

Tackling Today's Data Center Energy Efficiency Challenges – A Software-Oriented Approach

*James Parker, PE, CEM, CMVP, Square D Engineering Services
Hugh Lindsay, Schneider Electric, Power Monitoring & Control
Bill Brown, PE, Square D Engineering Services*

In Conjunction with the Square D Critical Power Competency Center

1. Abstract

Energy consumption is a growing concern for data centers. Advances in server equipment technologies and increased demand for computing power have increased load densities in the computer room, which in turn has caused corresponding increases in data center power consumption. Energy efficiency measures are thus of high importance for data center designers, operators, and owners. This paper outlines a software-based approach to the data gathering, trending, and analysis that are necessary to apply successful energy efficiency measures in data center environments.

2. Introduction

Driven by the sheer amount of consumed electrical power and related year-over-year increases, growing scrutiny is being placed upon electrical efficiency in data center applications. In the year 2000, servers and their associated cooling equipment in the United States (not counting network and storage equipment) consumed approximately 23 billion kWh [1]. By 2005, this figure had risen to 45 billion kWh, accounting for 1.2% of all 2005 U.S electricity sales [1]. The total magnitude of data center energy consumption, including network and storage equipment, is now estimated to be between 1.5% and 3% of all electricity generated [2]. This increased energy consumption also comes with an increased price tag: Currently, the typical 3-year cost (operating expenses + amortized capital expenses) of powering and cooling servers is approximately 1.5 times the cost of the server hardware itself, with projections of up to 22 times by the year 2012 [2].

The reason for increased energy consumption in the data center is, simply put, an increased demand for computing power. Typical server computing throughput per watt has increased by a factor of 16 from 1999 to 2007 [3]. An increased need for network storage is also a factor.

An exacerbating factor for the increase in energy consumption is the trend toward increasing load density in the computer room. At the end of 2005, the average enterprise data center gross computer room heat density was 32W/ft.², an average increase of 39% over 1999 levels [4]. This is a result of advances in semiconductor technologies that yield smaller chips, along with advances in server technology, such as the advent of blade servers, that together allow more computer equipment to be packed into the same physical footprint. Since the heat generated by computer equipment must be removed in order to avoid overheating, increased computer load density results in increased heat density, which becomes a challenge for the HVAC equipment design. This trend is expected to continue.

2.1. Reducing Energy Costs Through Increased Energy Efficiency

Data center energy consumption reduction may be carried out in a number of ways. Reducing the IT electrical load is an obvious way of reducing energy consumption, since the IT load also affects the amount of cooling power required to remove heat from this equipment. However, as described in [3], manufacturers of semiconductor chips and the servers that use them are already increasing the computing throughput per watt such that it roughly doubles every two years. The increasing power consumption of this equipment is therefore an indicator of the increased demand for computing power, as mentioned above. Reductions in IT electrical load must therefore be made in ways other than hardware efficiency improvements.

One such technique is server virtualization, in which the server environment is controlled based on *perceived* server activity. Server virtualization can itself be viewed as a growing trend that includes storage virtualization, network virtualization, and workload management, elements which can help to eliminate server sprawl and make more efficient use of available server resources.

Large gains in efficiency may also be achieved by increasing the cooling efficiency of the environmental system: Free cooling, water-side economizers, liquid cooling at the computer equipment rack level, and “fine tuning” of the system. The technical details of these energy reduction/efficiency methods are not the subject of this paper, which will instead focus on the methodologies and tools available to evaluate energy saving opportunities and the technologies that can enable significant insight into (and validation of) the results of energy conservation measures.

An important point when considering the reduction of data center energy consumption is that reliability cannot be compromised. When making changes, transient conditions such as the loss of utility power and subsequent loss of cooling until standby generators can be started and brought on-line must be taken into account, in addition to normal operating conditions. The need for high reliability during such transient conditions can lower the maximum energy efficiency gains that could be achieved in the normal operating condition; however this is unavoidable given the mission-critical nature of most data center applications.

2.2. Measuring the Success of Energy Efficiency Implementation – The Software-Oriented Approach

What is often lacking in current literature on this topic is the need for an infrastructure that allows for the successful tracking of energy efficiency measures. This infrastructure allows tracking of energy consumption as a function of time down to the desired level, giving a clear measure of the effectiveness of energy reduction techniques. If energy reduction techniques are staged so that they are implemented at different times, this infrastructure gives insight as to the effectiveness of each technique.

Such an infrastructure consists of:

- 1.) Appropriate measuring and recording devices and their related instrumentation hardware
- 2.) Network communications between measuring/recording devices
- 3.) Software that interprets the results of the observed measurements

The importance of the software cannot be overstated. In fact, the “interpretation” of measurements can be quite sophisticated, and may be expanded to include not only a report of current system status and key performance indicators, but also advanced analytics, budgeting, and load forecasting. In this role the software automates many of the tasks that are necessary to document the success of energy efficiency efforts. It also automates many of the tasks that are required for continued reliable operation of the system and planning for system changes such as IT hardware change-outs. Using software in such a role is referred to as a *software-oriented approach*, and is the topic of this paper.

3. Strategic Energy Management

Strategic Energy Management is both a business process and a framework for establishing, monitoring and verifying results of plans for improved energy efficiency. In its simplest format, it defines a plan for evaluating current performance, identifying opportunities for efficiency gains, implementing actions to achieve savings and methodologies for tracking the real results of those actions.

Standardized process examples are available from organizations like Energy Star and the California Energy Commission and are intended as guidelines for the development and communication of a corporate energy efficiency plan.

