



Going Green in Data Centers

Organizations should work together to coordinate an all-encompassing effort that will address data center energy efficiency so that important ideas are not overlooked.

BY WILLIAM J. KOSIK, PE, CEM, LEED AP, *Managing Principal, EYP Mission Critical Facilities Inc., Chicago*

The August release of a study by the U.S. Environmental Protection Agency (EPA) on data center energy efficiency adds fuel to the fire in the research and development of new ways to reduce energy use. The findings on the energy use of data centers, summarized at www.energystar.gov/datacenters, are staggering:

- Data centers consumed about 60 billion kWh in 2006, roughly 1.5% of total U.S. electricity consumption.
- The energy consumption of servers and data centers has doubled in the past five years and is expected to nearly double again in the next five years to more than 100 billion kWh, costing about \$7.4 billion annually.
- Federal servers and data centers account for approximately 6 billion kWh (10%) of this electricity use, at a total electricity cost of about \$450 million per year.
- Existing technologies and strategies could reduce typical server energy use by an estimated 25%, with even greater energy savings possible with advanced technologies.

Several organizations are working to reduce energy consumption in data centers. Capitalizing on this momentum is critical to the success of this endeavor. However, it is equally critical that there is a coordinated, all-encompassing effort that will address the different areas so that important ideas are not overlooked.

This effort has similarities to the U.S. Green Building Council's (USGBC) LEED rating system. The rating system was developed and is maintained "via a robust consensus process that has been refined over more than a decade of leadership experience. The key elements of the USGBC's consensus process include a balanced and transparent committee structure; Technical Advisory Groups to ensure scientific consistency and rigor; opportunities for stakeholder comment and review; member ballot of new rating systems and certain changes to existing rating systems; and a fair and open appeals process."

Using LEED as a model, it will be possible to dovetail in the following initiatives that are already underway (these are just a sampling of

current initiatives—the EPA study provides comprehensive information on others):

- The Green Grid already has started down the path in measuring, benchmarking and developing best practices for energy-efficient data centers.
- Lawrence Berkeley National Labs is developing more robust and user-friendly assessment tools to identify areas of improvement in data center energy use.
- Major computer manufacturers publicly have announced various programs that achieve energy savings. These programs are available to consumers.
- The EPA is currently exploring “whole building energy use data for stand-alone mission-critical facilities, with the goal of creating a building-level metric.” The first portion of this work is expected at the end of this year.
- The Uptime Institute also published its opinion on the EPA report. In addition to other key areas of study that are suggested, it proposes that a green data center have four separate performance factors:
 - IT hardware productivity
 - Maximum computational performance per unit of internal power
 - Efficient delivery of power at the plug to IT hardware components
 - Efficient site infrastructure.

The missing link

While either tangentially touched upon or implied in the initiatives, there is little discussion on the environmental impacts that are unique to data centers. In general, in addition to consuming tremendous amounts of electricity, large enterprise data centers have the following attributes that make them ideal candidates for “greening”:

- Typically one-story, large-footprint buildings
- For security reasons, often built on undeveloped, “greenfield” sites
- Often built in areas where electricity costs are low
- Typically sited away from public transportation
- Can consume vast quantities of water for cooling
- Can produce large amounts of water discharge into the municipal sewer system
 - Use large amounts of chemical for water treatment
 - Require on-site storage of fuel for generators
 - Require on-site storage of water for cooling
 - Designed with multiple levels of redundancy in systems and equipment
 - Very high initial and ongoing operational cost
 - Need to be flexible and scalable without any shutdown or interruption to operations
 - Have massive amounts of on-site, diesel engine generators used only in case of utility failure
 - Use 100% outside air for battery room ventilation
 - Use large banks of batteries to ride through transfers to generator power
 - Often will have “mirror” sites with duplicate facilities, each acting as

a backup to the other in case of a major regional catastrophe

- Have rigorous maintenance, operation and change management procedures
- Typically not many people are in the building on a permanent basis
- Buildings are internally loaded where cooling is required non-stop
- High concentrations of IT hardware and cabling where off-gassing of VOCs and other indoor pollutants may be of concern
- Business needs will often dictate amounts and types of computer equipment, with a continual cycle of upgrade and/or replacement of relatively new equipment
- Power use monitoring typically is done at a very granular level, down to the circuit level in some cases
- Monitoring, control, trending and warning systems tend to be very robust and sophisticated
- Design, construction and commissioning schedules are extremely aggressive
- Projects are often highly confidential and information regarding technical details of the facility are not allowed to be disclosed
- Require very strict indoor temperature, humidity and filtration levels to maintain an acceptable environment for the IT equipment.

