

ENERGY & POWER

MANAGEMENT

Serving the commercial, industrial and institutional markets

Taking the Long View

A systematic approach to data center designs

BY WILLIAM J. KOSIK, PE, CEM, LEED AP, EYP MCF

Ongoing technological advances in high-performance computing continually increase the importance of reliable power and cooling for data centers. As the electrical power consumption of servers continues to grow, the cost of energy and the difficulties associated with cooling very high-density data centers has driven the need for innovative yet cost-effective design solutions for electrical and cooling systems. Because energy costs can exceed over one-half of the total maintenance and operational costs of data centers (and can equal the cost of the server itself within the first 3 years of operation), it is increasingly important to implement power optimization strategies wherever possible.

Additionally, an important issue to consider regarding energy consumption in data centers is related to the “eco-footprint” the facilities cause. A typical 20,000 square foot (ft²) enterprise data center at 100 watts per square foot (w/ft²) will have a peak cooling demand that is comparable to a 200,000 ft² commercial office building, and total annual energy consumption comparable to that of a 400,000-ft² commercial office building.

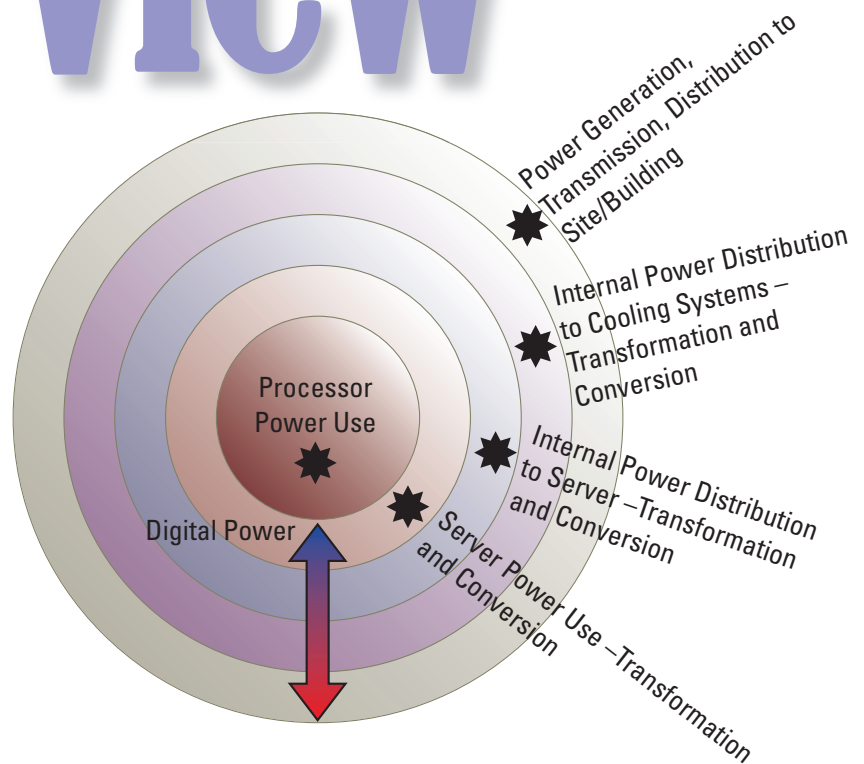


Figure 1. Depiction of power delivery from source to core

The best way to approach delivering the maximum amount of power to computer equipment in the most efficient way must be examined holistically.

Holistic View

To reduce the overall power consumption of a data center or to maximize the power available to the computers, designers and users must consider the interdependencies between the

various power consumers (see figure 1).

Depending upon the type of fuel (nuclear, coal, natural gas) and the equipment used for electrical power generation (steam turbine, gas turbine, reciprocating engine), bringing power to the data center from the source of generation will consume more than 50% of the available heat in the fuel. This major loss does not include additional losses (albeit modest by comparison) involved with sending electricity

