About this research note:

Operate & Optimize notes provide recommendations for effective and efficient IT practices that help improve the performance or reduce the cost of technologies already deployed in the enterprise.

Top 10 Energy-Saving Tips for a Greener Data Center

Publish Date: April 11, 2007

Even in small and mid-sized enterprises (SMEs), the power and cooling of data center equipment can easily cost tens or hundreds of thousands of dollars a year. Employ best practices to boost data center energy efficiency and lower costs.

INFO~TECH
research group
Passionate About Research.
Driven By Results.

Executive Summary

Rising energy costs coupled with growing computing and storage requirements, new advancements in data center technology, and changing attitudes towards the environment are causing enterprises to rethink their data center strategies. Key topics discussed include:

- The rising cost of energy and energy inefficiency in data centers.
- Data center cooling capacity and heat density.
- Emerging energy standards.
- Best practices for achieving a more efficient and cost-effective data center.

As energy costs begin to compete with the cost of maintaining the underlying hardware, traditional thinking around power consumption, heat distribution, and server hardware will change. The new paradigm calls for maximizing efficiency and “greener” thinking.

This note is optimally focused for those enterprises on the “larger” end of the small and mid-sized enterprise (SME) demographic. These enterprises do not have dedicated data center facilities and are more likely housing server racks in a shared facility or room within the enterprise’s corporate offices. For truly “small” enterprises with only a few servers, there are still some great value nuggets, but realistically, power and cooling issues don’t really become an IT priority until the enterprise has at least a couple of server racks. Spending wise, this note targets those enterprises with an IT budget of $1 million or more.

Optimization Point

Data centers operate in an “always on” environment and consume copious amounts of energy. In fact, most estimates suggest power and cooling costs consume approximately 20% of data center budgets. On a macro scale, the US government estimates that US enterprises spend $3.3 billion per year to power their data centers.

At large enterprises, data center energy consumption is becoming a serious issue. Executives are taking note of million dollar electricity bills, demand for power and cooling is quickly outstripping existing capacity, and corporate mandates to become “greener” are bolstering energy efficiency initiatives across the enterprise. Small and mid-sized enterprises (SMEs) are also seeing rising costs that, in proportion to size and revenue, are very significant. SMEs face similar challenges in terms of growing computing and storage requirements, rising energy costs, increasing demand for cooling, and corporate environmental directives, but are often limited in what they can do to solve the problem.

Smaller enterprises don’t have the luxury of designing brand new data centers; in most cases, they are required to make do with the facilities they have. Not only that, but in many ways, they lack the energy management maturity of larger enterprises (whether due to lack of tools or in-house expertise). Still, there are a variety of energy-saving best practices that are universal and apply to enterprises of all sizes.

Key Considerations

The Data Center Energy Imperative

According to the Lawrence Berkeley National Laboratory, the average data center has 10 to 30 times the energy requirement (and cost) of equivalent sized office space. This, coupled with the fact that US commercial electricity prices have increased over 20% in the past five years, means that data center energy consumption is finally attracting the attention of enterprise executives.

Unfortunately, most data centers are also very inefficient when it comes to energy management. While this is costing enterprises a significant amount of money, the good news is that IT leaders have an opportunity to improve the current situation without expanding facilities or investing significantly in new capacity.

One of the most popular power/cooling metrics used today (a variation of which is used by The Green Grid and The Uptime Institute) is overall data center energy efficiency. The metric measures the total energy load delivered to the data center versus what is actually consumed by critical IT loads.

In a best case scenario, a data center would score in the range of 1.6, meaning that for every 1.6 kilowatt (kW) of total electricity going into the data center, 1 kW is consumed by core data center infrastructure, servers, and storage. The remaining 0.6 kW is energy is overhead.

However, according to The Uptime Institute, most data centers are operating at a ratio of 2.5 or higher (meaning at least 1.5 kW of energy overhead per kW of IT load).

Excessive energy overhead can be generated by a number of factors, including over provisioning, poor airflow and ineffective cooling, inefficient power supplies, switching, unnecessarily redundant systems, inefficient/older equipment, and lighting.

The difference between a best case scenario and a worst case scenario can be dramatic in terms of cost. For example, the annual power and cooling cost of a 20kW rack with an efficiency ratio of 1.6 would be $28,032; however, the same rack with an efficiency ratio of 2.6 would carry a cost of $43,800 (assuming $0.10/kW-hour). The opportunity for savings in this example would be more than $15,000 per rack per year.