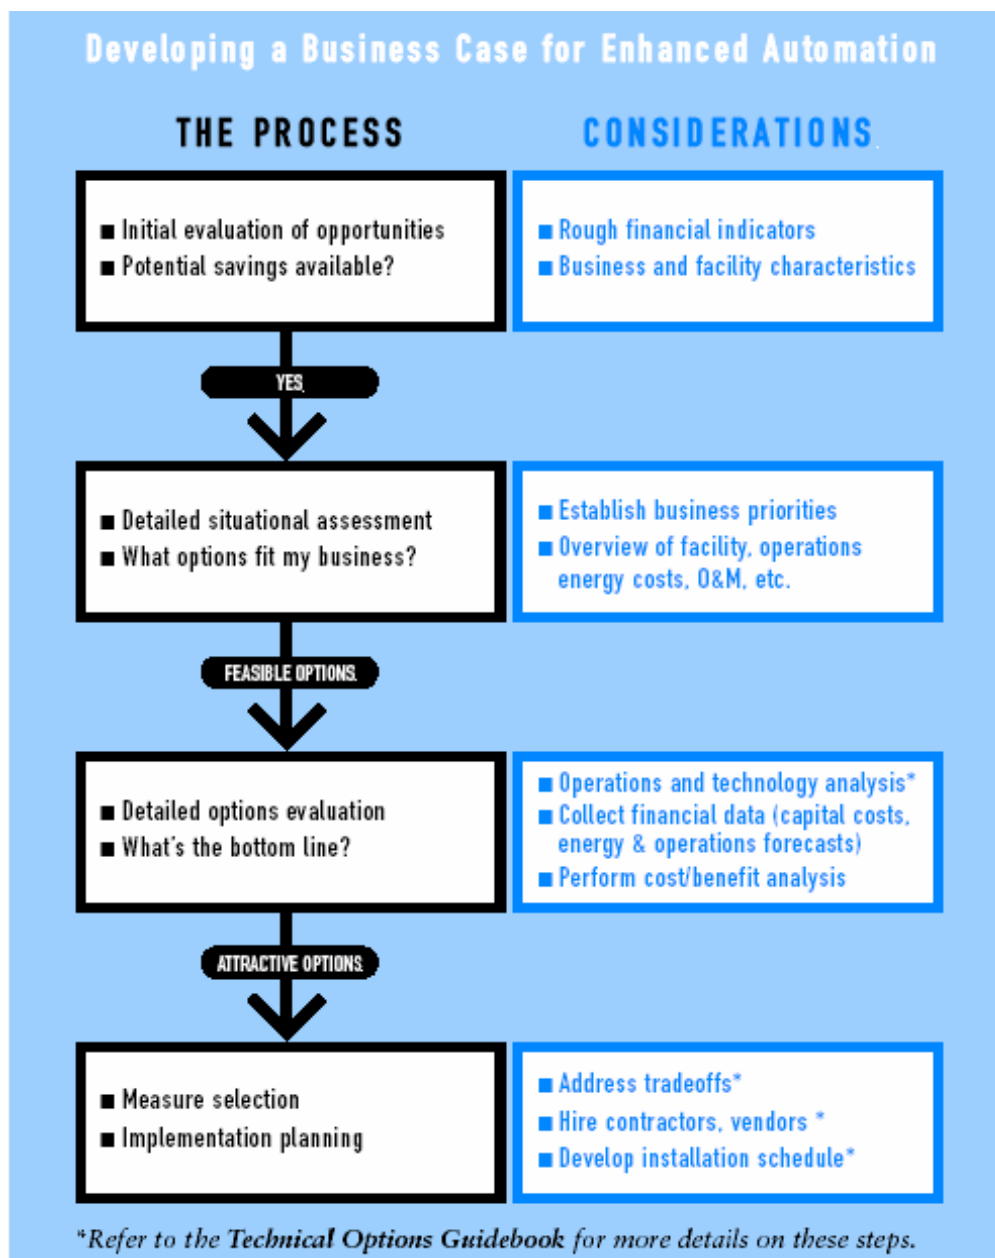


Fig. 1 Sample Strategic Energy Management process diagram. Enhanced Automation - Business Case Guidebook, California Energy Commission, 2004

A key element to the success of any energy efficiency project is accurate and relevant data. This data, captured and presented in an intuitive and flexible software interface, provides significant insight into efficiency opportunities, supports the business case for investments, and reports on the successes of the actions taken to improve performance. The most powerful of such software interfaces are able to mask engineering complexities and represent the data as energy metrics and key performance indicators (KPIs).

3.1. Defining Energy Metrics

Data center infrastructure is under constant pressure to accommodate growth and change. Such pressures come from:

- Evolutionary upgrades to legacy systems
- Additional server and storage requirements to support business growth
- Increasing power and cooling requirements to support new server technology
- Variable server loading that result from the shift to On-Demand computing.

Each new change requires careful consideration of the infrastructure resources that are available to support this growth and of the potential of that change to impact reliability and increase risk. From an energy efficiency perspective it is equally imperative to understand and factor these changes into the calculations of performance, otherwise these calculations could lead to a false sense of success or failure of any particular or overall energy conservation project.

Strategic Energy Management programs and coincident Measurement and Verification practices define methodologies for tracking success. Measurement and verification practices, like those defined by the Efficiency Valuation Organization in the *International Performance, Measurement and Verification Protocol*, defines energy savings calculations as [5]:

$$\text{Energy Savings} = \text{Base Year Energy Use} - \text{Post Retrofit Energy Use} \pm \text{Adjustments}$$

The 'Adjustments' term is a reference to the need to compare the two time periods against a set of variables that might otherwise influence calculations. Variables such as weather, facility expansions and even processing load can significantly effect consumption and efficiency calculations.

Modern monitoring technologies and analytic software systems are key to automating these calculations (which are often complex) and can provide immediate insight into the current and predicted future energy performance. Specifically, these systems become a layer of Energy Information Technology (Energy IT) that:

- Measures and calculates and Energy Baseline
- Establishes and tracks Energy Performance Benchmarks
- And, creates active whole building and sub-system level Energy Performance Models.