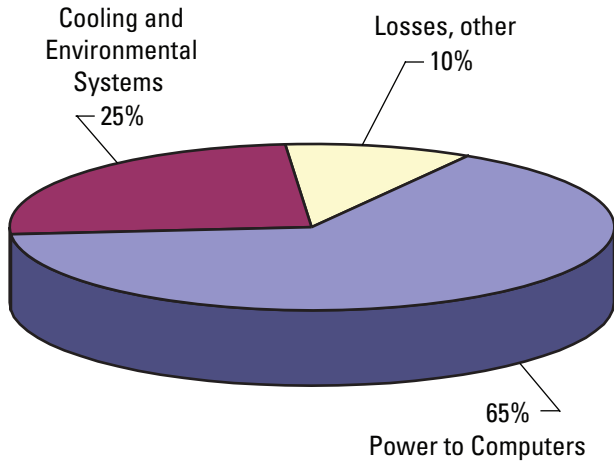


Figure 2. Typical breakdown of power requirements for data centers

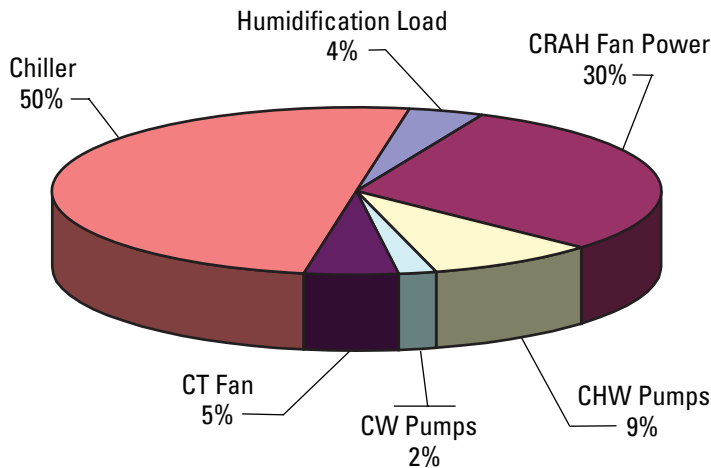


Figure 3. Typical breakdown of power requirements for cooling and environmental systems in data centers

across the power grid to its destination.

Looking internal to the data center, expressed by a ratio of total power required for the data center (which includes power required for cooling and environmental systems) to power required for the computer systems ($P_{TOTAL}/P_{COMPUTERS}$), the efficiency of internal electrical distribution to cooling and environmental systems will generally range from 1.5 to over 2.0. This is an important issue to consider especially as it relates to data centers with a fixed amount of available power.

In addition, there can be significant losses from the internal power distribution to computer equipment, depending upon its topology (see figure 2). The major contributing factors are:

- loading on the UPS system (if one is present)
- type of power conversion equipment
- type of power supply
- type and voltage of power distributed to the computer equipment

Efficiencies will range from approximately 50% (traditional static UPS and power conversion) to upwards of 75% (high-voltage dc power distribution). For every kilowatt meant for the computer equipment sent through the electrical distribution system, the usable power that actually

gets to the computer equipment will vary from 500 to 750 watts.

Computer power consumption is the key point in the overall reduction of power for the data center. Because server power influences both power and cooling, reducing power consumption also reduces the overall power consumption of the data center (see figure 3). An example of this movement is Intel's Energy Efficient System Architecture (EESA) initiatives that take a far more holistic approach to all aspects of the operation of the server in an attempt to find the perfect balance of power consumption and performance (see figure 4).

Many chip and server manufacturers have tried to address this issue in order to optimize the power/thermal capabilities of their computers.

New Power Consumption Metrics

While the server power consumption metrics (watts per square foot and watts per cabinet) are readily understood, generally there is little discussion regarding the metrics for total power consumption of the data center where those servers are housed. It is important to break down the power consumers in order to understand where power efficiency strategies will have the greatest impact. This leads to a question: Is designing for watts per square foot over the raised floor for computer equipment enough or should the industry start to speak in terms of overall power demand of the system as a whole?

Bill Tschudi, program manager at Lawrence Berkeley National Laboratory, has led multiple studies and projects researching data center energy usage and efficiency, including looking at LBNL's own supercomputing facilities. "The focus in supercomputing, in particular, has typically been on acquiring the most powerful computing system one can get," says Tschudi. "Energy has been a secondary consideration. The computer manufacturers and chip makers are beginning to get the message about power efficiency because of the difficulty in providing large amounts of energy and cooling to large data center facilities. Predicting future electrical and cooling needs is a challenge as the technology evolves, however the eventual move to liquid cooling should dramatically improve energy efficiency."

Tschudi and his team have recently collaborated with PG&E to publish a data center design guidelines sourcebook outlining 10 energy-efficient design strategies and are planning on adding a dozen additional best practices (<http://hightech.lbl.gov/datacenters-bpg.html>). They have also been involved in demonstration of dc power distribution for data centers (<http://hightech.lbl.gov/dc-powering/>).

Size Matters

When investigating different air conditioning systems, the different metrics, raised floor area, power density (watts per square foot) and power draw per cabinet (usually expressed in watts or kilowatts per cabinet), will help determine the most effective cooling and power distribution system for the data center. For example, a data center with a 5,000-ft² raised floor area and a power density of 50 W/ft² could be cooled by direct expansion (DX) split-system computer room air handling units. Higher energy efficiency could be obtained by using water-cooled compressors (to lower the condensing temperature). At this size it probably does not make sense to use a central chilled water system due to the cost and complexity. However, take that same raised floor area and increase

the density to 400 W/ft² and a water-based system becomes necessary. Also chilled water or ice storage is required to continue water flow in case of chiller failure or delay in chiller restart. So as the variables of raised floor area, power density, and power draw per cabinet change, so does the appropriate system selection.

