**WHITE PAPER:** OPERATIONS ENERGY MANAGEMENT

# Operations Energy Management: From the Data Center Through Facilities

OCTOBER 2008

Chris Stakutis
CA EMERGING TECHNOLOGIES

# Table of Contents

# Executive Summary

## Challenge

Energy costs around the world are causing facility and IT managers to pay close attention to their energy uses. A typical data center can use 30 times more energy per square foot than ordinary office space. Large facility equipment such as HVAC systems and chillers can account for the majority of power usage in some industries. Can technology be used to dramatically reduce energy usage and get this precious cost once and for all under control?

## Opportunity

Historically, most businesses were not overly concerned with the cost of energy because energy was abundant and inexpensive. This, of course, has changed today, and facility owners and managers are looking towards technology to help them get a grip on spiraling energy costs. Opportunities exist for innovative technology such as pervasive monitoring and wide-area controls integration to bring a new level of insight and control to operations energy management.

## Benefits

The benefits of this new approach to energy management are clearly cost and sustainability. With cost improvements, business can stay more competitive and profitable, especially in intensive energy use industries such as IT and retail. Just as important is improving energy usage as this is critical to the longevity of the world. The most prevalent energy source is harmful carbon based fuel. Reducing energy use is just good business all around, and technology can once again help to improve our world.

SECTION 1

# Energy Costs Are Out of Control

### Despite Great Strides in Lower-power Equipment, Power Use Is Doubling

Manufacturers from refrigeration systems to HVAC systems to computer chips and circuit boards have been diligent about increasing their efficiency, but no matter, our appetite for increased computational horsepower and facility comfort outpace these improvements. Specifically within today's data centers, power usage is growing at an alarming 20 percent rate, which means doubling the use within 5 years. Explosive growth of commercial outlets and the global reach of huge retail giants (coupled with Asia expansion) have created opportunities for energy management at an unprecedented scale. Even within a data center, the sheer number of elements consuming power is at a scale that the IT industry never anticipated.

It is clear that something must be done to rein in rampant power use, and the first step is to identify the main components of energy usage.

### Cooling Is a Major and Often Poorly Managed Consumer

Let's start with examining cooling within the data center. Upwards of 25 percent to one-third of all power used in the data center is for cooling. Despite equipment running more efficiently each year, the overall heat is growing because of the overall need for higher computational elements. As such, cooling costs will continue to rise accordingly.

However, today's typical data center is overcooled by approximately 10 degrees to 20 degrees Fahrenheit. Data center managers have been trained for generations to keep the data center operating at a chilly 60 degrees. This is completely unnecessary; all of the major computer and chip manufacturers have published agreed upon temperature guidelines that allow for air input temperatures to be close to 80 degrees Fahrenheit. In fact, the air input (to the front face of the equipment) can be has high as 90 degrees for fairly long periods (days) without adverse effect.

Why then does every data center run at 60 degrees? Partly because the industry has been improperly trained and largely because of fear of renegade air flows that could truly result in extreme and dangerous temperature levels. When energy was inexpensive, it was safe to simply overcool. Obviously, energy is no longer inexpensive, and a change in setpoint from 60 degrees to 70 degrees or 78 degrees could cut the cooling costs by 80 percent, or even eliminate it all together (with cooling through economizers because in many geographies the average air temperature is well below the equipment ratings).

Outside of the data center, cooling is a critical cost that is now found in many retail settings, such as those that involve chilled or frozen food. A typical store can have hundreds of refrigeration points and quite likely most are operating at a subpar level. When energy was inexpensive, there was not much need to closely monitor all of the operational characteristics of thousands or millions of elements across a retail chain. Today, there is a compelling case to closely manage and monitor such systems and avoid subpar operations.

FIGURE A

According to the US Environmental Protection Agency (EPA), the energy used by US domestic data centers doubled between 2000 and 2006 and, assuming historical trends continue, is expected to double again by 2011.