The benefits of these systems has been proven in a number of independent research studies, including a study by US Department of Energy personnel [6] which concluded that energy efficiency projects that included active monitoring achieved a 45% greater level of savings as compared to non-monitored projects.

3.1.1. Developing an Energy Baseline

An *energy baseline*, simply stated, is the "starting point" for energy efficiency efforts. Establishing a *baseline* allows comparison of the performance of a system or process at one time against itself at a later time. For energy consumption, such a baseline can be used to evaluate the effectiveness of an investment in energy savings projects as a verification of savings.

Essentially, an energy baseline is a measure of the amount of energy that would normally be consumed without implementation of energy conservation measures. Establishing a baseline gives a fixed reference against which the efficacy of energy efficiency projects can be compared. As stated by the EPA's Energy-Star Program: "Understanding current and past energy use is how many organizations identify opportunities to improve energy performance and gain financial benefits...Assessing performance is the periodic process of evaluating energy use for all major facilities and functions in the organization and establishing a baseline for measuring future results of efficiency efforts" [15].

A baseline can be established via historical data for an existing data center that has been running for two years or more. Formation of the baseline for such a facility can best be established through the use of time interval based energy metering and data logging. For a new “greenfield” site, establishment of the baseline usually requires an energy model for the facility, as described in section 4.

In data center environments, given that often 80% or more of the energy use is linked to IT and HVAC loads, it is important to develop a baseline for both the whole building as well as for the constituent load elements. This disaggregated view is essential to identify the most immediate opportunities to target for improved efficiencies (Fig. 2).

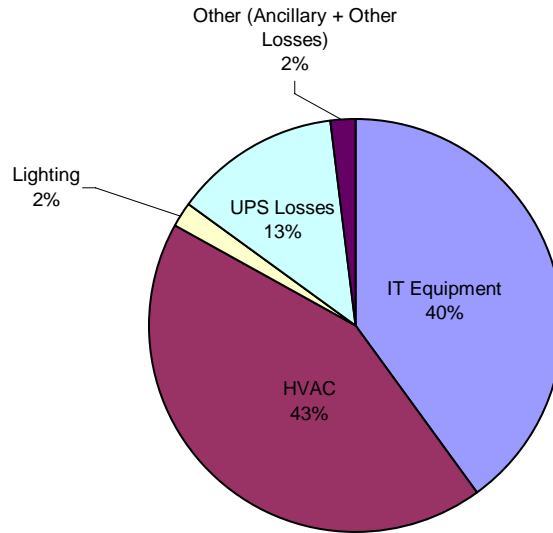


Fig. 2 Representative Data Center Gross Electrical Load Makeup

A further benefit of these disaggregated calculations is an opportunity to categorize energy consumption into groups that segment electrical distribution losses from end-use loads:

TABLE I.
BREAKDOWN OF ELECTRICAL LOSS TYPES AND END-USE LOADS IN THE DATA CENTER

Electrical Losses	End Use Loads
Transformer Losses	Computing
UPS Losses	Storage
Distribution Losses	Networking
Other	Cooling - Chillers / RTU
	Cooling - CRAC
	Lighting
	Office space
	Other

3.1.2. Benchmark historical and current performance

Once baseline measures have been established, it is important to next develop and evaluate the baseline data relative to the relevant performance benchmarks. Benchmarks provide intuitive

insight into the comparative performance of a facility or a component within that facility. Benchmarks can be based on industry standards as identified by a host of credible industry associations and research organizations, such as the Uptime Institute, Lawrence Berkley National Labs (LBNL), and the newly created Green Grid consortium. In addition to these whole-facility metrics, organizations like the Environmental Protection Agency's Energy Star organization and the American Society of Heating Refrigeration and Air-Conditioning Engineers (ASHRAE) offer component-level efficiency ratings.

Another strategy to better understand performance and results is to establish a set of internal metrics and self-referencing benchmarks. These metrics, or Key Performance Indicators (KPIs), can be tracked to characterize site-specific performance, or additionally can be leveraged for multi-site comparisons. When used in conjunction with permanent metering at all important points from service entrance to racks, loading can be calculated in both Watts and real cost in real time at every level. This can include evaluating loading by physical hierarchy of equipment (utility feed, generator, UPS, PDU, circuit) or by logical business hierarchy (building, floor, zone, application, row, rack, etc.). Graphic load profile analyses can help quickly identify areas approaching maximum limits and other areas of underutilization. In this way, a more comprehensive and accurate performance profile can be developed and energy efficiency can be actively managed.

3.1.3. Data Center Industry Benchmark References

3.1.3.1. The Uptime Institute

The Uptime Institute is a data center industry think tank organization, made up of a consortium of member companies responsible for the products and operations of data centers. In addition to the knowledge communities, data center management courses, and other services, the Uptime Institute researches and develops industry standards and defines benchmarks for data center operation and performance. From the Uptime Website [www.uptime.com] ,

“Using benchmarking, abnormal incident data, and industry Best Practices collected from members of its knowledge communities, The Uptime Institute, Inc. (The Institute) has distilled uptime management into scientific disciplines and practices which can be confidently applied.”

A recent example of this research include development of a benchmark for overall Data Center energy efficiency, called the *Site Infrastructure Energy Efficiency Ratio (SI-EER)*.

The SI-EER is defined in [2] as the ratio of power consumed by the entire facility to the power consumed by the IT equipment. A typical value of SI-EER, as measured over several data centers, is reported in [2] as 2.5. This means that for every watt of IT load, an additional 1.5W is expended. This additional 1.5W is comprised of power for environmental control of computer rooms (including cooling and humidification/dehumidification), losses in the electrical system (including UPS conversion losses), losses in the environmental system, and other building loads. According to [2] the theoretical minimum for the SI-EER is 1.6, a 36% energy consumption reduction for a site with a starting SI-EER of 2.5. For many sites, this would represent an even larger reduction in energy consumption since the SI-EER is sometimes above 2.5 [2].

