instead of the dynamic data center configuration. The designed (static) versus actual (dynamic) nature of a data center must be considered. Improvements will come through incremental step changes in infrastructure over time. Also, it is important to keep in mind that as the load changes in the data center, the operating point of the subcomponents on their efficiency curves will change." [Green Grid, 2008]

Recently, a software technology platform called Power Analytics has made it possible to have true insight into the operational variations between the designed versus actual nature of a data center. Power Analytics is driven by sophisticated mathematical modeling which synthesizes the performance specifications of all the electrical interconnections and equipment of the electrical design model (CAD model) of the data center into a digital image of the actual facility. The power and uniqueness of the platform is derived from this complete encoding of the design specifications from the original, as-built power infrastructure into the model "designbase".

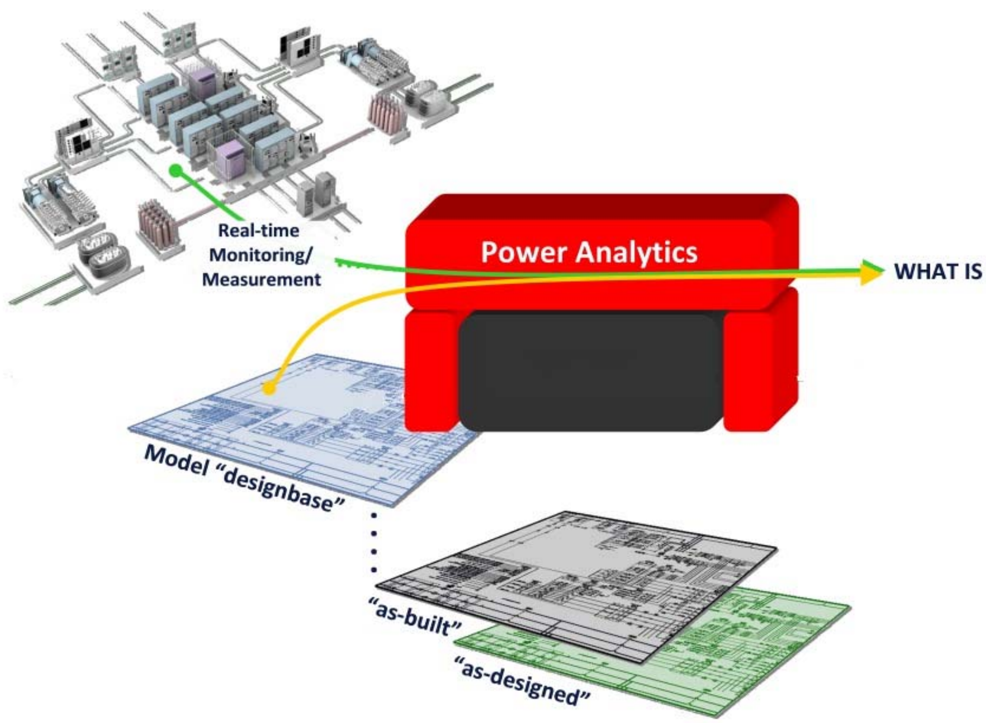

Figure 2. Comparison of Real-time Operating Parameters with Model "Designbase"

During a data center's normal operation, Power Analytics dynamically and continuously compares the encoded model, "designbase," of the facility with the operating parameters and interoperability of each piece of hardware, from the main switchgear, through the facility's power and cooling infrastructure, to the ICT critical compute load. (See Figure 2) Power Analytics can accurately corroborate the operating parameters of a "live" component with their "as-designed", or other model specifications.

A successful energy management program is built upon having accurate information regarding the consumption of energy - ***where and when it is consumed, and the operational and environmental costs.*** Power Analytics can report accurate, real-time energy usage based on the amount of ICT equipment in racks, the capacity and redundancy of the power distribution and cooling equipment infrastructure, and the variations in compute loading, Further, this data can be compared to the "as-designed" energy usage calculated by the analytic system to give insight into system instability, unbalances, capacity restraints, or overloads.

Is this capability essential for an effective energy efficiency program? Consider the following.

- Simply determining DCiE weekly or monthly will provide the data to show the effectiveness of measures to improve efficiency. However, what if server consolidation / virtualization, changes to the cooling / power infrastructure, or other energy efficiency measures result in the instability of the data center?
- By increasing the granularity of the energy measurements, a hierarchy of subsystem power losses can be revealed. Correlated with the efficiency metrics, a ranking of a contributor's power loss by magnitude can guide changes to lower the energy consumption of the subsystems. How can this be done and still support the requirements of the IT equipment, and meet service availability?
- Can the energy efficiency status of a multiuse facility be determined accurately?
- How can the environmental benefit and energy cost impact of efficiency metrics be determined?
- Can new IT system designs, network architecture, and equipment options, developed to enhanced the availability, reliability, and productivity of information factories, be selected and justified to improve business performance goals while using less energy?
- How can the productivity and utilization of IT assets be maximized?