ENERGY USED BY US DATA CENTERS WILL DOUBLE AGAIN BETWEEN 2007 AND 2011

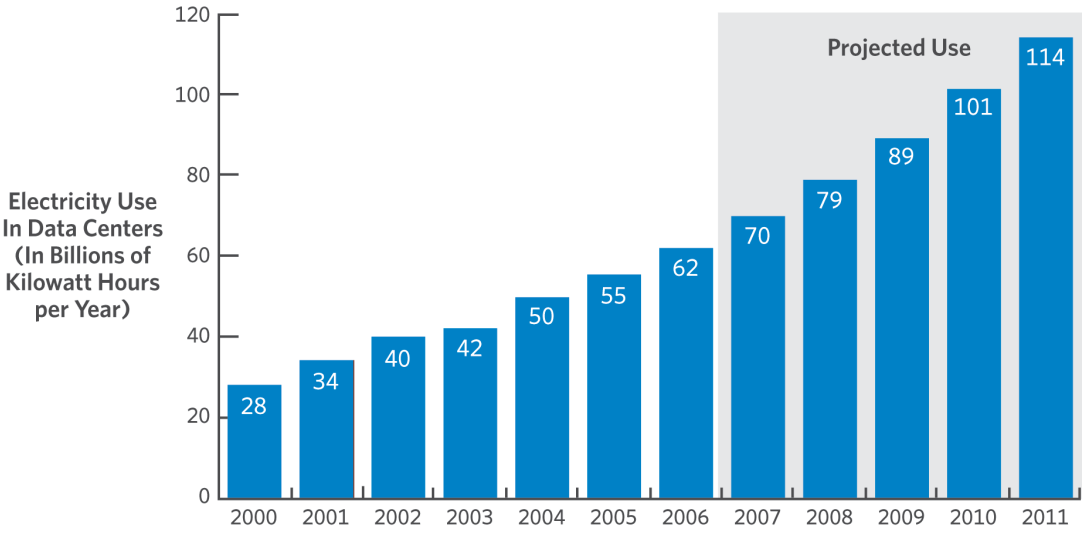


Source: US Environmental Protection agency, Yankee Group Research, Inc., Green Software Can Curb IT's Energy Crisis, September 17, 2008

## Ghost Servers and Underutilization Are the Culprits

Servers consume 50 percent of the power within data centers, and if the cooling component is dramatically reduced or eliminated, the server power use would be closer to 75 percent. In general, this is good — servers should be the main use of power within the data center since they are providing the work function. However, there is huge waste in computational horsepower making this 75 percent energy user a prominent foe.

Close to one-fifth of the servers in a data center are no longer used and should be removed. Often this is a result of decommissioning that somehow misses the final step of power-off and removal or an elongated decommissioning process that has no definite end. IT owners of such systems move on to other capacities and forget about defunct application processes, or owners believe that someday they will rediscover a purpose for this equipment.

As alarming of a problem that those ghost servers cause, it is pale to a more stunning statistic, which is that 85 percent to 87 percent of all computing horsepower in a data center is not used. Data centers are far over provisioned than previously thought, and there are a number of causes for this. First, application owners add more powerful servers without fully releasing the originals. Second, general paranoia about headroom for end-of-quarter or holiday seasons causes a significant over buy. Third, IT and application owners often plan for future uses that do not occur or occur too far in the future when jobs or technology have changed again. The result is 12 percent to 15 percent computer utilization.

Doing the math, 85 percent waste of 75 percent of your power and energy spending means that about 50 percent of your energy costs in the data center are pure waste and could be completely avoided. Said differently, you could double your computing needs with a zero change in energy cost if this cost was better controlled.

SECTION 2

# Technology Empowering the Manager

## First: Instrument

We are a society that is fascinated by data and statistics, and the Web and Internet have shown us the power of collecting massive data and understanding the information therein. We can now collect data about power usage to a level sufficient for truly understanding our energy use and hopefully mapping that back to business value for decision making.

Although data centers have always been physically large, they have not historically had the number of components that we find there today. Density is king and whereas in the past a row might be a single mainframe, today each rack in a row is capable of holding hundreds of blade servers and dozens of other power consumers. A typical data center now has over a thousand circuits and each circuit can fan out to a dozen or more outlets within the rack.

While it may seem daunting to have potentially 10,000 power sensors reporting on current and energy draw, intelligent software can roll up views to give the IT director a fighting chance at responding to some very basic questions. Some of these include: Which side of my data center is drawing the most power? Which row is drawing the least? How does my power use map back to the business applications running on those elements?