3.1.3.2. Lawrence Berkley National Labs (LBNL)

As part of the research and development activities under the “High Performance Buildings for High-Tech Industries” initiative, LBNL has proposed an overall energy performance benchmark for Data Centers [10]:

- Energy Utilization Index (EUI)

The LBNL Energy Utilization Index identifies a method for calculating energy performance for a whole facility. The calculation is relatively straightforward:

- o $EUI = kWh / Sq Ft / Year / Per Facility$

LBNL further recommends a need for more comprehensive benchmark standards as a method for implementing improved energy efficiency controls.

3.1.3.3. Green Grid

The Green Grid is a consortium of information technology companies and professionals who are singularly focused on improving energy efficiency in Data Centers. While this organization only recently began, its initial publications highlight the need to establish relevant benchmarks and other metrics as a means of referencing energy performance within the industry. In February of 2007, the Green Grid published a white paper titled “Describing Datacenter Power Efficiency” [16]. Similar to the Uptime Institute SI-EER metric, this whitepaper describes the following benchmark references:

- Power Usage Effectiveness (PUE) and Datacenter Efficiency (DCE)

The PUE metric is essentially the same as the SI-EER metric, with a coincident theoretical performance potential of 1.6, calculated as:

- o $PUE = Total Facility Power / IT Equipment Power$

The DCE metric is expressed as the reciprocal of the PUE calculation.

- o $DCE = IT Equipment Power / Total Facility Power$

The Green Grid further proposes a future metric, the Datacenter *Performance* Efficiency (DCPE). While more difficult to determine, this metric expresses a clearer representation of the “Useful Work” by factoring some measure of thermal and data processing efficiency into the equation:

- o $DCPE = Useful Work / Total Facility Power$

3.1.4. Component Level Metrics

Once these whole facility benchmarks have been established and referenced, the natural evolution is to further measure and benchmark performance of the key components of the data center. The following are sources for component-level metrics:

- ASHRAE already offers a host of useful metrics for monitoring thermal performance.
- Energy Star is developing new standards and metrics for profiling the efficiency of IT equipment
- The Green Grid is developing a series of component level metrics.

3.1.5. Self-Referencing Key Performance Indicators (KPIs)

As the industry develops and evolves standardized benchmark references, it is also possible to use monitoring equipment and analysis software to develop and monitor current performance using internally developed metrics or Key Performance Indicators (KPIs). The detail of these metrics is limited only to the availability of data to support the calculations. KPIs can be compared within and/or across facilities. Examples include:

Energy Characterization of Load Types:

- kWh over time for IT Loads

- kW per CFM for Cooling Loads
- Watts per square foot by rack, work cell, usable floor space
- Comparisons of Processing vs. Storage. vs. Telecommunications energy use over time.

Energy Characterization by Cost:

- Total energy cost per square foot
- IT energy cost, by rack, by zone, by floor.

By monitoring the metrics on a regular basis, it is possible to gain significant insight into historical energy performance as well as advance indications of performance degradation. Further analysis and utilization of energy modeling techniques together with monitoring metrics can shift energy efficiency efforts from reactive model to proactive management.

4. Energy Modeling

Energy modeling is the use of calculations, typically performed via computer, to predict the energy use of the system in question as various parameters are changed. For data centers, energy modeling is a powerful tool that deciphers and reveals how to make an entire facility, or groups of facilities, more efficient.

4.1. Why Model?

The fundamental reason to model energy use in the data center is because it allows the evaluation of “what if” scenarios that would otherwise simply be experimentally evaluated. Such experimental verification is, in general, expensive and inflexible. Modeling, on the other hand, is much less expensive and is very flexible to changes in operating parameters.

Another fundamental reason to employ energy modeling is to establish the baseline – the “starting point” of energy efficiency efforts – as described in section 3.1.1.. For an existing data center that has been in operation for two or more years, historical energy use records can be used to establish this baseline. However, historical data is not available for a new “greenfield” data center, and here a building simulation model of the facility can be used to establish the baseline. Building simulation models can accurately predict energy use pattern hour-by-hour for a base year of operation. This model can be used to compare the new data center energy use over the first few months after commissioning and operation.

Another reason to model is for the ability to predict future energy use: Tracking changes in operating parameters such as equipment change-outs or raised-floor additions, and using the historical energy-use data of older equipment to help predict the lower energy usage of planned energy efficient equipment upgrades. Such predictions are invaluable for evaluating the effects of these changes on the over-all energy efficiency of a facility.

Yet another reason to model is to give “before the fact” validation of capital investments related to energy savings. These are based upon the building energy simulation model for a base case and then compared to various alternative control efficiency strategies or efficient equipment investments. Such validation can be crucial to justify spending, particularly on large-budget projects.

4.2. Types of Models

Simple one-time readings can be used for very predictable changes in hardware or operating practices. For example, the energy usage of lamps is simple to measure and their load profile is easy to understand once the operating schedule is known. A simple change might be to install a lower wattage lamp or to turn off existing lamps with a new schedule. A savings model in this case is straightforward and can be predetermined to be a certain amount without instituting a complex modeling scenario.

More complicated pieces of equipment such as chillers or HVAC systems can have performance that varies depending upon percent of rated load, supply water temperature or supply air temperature, internal latent loads, and outdoor influences like dry-bulb or wet-bulb temperature. When major pieces of equipment are changed out, one may need to measure and model energy use over a season to determine energy consumption before and after the change to verify true energy cost savings.