Actual energy costs have rarely been incorporated into ICT availability management because the primary mission of the facility was availability. While the application of rate structures, utility penalties and overall energy cost has their roots in more traditional energy billing and allocations the application of the technology was almost never included in the critical space. Previous applications of actual energy costs were a function of energy billing systems, residential users and mix use campuses. The Power Analytics platform is the ideal architecture for Energy Management in the ICT or critical space because it does so without compromising the primary mission of availability management while addressing the important costs issues.

- What is my actual cost for energy for the critical space based on my billing structure?
- How does that cost fluctuate on an hourly, daily, monthly or annual basis?
- How is energy cost impacted by SLA agreements?
- How will changes we are planning affect energy cost *and* availability?

The Power Analytics platform answers these questions and more by combining availability management and energy management via ICT Energy Management and the Power Analytics Blackboard.

ICT Energy Management and The Power Analytics Blackboard

The Power Analytics Blackboard, a module of the Power Analytics software platform, functions by capturing a "snapshot" of the current ICT facility's input power parameter settings to drive the model "designbase" in simulation of the operation of the encoded data center. By altering selected model specifications, e.g., circuit breaker settings, the impact of real-time, "what if" system changes (server loading, power distribution and cooling, capacity shifts, equipment substitution, etc.) can be projected and analyzed. The facility's current performance and costs, i.e., "*what is,*" can be compared to performance and costs based on "*what if*."

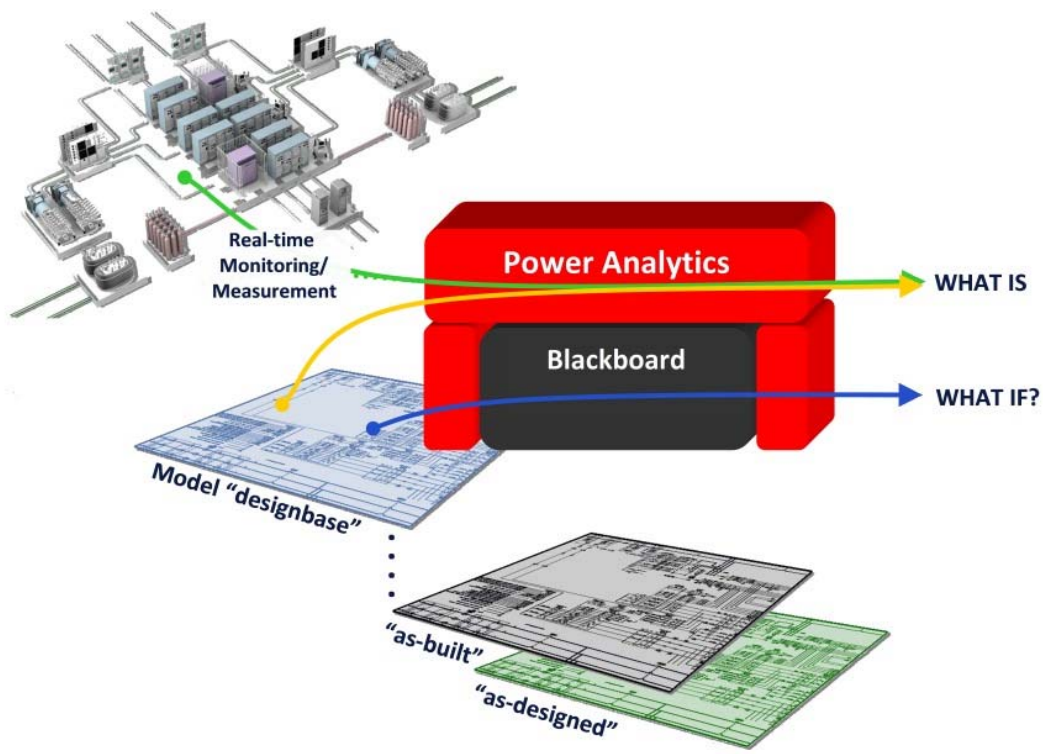

Figure 3. Blackboard - "What-If" simulation using model parameters

Virtualization, server consolidation, decommissioning, improved air flow and cooling efficiency, more efficient power distribution, rationalized capacity, and other energy management measures involve the risk of unintended consequences. The simulation of a system's performance in a virtual environment can safely investigate the effects of changes that might have an impact on the live system without the risk of actually doing live testing. Moreover, during the "what if" simulation of an energy reduction scheme the data center's performance can be watched to make sure the IT equipment is supported, and system availability and reliability are maintained.