Today the typical data center manager does not even have the single total power metric of his data center, let alone any insight into the components or sections that are drawing the power. This sort of information is essential in the modern world because energy costs are so high, and so much attention now needs to be placed on efficiency and resource utilization. The very first step in any operations energy management task is to measure where you are, and to the finest level available. Thankfully, the power-metering technology is now here and is extremely inexpensive and very easy to install non-disruptively.

But don't stop at the data center. It is extraordinarily easy today to also monitor building feed power and natural gas consumption and water usage. Even your backup generators (which can burn a tremendous amount of diesel when activated) can be integrated and monitored. Equipment from Honeywell (Tridium), Opto-22, RDM, and many others can take these very standard electrical pulses and signals and convert them to meaningful data that can be integrated into your monitoring tools. So, in the intense energy areas such as your data center, monitor power-down to the circuit or outlet level, but in a macro sense, monitor overall building energy consumers as well.

## Second: Visualize and Monitor

Now that you have your data center instrumented, the next step is to take that data into a modern tool that can help you easily visualize it. What sort of data and views are essential?

- Total carbon footprint (all properties and all elements summed up)
- Per facility roll-up totals
- Per large consumer (for example, data center) roll-up
- Within the large consumers, a hierarchy detail that can be mapped back to business processes and values.

While your monitoring and presentation software will not be as heavily used as say Microsoft Word or PowerPoint, it should still be intuitive, very graphical, and quick to see summary information. It should also be quick to drill to find out problem areas. The software needs to be useful to a CIO, not just to a lower-level facilities person, and CIO's want quick answers, useful views and data, and with hardly any effort.

The more powerful the software, the higher the utility to the IT manager and the director. Key is being able to drill down through various layers of macro-power usage, all the way down to not just a rack (within a row), nor just a server within a rack, but down to a virtual server within a server and then the applications and business processes running therein. That is the holy grail — deep business-mapped understanding of power usage.

### Third: Make Changes

With instrumentation in place and presentations showing you views of your energy consumption, you are now ready to make changes and track the impacts. Within a data center, the most important and typical changes that should be considered are:

#### HVAC AND COOLING

- Hot and Cold isles Too many data centers today have not fully embraced the hot and cold isle configuration. It is not efficient to try to get uniform cold temperature throughout the data center. It is far more efficient to blow the chilled air only in front-facing server rows and have the rows oppose one another (thus having a hot isle in the middle). With some HVAC minor rework or curtains, you should be able to have less spillover air and higher efficiency. Be sure to have ample temperature probes, especially at the row-ends and upper edges.
- Run warmer If you have decent hot and cold isle arrangements, you should now be able to run your data center potentially 10 degrees to 15 degrees warmer. The input air to the face of each server should be at 78 degrees or less (but no reason for it to be 70 degrees or 65 degrees or 60 degrees). The air coming out of the back can be 20 degrees warmer (100 degrees potentially) and this is acceptable as long as you are carefully monitoring to know that you are not exceeding threshold in any corner cases.
- Blanking panels Your racks should appear to be 100 percent full from the front. Any unused slots must have blanking panels, otherwise the airflow will not be efficient and you may not be pulling the right amount of air over the real elements.
- Curtains Consider installing curtains in the data center to help guide the airflow. These are quite inexpensive and easy to set up.
- Introduce variable speed drives on all blower, HVAC, CRAC, and CRAH systems.
- Economizers Introduce air economizers and attempt to run the majority of hours on outside air.

#### SERVERS

- Ghost servers Locate your ghost servers and immediately remove and decommission.
- Virtualize Compact all those 12 percent utilized servers onto single physical machines to achieve dramatic energy savings. Virtualization is a very straightforward and fairly simple step to help rein in much of the server CPU waste that is endemic. However, one challenge is that lightly-used images sometimes come alive and need ample headroom for such intervals.

- **Automation** Run Data Center Automation tools to help streamline all processes and automate the migration of server images to the best hardware while automatically powering off vacant VM managers.
- **Bill-back** Implement policies that ensure server owners are charged per their energy usage.