When multiple energy changes of significant value are completed at a site in a relatively short period of time, whole-department or whole-building energy consumption can be used to verify savings. Energy use for one year before can be compared to energy use after to verify savings. This data can be adjusted for other factors such as occupancy changes, weather, etc., as a means of confirming actual savings.

The final model is a comprehensive building simulation model. Various comprehensive building simulation software suites are available to model a facility hour-by-hour to determine baseline energy consumption. Various alternatives to the baseline can be simulated comparing each alternative with the baseline. The expected investment cost of each alternative can then be compared with the modeled savings projections of the alternative.

4.3. Required Data for Modeling

Successful energy modeling requires certain types of data which, in general, must be complete and accurate. The following are considerations when establishing what data is required for a precise energy model:

- Determine the appropriate level of detail: The level and scope of data collection will vary from organization to organization. Some may choose to collect data from sub-meters on individual processes while others may look only at utility bills.
- Account for all energy sources: Inventory all energy purchased and generated on-site (electricity, gas, steam, waste fuels) in physical units (kWh, MBtu, Mcf, lbs of steam, etc.) and on a cost basis.
- Document all energy uses: For the sources identified above, assemble energy bills, meter readings, and other use data.
- Gather at least two years of monthly data, or gather on a more frequent interval if possible. Use the most recent data available.
- Collect facility and operational data: To be able to normalize and benchmark, it may be necessary to collect non-energy related data for all facilities and operations, such as building size, building construction, site orientation to the sun, shading, operating hours, etc.

Energy data may reside in the accounting department, be held centrally or at each facility, or can be acquired by contacting the appropriate utilities or energy service providers. Of course, implementation of the “hardware monitoring” and “software monitoring and analysis” layers discussed in sections 5 and 6 below provide the ideal source for this data.

4.4. Model Normalization

Models or actual test results can produce unexpected or inaccurate results unless the data is normalized. Usually, the site will want to normalize for occupancy and occupancy type, such as servers versus storage versus network, bandwidth used, actual data throughput, weather, major construction modifications and so on.

For example, annual energy use for an existing air conditioning system can be significantly higher one year compared to the next if the weather was much hotter in a particular year. Comparing an old air conditioning system's energy use during one year to a newly installed efficient air conditioning system during the next hotter year could show that the energy use did not change at all if the data is not normalized for weather. The new system had to work longer and harder to meet the higher cooling load and reject heat under more challenging ambient conditions.

5. The Hardware Measurement Layer

A key aspect of any successful energy management strategy is the presence of accurate energy usage data. This data allows the construction of a baseline, as discussed in section 3, construction of appropriate energy usage models, as discussed in section 4, and implementation of the established metrics and energy usage models in the analysis of efficiency measures vs. established benchmarks, as discussed later in section 6. The hardware used to obtain these measurements is referred to as the *hardware measurement layer*.

5.1. What do I measure, and where do I measure it?

The choice of what quantities are to be measured has a large impact on the effectiveness of the software-oriented approach – after all, the results from the analysis software are only as good as the supplied data. The choice of measurement hardware is made somewhat easier in that one modern multi-function power monitoring device can measure many different quantities. Certainly, current, voltage and power/energy measurements at each point of interest are the base requirements. Where to apply such measurements is, in part, dependent upon the definition of the metrics as discussed in section 3 above. Additional functionality can be obtained by specifying more advanced features such as status inputs and waveform captures. The requirements and benefits of sequence-of-event recording are described in [8]. From the standpoint of energy efficiency, Fig. 3 shows the minimum power/energy measurements locations for a typical data center.

5.2. Permanent vs. Temporary Monitoring Solutions

The choice of whether temporary or permanent monitoring solutions are used depends upon many factors, including cost, the availability of space, the size of the system, and whether some monitoring capability already exists in the infrastructure.

Temporary monitoring, in the form of portable data logging test equipment, can be used in small-scale systems for limited-scope energy management studies. This can be useful in evaluating such systems over short time periods of weeks to a few months, recording variations in data due to time, weather, and occupancy or production. The results can be extrapolated over time to estimate annual energy use and costs.

Permanently-installed monitoring equipment can also be used for the above, without the need for invasive temporary installation (and subsequent removal) of the equipment. If such equipment is linked to the data center's communications network, it is possible to provide real-time energy information as well. In the data center environment, such a system may already be in place but not used to its greatest potential.

For a large-scale system, permanently-installed power monitoring equipment linked to the facility's communication network is almost always required. The specifications for such hardware should be based not only upon energy management needs, but also upon the need for power quality and event recording information. High-end power monitoring devices can provide all the required monitoring functions for a given system location in a single device. One specification, however, which always is needed for energy management purposes is the ability of the monitoring device to communicate with the "software monitoring and analysis layer" described in section 6 below.