Layering in local climate conditions and utility rate structures make the equation even more complex. As owners evaluate potential locations for new facilities, they must consider a wide range of issues from climate to energy costs to power availability.

Locations that have mild weather without extreme heat, cold, or humidity are ideally suited for these types of facilities. The interplay of the building envelope with the internal heat gains generated by the computer equipment is a subtle but important concept when designing data centers. Providing heavy exterior wall and roof insulation in a heavily internally loaded-building can increase the overall energy use, especially in a mild climate where the outside temperature will often be less than the targeted internal design temperature. Mild climates are also ideal for using outdoor air for “free cooling” as part of the heat rejection system.

Using hourly weather data along with a robust energy simulation program to analyze multiple building system, building envelope, climate location, utility rate and usage profiles will allow for optimization of the parameters. This technique, called parametric analysis, is an objective, clear-cut method for analyzing multiple variable problems, and when combined with financial parameters, can be used to determine life-cycle costs of each option.

Costs of Electricity

The cost of electricity, including the cost of a kilowatt-hour, any demand charges and time-of-use charges is obviously important when selecting. Some utility companies offer rebates for installing energy-efficient equipment or systems that can help reduce the peak electrical consumption of a facility. “Each year, more than \$1.5 billion is made available by utilities, government agencies, and other sources to help pay for energy- and water-efficiency improvements to buildings,” says Rob Everhart, senior project engineer at RealWinWin, a leading provider in rebate administration services to corporate America. “In rebate-friendly territories we have seen up to \$0.40 per ft² for energy-efficient new construction, renovation, and capital improvement projects.”

Having the ability to shift load to off-peak hours or disconnect completely from utility power sources during periods of peak electrical usage is a highly desirable aspect when negotiating utility rates on the open market. Having a relatively constant electrical consumption over the course of the day, like a data center, allows electricity generators to more accurately predict their capacity requirements. This profile provides an edge in negotiating the best utility rates on the open market. Says Kris Childress, technical director for GEV Corp., a Chicago-based energy consulting and brokering firm, “Newly competitive markets provide customers with great opportunities to reduce and manage electricity costs. But you should not overlook either the complex regulatory environment in which those markets still operate or the complexity of electricity as a commodity. Electricity supply contracts are correspondingly more complex, and a straightforward economic analysis that may be sufficient for other commodities is not adequate for electricity. This complexity rewards customers who obtain expert assistance to secure and implement optimum electricity supply arrangements.”

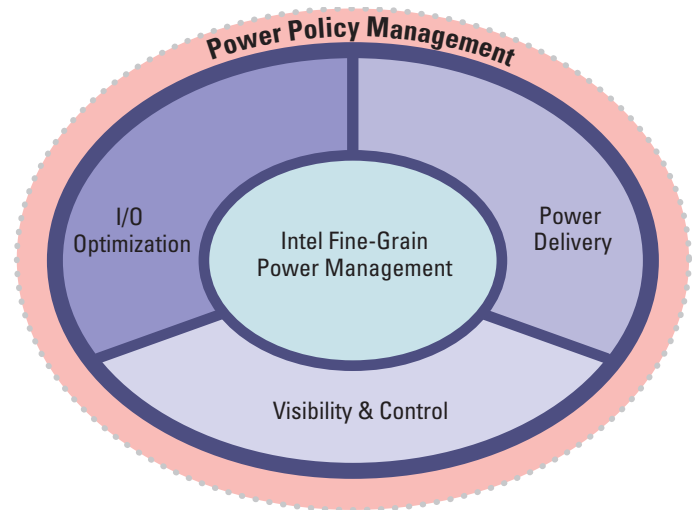


Figure 4. Diagrammatic representation of Intel's energy-efficient system architecture (EESA) operating within their power policy management technology operating parameters

Minimizing Power

It is helpful to break it down the cooling and environmental systems in a data center into sub-systems in order to fully understand the components, when analyzing the energy use of in data centers. The primary building blocks are:

- Heat rejection (cooling towers, dry coolers, evaporative coolers)
- Refrigeration (water-cooled, air-cooled, refrigerant-cooled)
- Air systems (CRAHs, central station air handling units)
- Air distribution (hot-aisle/cold-aisle, heat containment)
- Specialized cooling (water-cooled cabinets, low-pressure refrigerant systems)
- Pumping (condenser water, chilled water)

The chiller is the second highest energy consumer in a data center, and so examining the refrigeration cycle can yield energy savings. The refrigeration cycle varies the pressure of refrigerant to control the evaporation and condensing temperatures. In this cycle, lowering total entropy decreases total work, or the energy that is required. As the air temperature moving across the condenser coil increases, the compressor needs to create higher pressures, thereby using more energy. When water is sprayed over the coil (as in an evaporative cooler), the air temperature decreases, and less pressure is required to cause the refrigerant to condense. Similarly, the higher the air temperature moving across the evaporator coil (chilled water coil in this example), the lower the pressure needs to be to cause the refrigerant to “boil” or evaporate.