Cooling Is Half the Battle

Optimizing cooling presents the single largest area of opportunity for IT to save energy (see Figure 1). Due to the increasing power density and heat generation of newer equipment, cooling and air conditioning energy costs now surpass the cost of powering servers.

Figure 1. Typical Data Center Energy Consumption

Source: American Power Conversion, Cisco, and Emerson Network Power

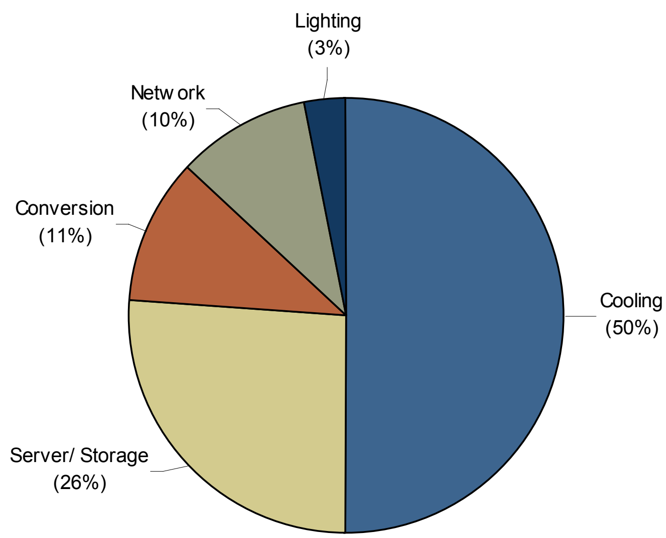

Future data centers will require considerably more power and cooling capacity than is available in most existing environments. According to both The Uptime Institute and HP Labs, the heat density of data centers is currently increasing at a rate of about 15% to 20% per year. As devices get smaller, faster, and cheaper, enterprises are packing more machines into rooms that were not designed to accommodate the increased cooling burden. So, even though individual devices may produce less heat, the aggregate number is driving up cooling requirements and energy consumption.

As an example, data centers three or more years old housed racks that consumed 5 kW to 10 kW of electricity or less. However, newer racks are easily in the range of 20 kW or more. According to IBM, some high-density server racks (such as blades) draw as much as 30 kW of power. This translates into an annual per-rack power and cooling cost of between $35,000 and $60,000, depending on the efficiency of the data center.

Since both heat density and computing requirements are growing, but server rooms are most likely not expanding, heat density will increase not only at the rack level, but also over the entire room. Many existing server rooms were designed based on an assumption of an average heat load (across the room) in the range of 40 to 70 watts per square foot (W/ft^2). These data center environments are not equipped to handle heat loads being generated by many of the newer high density racks that can boost power density to 300 W/ft^2 or more.

Rising heat density is inevitable and not necessarily something that is within the enterprise's control. However, knowing heat densities and identifying heat concentrations for various areas is necessary in order to deploy the appropriate cooling strategy. Higher heat densities also require more cooling redundancy because denser systems increase in temperature much faster, thus reducing the window during which cooling systems can be offline.

Standards are Coming

With standards and ratings, IT decision makers will be able to make better purchasing decisions and better evaluate products based on an energy-inclusive TCO analysis. While vendors and industry regulators are still a bit behind on the power and cooling issue, they are catching up. Several parties in the US, including the Standard Performance Evaluation Corporation (SPEC), the Environmental Protection Agency (EPA), and the US Legislature have recently begun to evaluate server energy efficiency metrics in an attempt to create standards.

The US Senate and House of Representatives recently passed bills (House Bill 5646 and Senate Bill 3684 respectively) that promote data center energy efficiency. The SPEC (which includes representatives from AMD, Dell, EMC, Hewlett-Packard, Intel, IBM, Oracle, Sun, and many others) is ahead of the game and has already begun publishing some preliminary benchmarks.

Another important working group seeking to lower the overall consumption of power in data centers is The Green Grid, which lists AMD, APC, Dell, Hewlett-Packard, IBM, Rackable Systems, SprayCool, Sun, and VMware among its founding members.