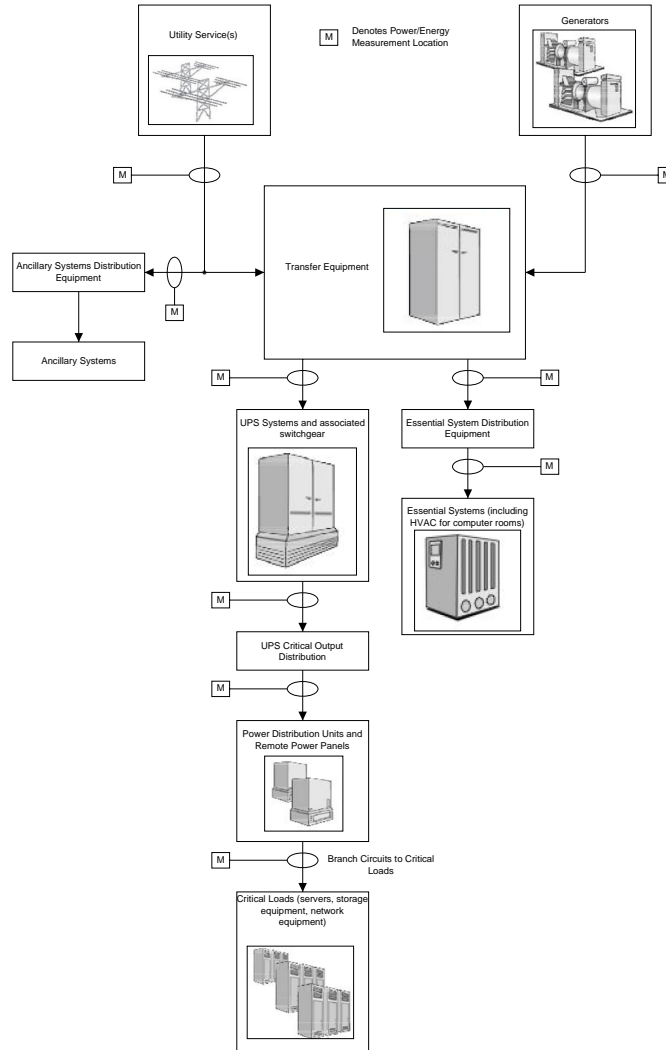


Fig. 3 Minimum Power/Energy Measurements for Assessing Energy Efficiency for a Typical Data Center

6. The Software Monitoring and Analysis Layer

Equally as important as energy usage data collection is the analysis of that data, and indeed all of the time, effort, and expense in implementing the hardware monitoring layer is targeted to making that data available for such analysis. The most efficient means of performing this analysis is via software, hence the term *software monitoring and analysis layer*. Such software consists of both real-time and long-term analysis components, together with a database that stores essential system information. A well-designed software solution should be able to perform energy analysis tasks automatically once the appropriate data is available and pertinent analysis directives are programmed, delivering the analysis in a report format that can be customized depending upon its intended audience.

6.1. The Database

A database serves as a central warehouse of information on the system being monitored. Such data includes both the content of the system (i.e., what equipment is installed), the data collected

by the hardware management layer, as well as the results of the analyses. This data can be accessed by other components of the software when necessary. The database should be flexible enough for the user to define new data types when required in order to accommodate unforeseen situations and to tailor the software solution to the facility.

6.2. Real-Time Monitoring and Analysis

For real-time analysis, key performance data is monitored in real-time and typically can be displayed in any way the user chooses. The real-time analysis of energy data is, in general, operating in parallel to that for power quality data, such as a sequence-of event recording system, and may or may not be part of the same software platform.

As a tool to flag gross changes in energy usage, real-time monitoring and analysis is invaluable. It helps establish energy usage patterns for equipment and identifies equipment and systems that are operating outside of these normal energy usage patterns. As a piece of equipment, system, or subsystem falls outside of its normal energy usage pattern, action can be taken for planned maintenance in an orderly fashion to keep the data center at peak efficiency. Real-time monitoring and analysis is also crucial for the real-time control of energy, as described in section 7.

During the initial phase of operation of a new data center, energy monitoring and management software can begin data collection to establish actual energy use patterns of the facility and many of its sub-systems. Data cleansing algorithms can provide automatic data correction of metered energy data and fill gaps in missing data.

Care should be used in attempting to compare real-time values of the energy metrics, established as described in section 3, against their associated established benchmarks. Unless the benchmark in question was specifically developed for use in real-time, the results will not be meaningful. This is due to the fact that the energy used (and its cost) accumulates over time, and it is this use of energy over time, when properly kept within the strategic performance goals of the facility, that gives the energy savings desired.

6.3. Long-Term Monitoring and Analysis

Long-term data can be used to provide accurate indicators of the energy performance of the data center. It is this data for which the energy metrics described in section 3, and their associated benchmarks, are typically developed.

In long-term monitoring and analysis, energy usage information is gathered and documented over time. Such data includes the specific data monitored by the hardware monitoring layer, plus:

- Monthly electricity, fuels, water and sewer, and other utility bills
- Production or data throughput over time
- Data related to weather, occupancy, unplanned shutdowns, and other influencing factors
- Significant events during the base period of analysis
- Hardware change-outs that can effect energy usage.

6.3.1. Baselineing and Benchmarking

Analysis is performed after long-term data has been collected,. The first step in this analysis is baselineing and benchmarking. In general, whole-building data should be collected for at least one-year prior to this process. As discussed above, baselineing determines the starting point from which to measure progress. Benchmarks, when applied to a single facility, compare the energy performance to peers and competitors over time. For a group of facilities owned by one entity, benchmarks can also be used to compare facilities to one another in order to prioritize which facilities to focus on for improvements. In performing baselineing and benchmarking, the actual

energy use over the calendar year, versus month of the year, versus day of week, versus time of day, etc., should be analyzed. Any periods related to holidays, extreme weather catastrophes, national or international events, etc., should also be taken into account.

At the end of the baselining process, key metrics as defined in section 3 are calculated and declared as the baseline. These baselines may then be compared against the appropriate benchmarks to determine what improvements are required.

6.3.2. Trend Analysis

Once baselines and benchmarks are established for the important energy metrics, data continues to be gathered over time and analyzed in order to understand energy usage patterns and trends. Such trends can be used to help identify improvement potential in key equipment, systems, and subsystems.

Energy use and performance patterns over time can be evaluated to:

- Determine the capability of existing equipment at present and for future growth
- Identify energy saving measures that can be implemented at low cost for cost savings
- Identify problems with HVAC or other utility systems
- Enable prompt action to repair or replace faulty or failing equipment
- Help identify what changes can and cannot be made due to effects on system reliability.

Using long-term data, the software should have the ability to model energy use while accommodating changes in weather, production, or occupancy. It should also be able to generate an overall model – as well as sub-models – for time ranges such as weekdays, weekends, holidays, shifts, etc. Flexibility in data collection should allow data to be collected and modeled over various interval time periods including 15-minutes, hourly, daily, weekly, and monthly.