We need a program that uses LEED as a starting point to identify data centers as a specific building type, addressing not only the nuances that make data centers difficult to achieve LEED certification, but also address the opportunities that are exclusively found in data centers.

Framework for analysis

The culmination of energy-reduction ideas could result in a rating system that takes a broad and deep look at all of the elements that affect the outcome and attempt to distill the ideas into a clear, objective assessment, not unlike the U.S. Green Building Council’s LEED rating system. If the Energy Optimization Scorecard approach is taken, it does the following:

- Addresses both financial and non-financial elements from a variety of perspectives, in an objective and unbiased fashion
- Provides information to corporate leadership to assist strategic policy formation and achievement
- Develops a set of measures that provides a fast but comprehensive view of the strategy and measurements
- Provides an optimized picture of overall performance highlighting areas that need to be improved or augmented based on anticipated changes in future state
 - Assists in clarifying vision and strategies and provides a means to translate these into action
 - Provides a comprehensive view that overturns the traditional approach of collecting and analyzing data as isolated, independent functions.

Outside the box

It's essential to broaden the topic beyond data centers. This is an area of focus that is dependent on the actual site of the data center. Each type of electric generation technology uses different types and amounts of fuel (natural gas, coal, oil) and each power producer uses varying types of renewable power generation technology, such as wind and solar. CO₂ emissions also need to be taken into consideration as potential sites are evaluated.

There are many other factors beyond the commercial building sector that produce global warming gases. The current strategy for long-term stabilization by 2055 comes in the form of wedges of reduction, where one wedge equals 1 billion tons of CO₂ (see Figure 1).

One of the wedge strategies comes from Princeton University's Carbon Mitigation Initiative. According to this template, the commercial building sector needs to reduce emissions from buildings and appliances by 25% by 2055. This will be particularly difficult for data centers as the demand for technology increases. However, this reduction is consistent with the findings of the EPA study that energy reductions ranging from 30% to 80% are achievable with varying degrees of investment.

An important environmental consideration rarely discussed is the amount of water that data centers consume in the cooling process. In addition to the water consumed, it is important to understand the amount of energy that is used in the overall water supply and discharge cycle for both the data center and the utility that is producing the electricity.

Balancing the relationship between water use and energy is a complex discussion (see Figure 2), but it is important to understand that, depending on the location of the data center, the strategy to use more water to offset electricity use—such as in an evaporative cooled air-conditioning system—may not be applicable.

For instance, according to the California Energy Commission's 2005 report, "California's Water-Energy Relationship," water-related energy consumes 19% of the state's electricity, 30% of its natural gas and 88 billion gallons of diesel fuel every year. Reducing the amount of water consumed also will reduce the amount of energy. So in addition to general water scarcity concerns, the interrelationship between water use and electrical consumption will impact the decision-making process on the overall strategy for data center planning.

Because the industry is poised to take a great leap forward down the path of addressing these disparate subjects, one logical thought is to try to achieve some type of balance or state of equilibrium.

Strictly defined, balance is "a state in which opposing forces or factors are of equal strength or importance so that they effectively cancel each other out and stability is maintained." However, the forces that come into play when evaluating reliability, cost, regard for the environment and maximizing the capability of IT systems are rarely equal—having the ability to stabilize them through achieving balance is not very practical.

Perhaps there is a better term—such as optimize—that helps define and shape the current efforts. Optimize is defined as "to find the best possible solution to a technical problem in which there are a number of competing or conflicting considerations, where nothing is out of proportion or unduly emphasized at the expense of the rest." This definition is undoubtedly a much better representation of the outcome that is envisioned.

So with all of the compelling arguments, pros and cons, and business impacts, it seems to be a daunting task to try to construct a decision-making envelope that contains all of the necessary criteria, standards, strategies, outcomes, technologies, and measurements required to enable an objective, evidence-based decision-making process.




Figure 1 Long-term stabilization comes in the form of wedges of reduction, where one wedge equals 1 billion tons of CO₂. The commercial building sector must reduce emissions from buildings and appliances by 25% by 2055; this will be tough for data centers as technology demands increase. Source: Princeton University's Carbon Mitigation Initiative