Improvement & Optimization

The following recommendations are geared towards those enterprises on the “larger” end of the SME demographic. These enterprises do not have dedicated data center facilities and are more likely housing server racks in a shared facility or room within the enterprise’s corporate offices. Truly “small” enterprises with only a few servers will also benefit from some of the best practices discussed, but realistically, power and cooling issues don’t really become an IT priority until the enterprise has at least a couple of server racks and an IT budget of $1 million or more.

Data center design and management is evolving to maximize energy efficiency and lower cooling requirements. Implement the following ten best practices to ‘do more with less’ when it comes to power and cooling:

1. **Right-size IT equipment.** The number one way to reduce asset acquisition as well as ongoing energy costs is rightsizing IT equipment to match near-term requirements. This is a step that can be taken by enterprises of all sizes – even those with only a few servers. Case studies show that average server utilization rates range from 10% to 20% of capacity. Equipment – including servers, storage, power supplies, UPS systems, air conditioners, fans, and air exchange systems – operating with significant overcapacity is a waste.

   In addition to lowering IT investment in fixed assets, The Green Grid estimates that rightsizing IT equipment can reduce energy consumption by as much as 50%. This is due to the fact that fixed losses in power and cooling are present regardless of how much IT load is actually being used.

   Where possible, smaller enterprises should invest in modular, scalable equipment and add capacity only as growth in IT loads demands it. Also, prioritize availability requirements for applications and services to balance uptime against costs. The ongoing energy costs of maintaining fully redundant systems can easily match or surpass the cost of simply duplicating the hardware. Even if IT has no intention of implementing chargebacks to business units, providing a cost allocation for energy consumption in proposals to business leaders will help convey the message that high availability comes with an associated “energy burden.” Significant energy savings can be achieved by ruthlessly prioritizing and only providing redundancy where it is really needed.

2. **Foster communication between IT and facilities.** For the most part, energy costs are invisible to the IT department because the electricity bill is managed by the facilities department, independently of data center personnel. This problem is exacerbated in small IT shops because in many cases the data center exists in a shared facility where power and cooling are metered at the building level and accounted for as corporate "general and administrative" expenses. This allows IT to remove itself from the problem and not worry too much about how much energy is wasted. The majority of IT departments don't know what the annual energy spend is for the data center and IT and facilities departments don't generally coordinate efforts unless there's a relocation involved.

   However, for energy savings to be realized, there needs to be better visibility into what the other side is doing. IT needs to understand how purchasing decisions affect facilities and what the baseline is in order to track improvement. Similarly, facilities must know ahead of time what's coming down the pipeline in terms of power, space, and cooling requirements. Virtually all energy efficiency-related metric gathering initiatives should be joint IT-facilities projects.

3. **Track and manage the energy consumption of IT equipment.** Tracking energy usage in the data center provides a baseline for calculating the ROI of new energy saving initiatives and helps identify areas of opportunity. Moreover, energy management products can reduce energy use during regular operation and can keep critical applications running during brown-outs.

   Best of breed products will track and manage power intake, monitor inlet/outlet air temperatures, and allow administrators to throttle server CPU performance to match requirements and save power. As an example, HP recently announced Insight Power Manager (originally introduced as a blade product). The software allows administrators to track and adjust how much power servers are using. IBM's PowerExecutive offers similar functionality.

   In order to monitor non-IT equipment such as air conditioning and UPS devices, involve the facilities department. For small facilities with shared air conditioning and electricity metering, these metrics may be difficult to obtain. In this case, consider investing in "smart metering" capabilities that separately measure usage in different areas and provide the desired level of granularity.

4. **Engage in longer-term capacity planning for power and cooling infrastructure.** The growth in electrical and cooling requirements may cause a shortage of power and cooling capacity in the enterprise's data center. A critical extension to right-sizing equipment is using energy consumption data to develop forecasts, assess future capacity requirements, and avoid unpleasant surprises. Since expanding or building a new data center may not be an option, take steps now to reduce energy consumption and improve current capacity utilization.

Tailor the data center cooling solution based on the current and projected heat density and type of infrastructure used. For example, a rack housing 21 2U dual core servers might only draw six or seven kW of power, whereas as a high performance, high density blade rack can draw 20 kW of power. Each scenario requires a different cooling strategy. To expand cooling capacity incrementally, consider precision cooling solutions such as spot cooling that target problem areas. If floor space is an issue, consider units that are either ceiling-mounted or integrated into the rack.