The software should also have tools available to forecast energy use at the site based upon the energy model. A benchmark can be established to compare one process to another, or, automatically normalized for weather variations, compare one facility to another located in a different vicinity.

Various graphical depictions of performance provide a visual means of evaluating performance and savings. The software should have the capability to correlate energy use to weather or other influencing factors. Also, management reports should be automated to convey information related to energy cost and performance history.

Another useful function is the ability to provide a graphical comparison of the cumulative variation of actual energy values versus modeled baseline values. If actual consumption perfectly matches the baseline, then this cumulative sum will be zero. If the actual consumption is higher than the baseline, it will be positive. If the actual consumption is lower than the baseline, the CUSUM will be negative. The report screenshot below shows an example:

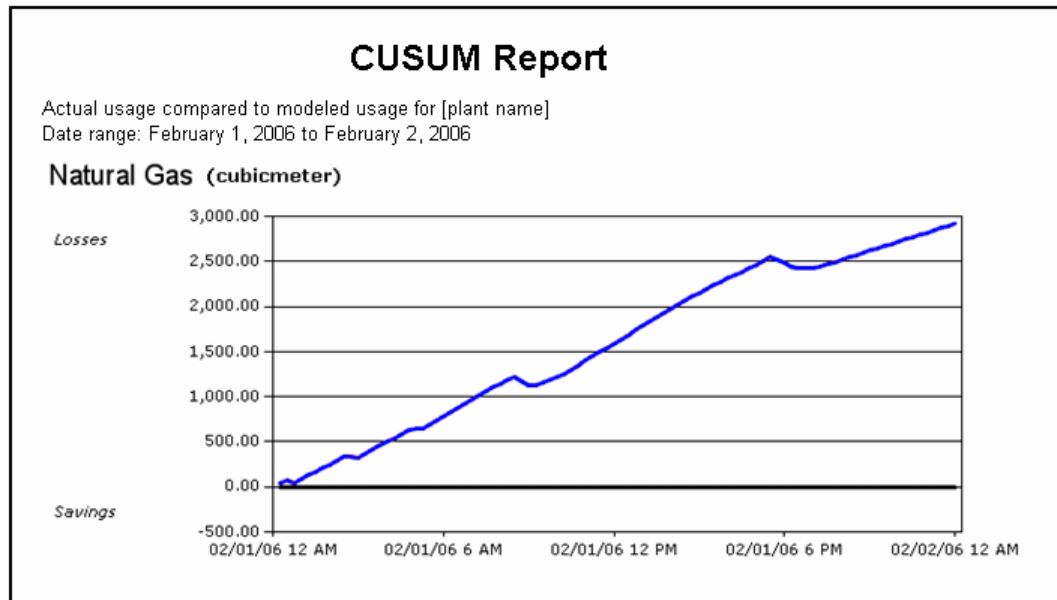


Fig. 4 Example of a “Cumulative Summation” report

By assessing energy performance over time, creative solutions can be developed to lower the over-all consumption of energy. Such solutions can [15]:

- Categorize current energy use by fuel type, operating division, facility, product line, etc.
- Identify high performing facilities for recognition and replicable practices
- Prioritize poorer performing facilities for immediate improvement
- Understand the contribution of energy expenditures to operating costs
- Develop a historical perspective and context for future actions and decisions
- Establish reference points for measuring and rewarding good performance.

7. Looking Ahead: Data Center Energy Control in Real-Time

As data centers look to cut energy costs by lowering energy consumption, methods that involve active control of energy costs in real-time are being explored. While some of these methods involve new technologies and others do not, their application to the data center environment is new. Such techniques move beyond static historical models to active evaluation of performance based on multi-variant inputs. In order to realize such control, both the hardware monitoring layer discussed in section 5 and the software monitoring and analysis layer discussed in section 6 must be very robust.

7.1. Component-Level Optimization

Many component-level opportunities exist to “fine-tune” energy performance. One such opportunity is the optimization of cooling tower performance. Larger HVAC systems use water-cooled chillers to provide building cooling. The chillers cool buildings by low-temperature water to the building air handlers. The chillers then transfer that heat picked up in the building loop to a separate water loop that goes outside the building and ejects this energy to the outside world via cooling towers. Cooling is accomplished partly by the direct conductive heat transfer to the outside air, but more significantly, a small portion of water evaporates to the outside air providing a large evaporative cooling effect. It is through this phenomenon that the water can be cooled to a lower temperature than the dry bulb (normal ambient temperature) of the outside air. Practically

all systems cycle the fans to maintain "basin temperature", a set temperature in the basin at the bottom of the tower. The colder the basin temperature, the less energy used by the chiller to accomplish its job of transferring energy between the building chilled water loop and the cooling tower loop. There is a point where more energy is expended in cooling tower fan horsepower making the tower water colder than is saved in chiller energy when the colder water gets back to the chiller. By measuring the wet bulb temperature and tuning the control to find this maximum savings point (rather than controlling the fans based on basin temperature alone) energy to the cooling tower fans can be reduced.

While this is not the only example of component level optimization, it does serve to illustrate how such optimization strategies work.

7.2. Demand-Side Management

Demand-side management is the reduction of energy costs by directly controlling what equipment is operating and how it is operating. Practically all such strategies factor in the real-time cost of energy when making decisions about the operation of the system.

7.2.1. Peak Shaving

Peak shaving is the use of on-site energy sources to replace some of the electricity required from the serving electric utility. In such a scenario, the on-site energy source drives a generator which can be brought on-line in parallel with the utility, supplying some of the electrical energy required for the facility. The utility rate structure must be taken into account when deciding when to bring the generators on-line; such times are typically during periods of peak demand where additional utility demand charges come into play. The capital cost of the generation and paralleling equipment, if they are not already installed, can be prohibitive unless the on-site energy source is cost-effective enough to offset this cost. Standby generators already exist in most data centers, and with some modification to the paralleling controls can be operated in parallel with the utility in a "base load" mode. In the ultimate version of such a scheme, the software monitoring and analysis layer determines when the generators should be brought on-line, how long they run, and at what power level.