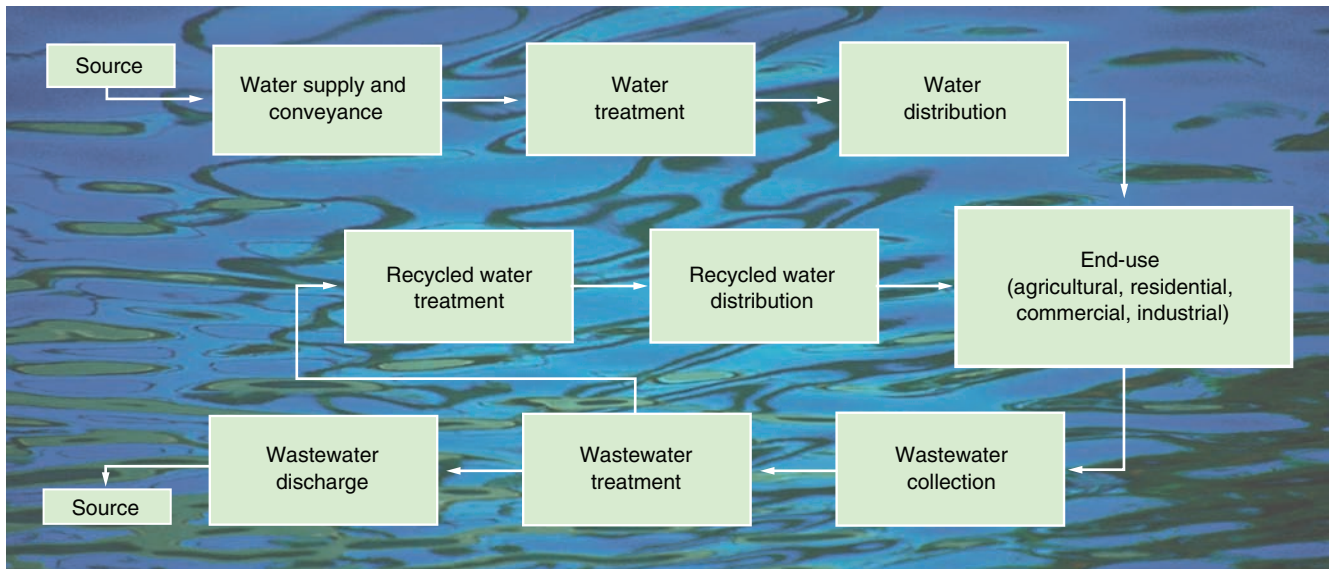


Figure 2 The interrelationship between water use and electrical consumption will impact the decision-making process on the overall strategy for data center planning. Source: California's Water–Energy Relationship, Final Staff Report

If one looks to the Balance Scorecard method produced in the 1990s by the Harvard Business School, they can glean some ideas on how to approach this complex problem. Because the focus will be on optimization, the analysis will be termed Energy Optimization Scorecard. If all of the issues can be aggregated into four primary, over-arching tenets that apply to data centers, developing an optimized a solution by analyzing and comparing the relative benefits and liabilities seems much more possible.

Four fundamentals

There are four fundamental beliefs that affect data center planning:

- **Operational continuity.** One of the most unique aspects of any building type, data centers and other types of mission-critical facilities have zero tolerance for failure or shutdown due to maintenance. While there are varying degrees of required system availability, this is an area that generally has very little ability for compromise. The design, operation and integration of the facilities systems (power and cooling) and the IT systems (network and computer hardware) is where the reliability can be influenced the greatest.

- **Lifecycle cost.** Not very different from other building types, getting the biggest bang for the buck is important. An important difference that must be recognized, however, is that utility costs for a data center become a very high percentage of the overall operational cost. Also, if alternatives are developed that reduce first cost by reducing reliability, the financial and reputation risks of a failure need to be accounted for in any lifecycle costing analysis.

- **Environmental impact.** Knowing that a data center on a square-foot basis will produce on the magnitude of 25 times more CO₂ emissions than a typical commercial office building and will use 25 times

more water, it is imperative to minimize the negative environmental impacts that are a byproduct of building and operating a data center. Also, understanding how the utility power that is feeding the data center is being generated also will have a big impact on the environment.

- **IT effectiveness.** The other component to be considered is the ratio of power that is dedicated to the computer equipment. This ratio, commonly known as the power usage effectiveness, is a measure of how much power is being provisioned for the IT operations and is affected by reliability level, climate, system type and equipment type. Another area of efficiency that needs to be explored is specific power consumption of the computers vs. the computational output of the computer. This metric is used in the supercomputing field in the form of watts per FLOP (floating-point operations) where there are standard performance benchmarking tests. This is an area that still needs considerable discussion before a solid metric can be presented.

By using parametric analysis to determine the interdependencies between the areas of impact and the how they affect the items in the energy optimization categories, an overall strategy can be developed. Because this analysis will have quantifiable results and can be done in an iterative fashion, several “what if” scenarios can be explored and an informed decision can be made.

Even though the industry has taken a major step forward in looking at energy-efficient data centers, there still is a lot of work to do. Given the very specialized and critical nature of these facilities, the technical strategies for improving energy performance will prove far easier than the issues concerning operational continuity and other business drivers that are specific to each market sector. Nevertheless, developing and implementing a robust, objective framework for analysis and decision making will ensure a fully optimized outcome. **cse**