5. **Explore server consolidation and virtualization to improve energy efficiency.** As a further continuation of rightsizing, use server/storage consolidation and virtualization as a way to eliminate redundant/inefficient equipment, improve capacity utilization, and reduce electricity consumption.

   » **Save on hardware too.** Info-Tech research suggests that server virtualization can reduce one-time and ongoing hardware acquisition costs by 40% to 75%, as well as reduce ongoing maintenance costs (incorporating power and cooling) by between 25% and 50%.

   » **Evaluate rebates.** To take savings a step further, utility company Pacific Gas & Electric (PG&E) (California only) offers rebates to IT shops involved in server consolidation and virtualization.

   » **Watch out for the cooling trap.** Consolidation and virtualization can contribute to increased heat density (more computing power in a smaller space) and raise concerns about cooling.

6. **Upgrade legacy equipment.** Older servers, switches, routers, and storage solutions tend to be less efficient and generate more heat than newer models. In addition, older hardware lacks the sensors and instrumentation to support energy management solutions. Holding on to legacy systems beyond their useful life can escalate energy costs both as a result of redundant or seldom-used applications that run unnecessarily and due to the poor efficiency of the old hardware (i.e. inefficient power supplies and fans, excessive heat output, and uneconomical design).

   Keeping up with new advances in energy efficiency will always be a challenge. Developing migration plans for all IT equipment from the get-go will allow for more predictable budgeting and ensure that IT is always running the most efficient configurations possible. Consolidating using blade servers can also help reduce physical space requirements while using less energy for the same processing power as rack mount servers. As an added bonus, many high-performance blade systems also include power management tools to help with tracking and monitoring energy efficiency.

7. **Revisit TCO calculations.** It is no longer acceptable to use equipment TCO models that omit considerations such as annual power and cooling costs. The real TCO of infrastructure components include energy costs, which can easily rival the cost of maintaining the underlying hardware. This becomes especially important when making the case for initiatives such as the replacement of legacy systems, adoption of blade servers vs. traditional rack mounted servers, multi-core processors, and server/storage virtualization and consolidation.

8. **Deliver more targeted and adaptive cooling.** Room-level air conditioners are imprecise and inefficient. For racks with an average heat output of 15 kW to 20 kW (or more), consider cooling solutions at the rack level that transfer heat close to the source as opposed to allowing it to diffuse throughout the surrounding area. Modern rack-level cooling solutions are also adaptive in that they can be automatically adjusted to meet changing heat outputs based on server load. This not only saves energy, it also allows space-constrained enterprises to squeeze more data center equipment into smaller spaces.

   One study from the Rocky Mountain Institute suggests that liquid cooling can reduce cooling energy costs by as much as 30% to 50% when compared against standard air-based room air conditioners. IBM, Sun, HP, and Egenera all recently announced liquid cooling technologies that are expected to be more efficient than traditional forced air methods. Refrigerant cooling takes water-based to the next level by making it safer in the event of a leak or ruptured pipe (refrigerant escapes as gas and will therefore not damage equipment if leaked). Several vendors offer air, water, and refrigerant cooling systems that are integrated directly into the rack or server.

   However, these are not cheap. As an example, ISR's SprayCool M-Series runs about $25,500 for a 1-rack setup supporting up to 36 1U servers. Egenera's CoolFrame for its Bladeframe adds $300 to $400 per blade to the price tag. In addition, these systems may require changes to the existing air conditioning systems and ducting. For more guidance on selecting a blade system, refer to the ITA Premium SE research note, "Blade Servers for Small Enterprises."

   » SprayCool's solution supports Dell, HP, and IBM hardware. In addition, US-based Avista, a Pacific Northwest utility, offers rebates (up to $3,600 per rack) to enterprises that install SprayCool-enabled servers.

   » 42U boasts up to 30 kW of cooling capacity per rack.

   » Liebert has several solutions including enclosures with integrated air or water-based cooling.

   » HP's Modular Cooling System can also deliver up to 30 kW of cooling capacity using chilled water.

9. **Improve airflow management.** Air distribution is key when it comes to efficient cooling. Proper data center air management minimizes the mixing of hot air emitted from equipment with the cool air supplied to it. Efficient air distribution can help increase the density capacity of the data center (W/ft^2), reduce heat-related malfunctions, and lower operating costs. The Uptime Institute found that on average 60% of the cold air supplied to the data center was wasted due to poor airflow. Similarly, PG&E suggests that poor airflow management can decrease the cooling capacity of a Computer Room Air Conditioning unit (CRAC) by 50% or more.