Energy efficiency metrics are valuable in revealing the efficacy and ranking of measures taken to reduce energy use. How, as mentioned above, can cost benefits be estimated? If real-time utility costs (energy and power demand rates) are recorded while Power Analytics technology scrutinizes the data center performance, costs at a measured performance state can be assessed. The utility costs verified by local rate table calculation can be correlated with "what if" performance simulations via Blackboard. (see Figure 4)

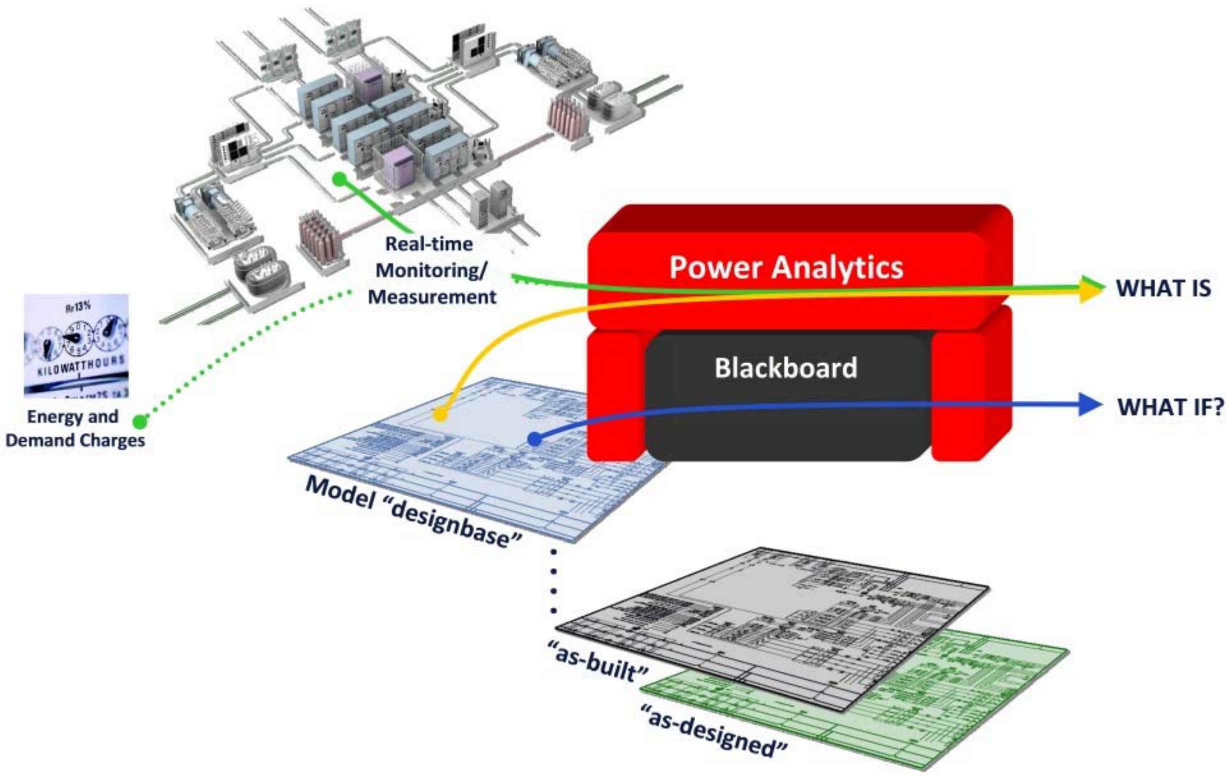

Figure 4. – Energy Use Costs correlated with Blackboard simulation

Furthermore, if Blackboard simulations are run with equipment costs correlated with power measurements taken at the individual equipment level, then the cost impact of energy efficiency measures made at that level can be determined. Bottom line: the economic benefit of virtualization, server consolidation, eliminating unproductive IT assets, reduced cooling load, more efficient power distribution, and other changes can be assessed before actually implementing the change.

Armed with the power of Blackboard simulation and Power Analytics, the energy efficiency of new ICT system designs can be explored, equipment delivering the most efficient performance can be selected and cost justified, and as productivity metrics are refined, the total cost of operating the data center can understood and verified.

Power Analytics - driven Energy Management…

The goals of the "greening" of the Data Center are to minimize energy consumption and reduce the emission of green house gases (carbon footprint) while maximizing IT performance. Energy efficiency metrics developed and proposed by noted industry organizations applied in real-time can be used to shape energy conservation strategies, and to determine the effectiveness of the myriad measures to reduce energy consumption… to increase energy efficiency in ICT facilities. The Power Analytics platform used with a virtualization environment such as the EDSA Blackboard Simulator is the means to understand the unforeseen risks to performance accompanying system changes, and by correlating the energy consumption of the data center with the production of green house gases, preview the total costs of operating the facility, economically and environmentally.

Energy Management "By Design"