FACILITY

- **Street power** Integrate pulse monitors on main building power feeds to provide for a total building consumption value. This is often the real bottom line of the majority of energy spending, and this number should be carefully and visibly monitored. It is often useful to have this metric be seen by the whole facility population (managers, employees, etc.) so that everyone can strive to improve the energy usage together.
- **Street gas** Perform the same step as in the paragraph above, but monitor the main other fuels that are coming into the building (typically gas for heating in cooler climates).
- **Diesel** Don't forget to monitor those expensive backup generators. In many hot climates, with Demand Response in place, many companies are obligated to run these generators for many hours every day (to reduce load on the electrical grid). Seldom do companies factor this energy cost into their total energy use.

Of course, closely monitor all the dashboards and reports during each change and ascertain if the right energy benefit was realized.

SECTION 3

## What Are the Benefits?

### Money Saved

Clearly, controlling energy usage will have a direct positive impact on the bottom line operational costs. With such high and escalating costs of energy today, this is the primary driver for most energy conservation programs.

The use of instrumenting and monitoring tools will help prove that the anticipated ROI's were indeed achieved. It is senseless to perform energy saving tactics without clearly measuring the benefits and then continuing to measure the benefits moving forward. Saving money is good for the organization's bottom line, and many energy saving approaches can reduce total energy spending by 20 percent to 50 percent. — a very significant savings.

### Sustainability

Conserving our natural resources so that future generations can be fruitful is at the heart of many initiatives. For too long as a society we have been sloppy in our use of all of our natural resources, and now in the face of both higher costs and limited global supply, we have to adjust our behaviors. Working in a sustainable environment and endorsing sustainable operations is paramount today.

### Public Appearance

Everyone is watching everyone else today. Every buyer in a supermarket, every observer of highway traffic, and every shareholder of their holdings is watching how companies and people behave regarding conservation. Companies need to be sure that they are not seen as negligent or wasteful and where possible, need to provide evidence of their positive sustainability programs. Metrics, data, and charts are all great materials to help the public understand your sustainability actions and effects.

SECTION 4

## Conclusions

There is no doubt that energy costs are an all time high and are now so appreciably high that it affects the bottom line profitability of many companies, especially those in the IT or voluminous facility spaces. Today, a wide variety of energy monitoring gear that can easily meter all of the energy users within an operation is now available. Thousands of points can now be monitored in real time and brought into tools for trending and analysis.

Most changes to improve energy usage will require some amount of cost, be it capital costs, for example, changing to variable speed fans or newer cooling equipment, or deployment of sophisticated automation software. Capital costs have to be balanced by a return and only through deep monitoring and tracking can proper analyses be made. Thankfully, both the monitoring gear and software are now readily available and easy to deploy and understand.

SECTION 5

## About the Author

**Chris Stakutis**
Vice President,
Emerging Technologies, CA

Chris Stakutis is currently Vice President of Emerging Technology with CA, focusing on new and pervasive devices that are providing volumes of information and connections. Chris is a renowned inventor with over 20 patents, 2 published books, and numerous articles. He is often seen speaking on the dramatic changes that pervasive technology is bringing.

Prior to CA, Chris was CTO of IBM/Tivoli's Advanced File System market initiative, which focused on bringing real-time file system management capabilities across the enterprise. Chris was the founder and CTO of SANergy, sold to IBM in 2000. SANergy was the first shared-SAN file system technology to use a split-data/meta-data approach to high-speed data sharing.

Chris's background is largely in real-time operating systems and process controls. At Mercury Computers, he was one of the lead architects for the massively parallel operating system used today by high-end defense and imaging applications. Prior to that, Chris was the lead architect at Precision Robots where he developed applications to operate a chorus of robots in a wafer handling environment. Many years at MIT/Lincoln-Labs gave Chris a broader understanding of the science of engineering coupled with the value of commercialization.

Chris has 20+ years experience in the field of computer science. He completed his BSCS from Worcester Polytechnic Institute in 3 years and his MBA from Babson College in a leisurely 10 years.

CA (NASDAQ: CA), one of the world's leading independent, enterprise management software companies, unifies and simplifies complex information technology (IT) management across the enterprise for greater business results. With our Enterprise IT Management vision, solutions and expertise, we help customers effectively govern, manage and secure IT.

MP332271008

Transforming IT Management