   A recognized approach to efficient air handling and air-flow management is implementing a cool aisle/hot aisle layout. This model is optimized when airflows are separated by enclosing aisles, thus allowing for more effective cooling. Rows of racks are arranged such that there are alternating cold (air intake) and hot (heat exhaust) aisles between them. However, in retrofit situations, where comprehensive redesign is not an option, consider simple low-cost improvements to optimize existing air supply/return configurations.

   » **Look for short-circuiting airflows.** The term “short circuiting” refers to the mixing of hot exhaust air and cool intake air. Make sure cable cutouts and holes in the perimeter walls (and the raised floors where applicable) are properly sealed to prevent this. Also, use blanking panels to fill unused rack space and prevent air mixing in hot aisle/cool aisle scenarios.

   » **Use flexible barriers where fixed ones are not available.** Specialized/custom rack products take cool aisle/hot aisle setups to the next level by supplying cool air directly to the rack air intake and drawing hot air out without letting it travel through the room at all. Depending on the current layout and ductwork, this can be a costly venture. Barriers made of clear plastic or other materials can be used to cost-effectively achieve similar results.

   » **Reposition poorly placed overhead air supplies.** Instead of looking to cool the entire server room (mixing hot and cool air), reposition overhead diffusers so that they expel cool air directly in front of racks where it can be drawn into the equipment. Up-flow computer room air conditioning units (CRACs) should be located near the end of a hot aisle and fitted with ducts that bring cool air to points above cool aisles as far away as possible from the CRAC. Down-flow units should be oriented to blow air down a cool aisle. Return vents should be located over the hot aisles using a dropped-ceiling return air plenum (or hanging ductwork return).

   » **Reposition poorly placed thermostats.** Place the thermostat in front of the equipment where air temperature is more indicative of equipment temperature. Alternatively, for servers that support Simple Network Management Protocol (SNMP) heat sensors, actual equipment temperature can be pulled directly from the management software.

   » **Control fan speeds based on need.** Fans set to operate at maximum capacity regardless of data center temperature waste energy.

Adjust the fan speed of CRAC air handlers in order to reduce the air volume supplied during times when equipment is running cooler (i.e. at part-load). According to PG&E, small reductions in air volume supplied result in a non-linear reduction in fan power; so a 20% drop in airflow volume results in a 45% to 50% reduction in fan energy consumption.

- **Investigate "free" cooling via economizers.** Airside economizers bring in cool outside air when weather conditions permit, thus reducing the burden on air conditioning units. While these systems are expensive to install (economizers can cost tens of thousands of dollars for the unit, and thousands more to install), they can significantly reduce long-term cooling costs for SMEs that have larger server rooms with several racks. Depending on outside conditions, an economizer can reduce data center cooling costs by more than 60%, according to a recent report from PG&E.

10. **Shop green.** Look for products with eco labels and those that include specific efficiency ratings. When writing RFPs for new equipment, include specific terms for heat output, cooling capacity, and energy consumption.
    - **Buy servers with energy-efficient power supplies.** Look to the EPA's Energy Star program to start offering ratings for server efficiency now that U.S. House and Senate bills encouraging EPA involvement have been passed. In the meantime, look for products containing 80 PLUS certified power supplies.
    - **Seek rebates from utilities.** In an effort to reduce new investments in electricity generation, a few US utility companies are creating incentives to try and counter the energy consumption of power-hungry data centers. Expect others to follow suit.
    - **Look for new innovative offerings.** With the recent push towards energy efficiency, many vendors are coming out with innovative new products.
        - **Cisco** started 2007 on a green foot. In addition to getting its own house in order (the company's internal green initiatives will save $20 million over a three-year period) and hosting Webinars on green data center design, Cisco also offers a range of products that claim to reduce power consumption and improve operating efficiency.
        - **Sun** has developed more efficient chips and servers and claims that its UltraSPARC processor is so efficient that it could reduce worldwide server demand by 50%.
        - **HP** and **IBM** offer software to adjust the power consumption of servers to match the load.

Bottom Line

Even in SMEs, the power and cooling of data center equipment can easily cost tens or hundreds of thousands of dollars a year. Employ best practices to boost data center energy efficiency and lower costs.