7.3. Load Shedding

Load shedding is, simply stated, the removal of unnecessary or low-priority loads. In the data center it is not possible to simply "turn off" computer and IT equipment at will, of course. In the data center environment the optimization of the number of running servers – server virtualization – is one way to achieve load shedding. Such optimization can run independently of the software monitoring and analysis layer. Another way of optimizing power consumption at the single-server level is to enable the energy-saving features of the server.

7.4. Energy Storage

Another method of energy control is energy storage. In an energy-storage strategy, energy is stored (i.e., the system is "charged") when energy commodity costs, such as electricity, are low, and the stored energy is maintained using a minimal amount of energy until these commodity costs are high. At this point stored energy is released into the system, offsetting the cost of the energy commodity in question.

An example of energy storage is thermal energy storage. In such a scenario, the system is charged when electricity prices are low, "charge" is maintained using a minimal amount of energy, and at the optimal time the stored thermal energy is released, allowing shedding of portions of the HVAC load. The software monitoring and analysis layer will determine the optimal time to store and release the stored thermal energy.

7.5. Load Shifting

Load shifting is the transfer of load from a location with a higher energy commodity cost to a location with a lower energy commodity cost. This would involve multiple sites which are geographically separated enough to cause differences in their energy cost structures, and between which such loads can be shared. For data centers, if sufficient communications bandwidth is available, processing load could be dispatched between facilities based upon the cost of energy, optimizing the dispatch to minimize the cost of expended energy. Such a scenario is conceivable, although to the authors' knowledge it has never been implemented. For the future, however, it is possible that such schemes could proliferate for large enterprises that have enough geographically-dispersed data center sites. This would maximize the use of the software management layer at each site, while an over-all centralized software management layer attempts to minimize the over-all energy consumption across the enterprise. While the use of such a strategy remains in the future, its roots are in the here and now, using the principles of the software-oriented energy management approach.

8. Summary

Energy consumption in the data center is growing; with this increased consumption comes an increased price tag, both in energy prices and the capital costs associated with the infrastructure that brings energy to the computer. Techniques for the reduction of energy consumption in the data center environment take many forms, and if these are implemented without an over-all strategy the maximum over-all energy savings will not be achieved. Such a strategy involves a large amount of energy usage data and a large number of computations, some of which must be performed for different sets of data. Doing this by hand is labor-intensive and may not give optimal results. Instead, the implementation of both a hardware monitoring layer and software monitoring and analysis layer allows the gathering of energy usage data and the analysis of this data to achieve maximum energy savings. Such a system can use elements of already-installed power monitoring hardware. Investing in such a system can allow data centers to realize energy savings today, and increase energy savings in the future as new energy-saving techniques are developed and implemented.

9. References

- [1] Koomey, J. G., "Estimating the Total Power Consumption by Servers in the U.S. and Around the World", February 2007, available at <http://www.koomey.com>
- [2] Brill, K., "Data Center Energy Efficiency and Productivity", 2007, available at <http://www.uptimeinstitute.org>
- [3] Belady, C. L., "In the Data Center, Power and Cooling Costs More Than the Equipment it Supports", *Electronics Cooling*, vol. 13, no. 1, February 2007.
- [4] Brill, K., "2005-2010 Heat Density Trends in Data Processing, Computer Systems, and Telecommunications Equipment", 2006, available at <http://www.uptimeinstitute.org>
- [5] EVO 30000-1:2006, *International Performance Measurement and Verification Protocol*, Energy Valuation Organization, 2006, available at <http://www.evo-world.org>
- [6] Kats G., Rosenfeld A., McGaraghan S.: "Energy Efficiency as a Commodity: The Emergence of a Secondary Market for Efficiency Savings in Commercial Buildings", 1997, available at http://195.178.164.205/library_links/proceedings/1997/abstract/ece97046.lasso
- [7] Lindsay, H., Westbrook, B., Tobin, T., "Maximizing Data Center Reliability and Utilization Using Enterprise Energy Management Technology", 2005, available at www.pwr.com
- [8] Brown, B., Kozlowski, M., "Power System Event Reconstruction for Modern Data Centers", 2006, available at www.criticalpowernow.com

- [9] Loper, J., Parr, S., "Energy Efficiency in Data Centers: A New Policy Frontier", The Alliance To Save Energy, 2007, available at www.ase.org.
- [10] Lawrence, E.O., "Self Benchmarking Guide for Data Center Energy Performance", Lawrence Berkley National Laboratory, 2006, available at www.lbnl.gov.
- [11] Van Gorp, J.C., "Using Key Performance Indicators to Manage Energy Costs", Power Measurement, date?, available at www.pwr.com.
- [12] American Society of Heating, Refrigeration, and Air Conditioning Engineers, Inc., *Datacom Equipment Power Trends and Cooling Applications*, Atlanta, GA: ASHRAE 2005.
- [13] Fichera, R., Koetzle, L, Powell, T, "Power and Cooling Heat Up The Data Center", Forrester Research, 2006, available at www.forester.com.
- [14] :Lehman, B., Ilkwocki, B., Battenfield, S., Peet, S., "Data Center Cooling Index", The Uptime Institute in conjunction with IBM, 2006, available at www.uptimeinstitute.org.
- [15] Selected Text from the EPA Energy Star website: www.energystar.gov.
- [16] The Green Grid, "Describing Datacenter Power Efficiency, Portland, OR: Green Grid 2007", available at <http://www.thegreengrid.org>