So elevating the supply air temperature by increasing the chilled water supply temperature and/or reducing the temperature of the air moving across the condensing coil will cause a significant reduction in energy consumption of the compressor (see figure 5).

Air Distribution

Air-conditioning design, whether it is taught in a university setting, at trade schools, or through industry associations, typically focuses almost exclusively on air-conditioning design for comfort cooling applications. Having this educational background, it seems that the engineers design-




Figure 5. Illustration of power reduction strategy for a refrigeration cycle

ing cooling systems for data centers rely on the same process and rules of thumb that are used in comfort cooling applications. This leads to incorrect design assumptions and approaches that often hinder the performance of the data center and increase power usage (see figure 6).

The typical variable air volume (VAV) systems for commercial office buildings use 55°F air supplied to the space. This air supply temperature is used mainly because it results in a room temperature within the acceptable temperature and relative humidity parameters as defined by the ASHRAE comfort standard. However, in a data center, supplying 5°F air into a room is not necessary if proper air distribution and heat removal/containment strategies are in place. Remember that the optimal environmental conditions at the face of the server are 72°F and 45% RH, well above the 55°F to 60°F that is commonly seen in data centers.

Into the Boardroom

The state of world affairs as it relates to energy security and public concern over long-term degradation of our environment has driven

ENERGY (kW/Ton)					
Chilled Water Temperature (°F)	ENTERING CONDENSER WATER TEMPERATURE (°F)				Integrated
	85	75	65	60	
45	0.565	0.463	0.376	0.338	0.369
47	0.535	0.439	0.353	0.315	0.346
49	0.509	0.414	0.332	0.289	0.321
51	0.483	0.397	0.309	0.271	0.301
53	0.459	0.367	0.288	0.262	0.287
55	0.432	0.346	0.267	0.256	0.275

Figure 6. Example of a 1000-ton centrifugal chiller power consumption per ton of cooling at varying chilled water and condenser water temperatures.

these topics to the front page of major mainstream media. Reduced energy use and the resulting reductions in water use and CO₂ emissions are issues not only debated amongst facilities and engineering directors, but are also topics of concern and discussion for those at the “C-level.”

There are many examples of Fortune 500 companies putting serious time and effort into global energy-efficiency programs. For example, one major financial client sets aside, on an annual basis, funds reserved exclusively for capital upgrades that reduce power use and can demonstrate reductions to on-going operational costs. This type of program can only come from the C level if it is to have any merit.

As Jeff Monroe, managing director of DuPont Fabros, a national developer and operator of large data center facilities, puts it, “DuPont Fabros methodically works towards maximizing data center efficiencies that translate into the most cost effective solution for the customer. Efficiency is a key driver in our business from design to development to operations.” This tack, while primarily focused on reduction of operational costs, also supports an overall goal of reducing power consumption of the cooling and power distribution systems, allowing for more power to be allocated to the computer systems and reducing the overall environmental impact.

Conclusion

Given the increasing electrical power requirements of data centers and the corresponding difficulties in cooling these high-density loads, a new opportunity presents itself to reduce the overall power consumption of data centers by blending both time-tested strategies and emerging technology into a comprehensive power optimization plan. It is not only important to look at the multiple aspects of the performance of the data center itself, but also the greater positive impact on the environment that these strategies could elicit. Understanding and analyzing all of the interdependencies that exist between the computer equipment, the cooling systems, the building envelope and the utility grid will yield a much more holistic approach and allow for prioritization of the different elements and customization of each individual situation. Finally, understanding that technology will continue to become of greater importance and use in people’s daily lives, the need for compute and storage servers and other types of computer equipment, often residing in large, centralized data centers, will also continue to increase. This will certainly drive future requirements to reduce overall power consumption, either as an operational expense reduction, and environmental initiative, or both. *e&pm*

About the Author: William J. Kosik, PE, CEM, LEED AP is managing principal of EYP Mission Critical Facilities’ Chicago office. Kosik specializes in energy analysis, building performance simulation, and HVAC design for technology intensive, high-performance buildings. He has consulted with USGBC to launch the LEED-CS Pilot Program and with the City of Chicago to develop energy-efficiency design standards and was asked to join the Facilities Panel for the US Defense Science Board Energy Strategy Task Force. Visit us online at www.eypmcf.com.